paloalto TECHDOCS

PAN-OS® Administrator's Guide

Version 10.1

docs.paloaltonetworks.com

https://t.me/learningnets

Contact Information

Corporate Headquarters:

Palo Alto Networks

3000 Tannery Way

Santa Clara, CA 95054

www.paloaltonetworks.com/company/contact-support.html

About the Documentation

- For the most recent version of this guide or for access to related documentation, visit the Technical Documentation portal docs.paloaltonetworks.com.
- To search for a specific topic, go to our search page docs.paloaltonetworks.com/search.html.
- Have feedback or questions for us? Leave a comment on any page in the portal, or write to us at documentation@paloaltonetworks.com.

Copyright

Palo Alto Networks, Inc.

www.paloaltonetworks.com

©2021–2022 Palo Alto Networks, Inc. Palo Alto Networks is a registered trademark of Palo Alto Networks. A list of our trademarks can be found at www.paloaltonetworks.com/company/trademarks.html. All other marks mentioned herein may be trademarks of their respective companies.

Last Revised

March 9, 2022

https://t.me/learningnets

Table of Contents

Getting Started 21

Integrate the Firewall into Your Management Network 22
- Determine Your Management Strategy 22
- Perform Initial Configuration 23
- Set Up Network Access for External Services 30

Register the Firewall 37
- Create a New Support Account and Register a Firewall 37
- Register a Firewall 39
- (Optional) Perform Day 1 Configuration 42
- Register the Firewall Line Cards 45

Segment Your Network Using Interfaces and Zones 46
- Network Segmentation for a Reduced Attack Surface 46
- Configure Interfaces and Zones 47

Set Up a Basic Security Policy 51

Assess Network Traffic 56

Enable Free WildFire Forwarding 58

Best Practices for Completing the Firewall Deployment 61

Subscriptions 63

Subscriptions You Can Use With the Firewall 64

Activate Subscription Licenses 68

What Happens When Licenses Expire? 70

Enhanced Application Logs for Palo Alto Networks Cloud Services 73

Firewall Administration 77

Management Interfaces 78

Use the Web Interface 79
- Launch the Web Interface 79
- Configure Banners, Message of the Day, and Logos 80
- Use the Administrator Login Activity Indicators to Detect Account Misuse 82
- Manage and Monitor Administrative Tasks 84
- Commit, Validate, and Preview Firewall Configuration Changes 85
- Export Configuration Table Data 87
- Use Global Find to Search the Firewall or Panorama Management Server 88
- Manage Locks for Restricting Configuration Changes 90

Manage Configuration Backups 92
- Save and Export Firewall Configurations 92

https://t.me/learningnets

Revert Firewall Configuration Changes....94
Manage Firewall Administrators....96
Administrative Role Types....96
Configure an Admin Role Profile....97
Administrative Authentication....104
Configure Administrative Accounts and Authentication....105
Configure Tracking of Administrator Activity....112
Reference: Web Interface Administrator Access....114
Web Interface Access Privileges....114
Panorama Web Interface Access Privileges....186
Reference: Port Number Usage....191
Ports Used for Management Functions....191
Ports Used for HA....192
Ports Used for Panorama....193
Ports Used for GlobalProtect....195
Ports Used for User-ID....195
Ports Used for IPSec....197
Ports Used for Routing....197
Ports Used for DHCP....198
Ports Used for Infrastructure....198
Reset the Firewall to Factory Default Settings....200
Bootstrap the Firewall....201
USB Flash Drive Support....201
Sample init-cfg.txt Files....202
Prepare a USB Flash Drive for Bootstrapping a Firewall....204
Bootstrap a Firewall Using a USB Flash Drive....206

Device Telemetry....209

Device Telemetry Overview....210
Device Telemetry Collection and Transmission Intervals....211
Manage Device Telemetry....212
Enable Device Telemetry....212
Disable Device Telemetry....212
Manage the Data the Device Telemetry Collects....213
Manage Historical Device Telemetry....213
Monitor Device Telemetry....215
Sample the Data that Device Telemetry Collects....216

Authentication....217

Authentication Types....218
External Authentication Services....218

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Multi-Factor Authentication 218
SAML 220
Kerberos 220
TACACS+ 221
RADIUS 222
LDAP 224
Local Authentication 224
Plan Your Authentication Deployment 225
Configure Multi-Factor Authentication 227
Configure MFA Between RSA SecurID and the Firewall 231
Configure MFA Between Okta and the Firewall 239
Configure MFA Between Duo and the Firewall 249
Configure SAML Authentication 258
Configure Kerberos Single Sign-On 263
Configure Kerberos Server Authentication 266
Configure TACACS+ Authentication 267
Configure RADIUS Authentication 270
Configure LDAP Authentication 274
Connection Timeouts for Authentication Servers 276
Guidelines for Setting Authentication Server Timeouts 276
Modify the PAN-OS Web Server Timeout 277
Modify the Authentication Portal Session Timeout 277
Configure Local Database Authentication 279
Configure an Authentication Profile and Sequence 280
Test Authentication Server Connectivity 284
Authentication Policy 286
Authentication Timestamps 286
Configure Authentication Policy 287
Troubleshoot Authentication Issues 291

Certificate Management 293

Keys and Certificates 294
Default Trusted Certificate Authorities (CAs) 297
Certificate Revocation 298
Certificate Revocation List (CRL) 298
Online Certificate Status Protocol (OCSP) 299
Certificate Deployment 300
Set Up Verification for Certificate Revocation Status 301
Configure an OCSP Responder 301
Configure Revocation Status Verification of Certificates 302

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Configure Revocation Status Verification of Certificates Used for SSL/TLS Decryption....302
Configure the Master Key....304
Master Key Encryption....307
Configure Master Key Encryption Level....308
Master Key Encryption on a Firewall HA Pair....309
Master Key Encryption Logs....309
Unique Master Key Encryptions for AES-256-GCM....310
Obtain Certificates....311
Create a Self-Signed Root CA Certificate....311
Generate a Certificate....312
Import a Certificate and Private Key....313
Obtain a Certificate from an External CA....314
Install a Device Certificate....315
Deploy Certificates Using SCEP....316
Export a Certificate and Private Key....320
Configure a Certificate Profile....321
Configure an SSL/TLS Service Profile....324
Configure an SSH Service Profile....326
Create an SSH Management Profile....326
Create an SSH HA Profile....335
Replace the Certificate for Inbound Management Traffic....345
Configure the Key Size for SSL Forward Proxy Server Certificates....346
Revoke and Renew Certificates....347
Revoke a Certificate....347
Renew a Certificate....347
Secure Keys with a Hardware Security Module....348
Set Up Connectivity with an HSM....348
Encrypt a Master Key Using an HSM....354
Store Private Keys on an HSM....355
Manage the HSM Deployment....356

High Availability....359

HA Overview....360
HA Concepts....361
HA Modes....361
HA Links and Backup Links....362
Device Priority and Preemption....368
Failover....368
LACP and LLDP Pre-Negotiation for Active/Passive HA....370
Floating IP Address and Virtual MAC Address....370

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

ARP Load-Sharing....372
Route-Based Redundancy....373
HA Timers....374
Session Owner....377
Session Setup....378
NAT in Active/Active HA Mode....380
ECMP in Active/Active HA Mode....380
Set Up Active/Passive HA....382
Prerequisites for Active/Passive HA....382
Configuration Guidelines for Active/Passive HA....383
Configure Active/Passive HA....386
Define HA Failover Conditions....391
Verify Failover....394
Set Up Active/Active HA....395
Prerequisites for Active/Active HA....395
Configure Active/Active HA....396
Determine Your Active/Active Use Case....402
HA Clustering Overview....418
HA Clustering Best Practices and Provisioning....421
Configure HA Clustering....423
Refresh HA1 SSH Keys and Configure Key Options....426
HA Firewall States....435
Reference: HA Synchronization....437
What Settings Don't Sync in Active/Passive HA?....437
What Settings Don't Sync in Active/Active HA?....439
Synchronization of System Runtime Information....443
Monitoring....447
Use the Dashboard....448
Use the Application Command Center....450
ACC—First Look....450
ACC Tabs....452
ACC Widgets....454
Widget Descriptions....456
ACC Filters....462
Interact with the ACC....463
Use Case: ACC—Path of Information Discovery....467
Use the App Scope Reports....474
Summary Report....474
Change Monitor Report....475
Threat Monitor Report....476

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Threat Map Report....477
Network Monitor Report....478
Traffic Map Report....479
Use the Automated Correlation Engine....481
Automated Correlation Engine Concepts....481
View the Correlated Objects....482
Interpret Correlated Events....483
Use the Compromised Hosts Widget in the ACC....485
Take Packet Captures....486
Types of Packet Captures....486
Disable Hardware Offload....487
Take a Custom Packet Capture....488
Take a Threat Packet Capture....492
Take an Application Packet Capture....494
Take a Packet Capture on the Management Interface....497
Monitor Applications and Threats....500
View and Manage Logs....501
Log Types and Severity Levels....501
View Logs....508
Filter Logs....509
Export Logs....510
Configure Log Storage Quotas and Expiration Periods....511
Schedule Log Exports to an SCP or FTP Server....511
Monitor Block List....513
View and Manage Reports....514
Report Types....514
View Reports....515
Configure the Expiration Period and Run Time for Reports....516
Disable Predefined Reports....516
Custom Reports....516
Generate Custom Reports....519
Generate Botnet Reports....522
Generate the SaaS Application Usage Report....524
Manage PDF Summary Reports....527
Generate User/Group Activity Reports....529
Manage Report Groups....531
Schedule Reports for Email Delivery....531
Manage Report Storage Capacity....532
View Policy Rule Usage....534
Use External Services for Monitoring....538
Configure Log Forwarding....539

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Configure Email Alerts....542
Use Syslog for Monitoring....544
Configure Syslog Monitoring....544
Syslog Field Descriptions....548
SNMP Monitoring and Traps....622
SNMP Support....622
Use an SNMP Manager to Explore MIBs and Objects....623
Enable SNMP Services for Firewall-Secured Network Elements....626
Monitor Statistics Using SNMP....626
Forward Traps to an SNMP Manager....628
Supported MIBs....630
Forward Logs to an HTTP/S Destination....639
NetFlow Monitoring....642
Configure NetFlow Exports....642
NetFlow Templates....644
Firewall Interface Identifiers in SNMP Managers and NetFlow Collectors....650
Monitor Transceivers....653

User-ID....655

User-ID Overview....656
User-ID Concepts....658
Group Mapping....658
User Mapping....658
Enable User-ID....663
Map Users to Groups....667
Map IP Addresses to Users....674
Create a Dedicated Service Account for the User-ID Agent....675
Configure User Mapping Using the Windows User-ID Agent....694
Configure User Mapping Using the PAN-OS Integrated User-ID Agent....708
Configure Server Monitoring Using WinRM....712
Configure User-ID to Monitor Syslog Senders for User Mapping....720
Map IP Addresses to Usernames Using Authentication Portal....730
Configure User Mapping for Terminal Server Users....736
Send User Mappings to User-ID Using the XML API....746
Enable User- and Group-Based Policy....747
Enable Policy for Users with Multiple Accounts....748
Verify the User-ID Configuration....750
Deploy User-ID in a Large-Scale Network....753
Deploy User-ID for Numerous Mapping Information Sources....753
Insert Username in HTTP Headers....757
Redistribute Data and Authentication Timestamps....759

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Share User-ID Mappings Across Virtual Systems....766

App-ID....769

App-ID Overview....770
Streamlined App-ID Policy Rules....771
Create an Application Filter Using Tags....771
Create an Application Filter Based on Custom Tags....772
App-ID and HTTP/2 Inspection....774
Manage Custom or Unknown Applications....776
Manage New and Modified App-IDs....777
Workflow to Best Incorporate New and Modified App-IDs....777
See the New and Modified App-IDs in a Content Release....778
See How New and Modified App-IDs Impact Your Security Policy....780
Ensure Critical New App-IDs are Allowed....780
Monitor New App-IDs....781
Disable and Enable App-IDs....783
Use Application Objects in Policy....784
Create an Application Group....784
Create an Application Filter....785
Create a Custom Application....786
Resolve Application Dependencies....790
Safely Enable Applications on Default Ports....792
Applications with Implicit Support....794
Security Policy Rule Optimization....798
Policy Optimizer Concepts....799
Migrate Port-Based to App-ID Based Security Policy Rules....806
Rule Cloning Migration Use Case: Web Browsing and SSL Traffic....813
Add Applications to an Existing Rule....817
Identify Security Policy Rules with Unused Applications....819
High Availability for Application Usage Statistics....822
How to Disable Policy Optimizer....822
App-ID Cloud Engine....824
Prepare to Deploy App-ID Cloud Engine....826
Enable or Disable the App-ID Cloud Engine....830
App-ID Cloud Engine Processing and Usage....830
New App Viewer (Policy Optimizer)....834
Add Apps to an Application Filter with Policy Optimizer....835
Add Apps to an Application Group with Policy Optimizer....838
Add Apps Directly to a Rule with Policy Optimizer....841
Replace an RMA Firewall (ACE)....844
Impact of License Expiration or Disabling ACE....845

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Commit Failure Due to Cloud Content Rollback....845
Troubleshoot App-ID Cloud Engine....846
SaaS App-ID Policy Recommendation....849
Import SaaS Policy Recommendation....850
Import Updated SaaS Policy Recommendation....852
Remove Deleted SaaS Policy Recommendation....853
Application Level Gateways....855
Disable the SIP Application-level Gateway (ALG)....857
Use HTTP Headers to Manage SaaS Application Access....858
Understand SaaS Custom Headers....858
Domains used by the Predefined SaaS Application Types....861
Create HTTP Header Insertion Entries using Predefined Types....862
Create Custom HTTP Header Insertion Entries....863
Maintain Custom Timeouts for Data Center Applications....865

Device-ID....867

Device-ID Overview....868
Prepare to Deploy Device-ID....871
Configure Device-ID....876
Manage Device-ID....879
CLI Commands for Device-ID....881

Threat Prevention....883

Best Practices for Securing Your Network from Layer 4 and Layer 7 Evasions....884
Set Up Antivirus, Anti-Spyware, and Vulnerability Protection....895
DNS Security....898
About DNS Security....898
Cloud-Delivered DNS Signatures and Protections....899
DNS Security Analytics....900
Enable DNS Security....903
DNS Security Data Collection and Logging....910
Use DNS Queries to Identify Infected Hosts on the Network....912
How DNS Sinkholing Works....912
Configure DNS Sinkholing....913
Configure DNS Sinkholing for a List of Custom Domains....914
Configure the Sinkhole IP Address to a Local Server on Your Network....916
See Infected Hosts that Attempted to Connect to a Malicious Domain....919
Data Filtering....923
Create a Data Filtering Profile....923
Predefined Data Filtering Patterns....926
WildFire Inline ML....929

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Configure WildFire Inline ML....929
Set Up File Blocking....933
Prevent Brute Force Attacks....936
Customize the Action and Trigger Conditions for a Brute Force Signature....937
Enable Evasion Signatures....941
Monitor Blocked IP Addresses....942
Threat Signature Categories....945
Create Threat Exceptions....953
Custom Signatures....955
Monitor and Get Threat Reports....956
Monitor Activity and Create Custom Reports Based on Threat Categories....956
Learn More About Threat Signatures....958
AutoFocus Threat Intelligence for Network Traffic....961
Share Threat Intelligence with Palo Alto Networks....968
Threat Prevention Resources....969

Decryption....971

Decryption Overview....972
Decryption Concepts....974
Keys and Certificates for Decryption Policies....974
SSL Forward Proxy....976
SSL Forward Proxy Decryption Profile....978
SSL Inbound Inspection....981
SSL Inbound Inspection Decryption Profile....982
SSL Protocol Settings Decryption Profile....983
SSH Proxy....985
SSH Proxy Decryption Profile....987
Profile for No Decryption....988
SSL Decryption for Elliptical Curve Cryptography (ECC) Certificates....989
Perfect Forward Secrecy (PFS) Support for SSL Decryption....989
SSL Decryption and Subject Alternative Names (SANs)....990
TLSv1.3 Decryption....991
High Availability Not Supported for Decrypted Sessions....993
Decryption Mirroring....994
Prepare to Deploy Decryption....995
Work with Stakeholders to Develop a Decryption Deployment Strategy....995
Develop a PKI Rollout Plan....997
Size the Decryption Firewall Deployment....999
Plan a Staged, Prioritized Deployment....1000
Define Traffic to Decrypt....1002

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Create a Decryption Profile....1003
Create a Decryption Policy Rule....1005
Configure SSL Forward Proxy....1009
Configure SSL Inbound Inspection....1015
Configure SSH Proxy....1018
Configure Server Certificate Verification for Undecrypted Traffic....1019
Decryption Exclusions....1020
Palo Alto Networks Predefined Decryption Exclusions....1021
Exclude a Server from Decryption for Technical Reasons....1022
Local Decryption Exclusion Cache....1023
Create a Policy-Based Decryption Exclusion....1025
Block Private Key Export....1029
Generate a Private Key and Block It....1029
Import a Private Key and Block It....1030
Import a Private Key for IKE Gateway and Block It....1031
Verify Private Key Blocking....1034
Enable Users to Opt Out of SSL Decryption....1036
Temporarily Disable SSL Decryption....1038
Configure Decryption Port Mirroring....1039
Verify Decryption....1042
Troubleshoot and Monitor Decryption....1046
Decryption Application Command Center Widgets....1047
Decryption Log....1051
Custom Report Templates for Decryption....1066
Unsupported Parameters by Proxy Type and TLS Version....1067
Decryption Troubleshooting Workflow Examples....1068
Activate Free Licenses for Decryption Features....1089

URL Filtering....1091

About Palo Alto Networks URL Filtering Solution....1092
How Advanced URL Filtering Works....1093
URL Filtering Inline ML....1095
URL Filtering Use Cases....1096
URL Categories....1099
Security-Focused URL Categories....1099
Malicious URL Categories....1100
Verified URL Categories....1102
Policy Actions You Can Take Based on URL Categories....1103
Plan Your URL Filtering Deployment....1106
URL Filtering Best Practices....1109
Activate The Advanced URL Filtering Subscription....1111

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Configure URL Filtering 1113
Test URL Filtering Configuration 1117
Verify URL Filtering 1117
Verify Advanced URL Filtering 1117
Configure URL Filtering Inline ML 1120
Monitor Web Activity 1123
Monitor Web Activity of Network Users 1123
View the User Activity Report 1125
Configure Custom URL Filtering Reports 1127
Log Only the Page a User Visits 1131
Create a Custom URL Category 1132
URL Category Exceptions 1134
Basic Guidelines For URL Category Exception Lists 1134
Wildcard Guidelines for URL Category Exception Lists 1134
URL Category Exception List—Wildcard Examples 1135
Use an External Dynamic List in a URL Filtering Profile 1137
Allow Password Access to Certain Sites 1139
Prevent Credential Phishing 1142
Methods to Check for Corporate Credential Submissions 1142
Configure Credential Detection with the Windows User-ID Agent 1144
Set Up Credential Phishing Prevention 1146
Safe Search Enforcement 1150
Safe Search Settings for Search Providers 1150
Block Search Results when Strict Safe Search is not Enabled 1153
Transparently Enable Safe Search for Users 1156
URL Filtering Response Pages 1162
Customize the URL Filtering Response Pages 1166
HTTP Header Logging 1168
Request to Change the Category for a URL 1169
Make a Change Request Online 1169
Make a Bulk Change Request 1170
Make a Change Request from the Firewall 1171
Troubleshoot URL Filtering 1173
Problems Activating Advanced URL Filtering 1173
PAN-DB Cloud Connectivity Issues 1173
URLs Classified as Not-Resolved 1175
Incorrect Categorization 1175
PAN-DB Private Cloud 1178
M-600 Appliance for PAN-DB Private Cloud 1178
Set Up the PAN-DB Private Cloud 1180
Enable SSL/TLS Handshake Inspection 1190

https://t.me/learningnets

Quality of Service 1195
QoS Overview 1196
QoS Concepts 1198
QoS for Applications and Users 1198
QoS Policy 1198
QoS Profile 1199
QoS Classes 1199
QoS Priority Queuing 1200
QoS Bandwidth Management 1200
QoS Egress Interface 1201
QoS for Clear Text and Tunneled Traffic 1202
Configure QoS 1203
Configure QoS for a Virtual System 1210
Enforce QoS Based on DSCP Classification 1217
QoS Use Cases 1220
Use Case: QoS for a Single User 1220
Use Case: QoS for Voice and Video Applications 1222
VPNs 1227
VPN Deployments 1228
Site-to-Site VPN Overview 1229
Site-to-Site VPN Concepts 1230
IKE Gateway 1230
Tunnel Interface 1230
Tunnel Monitoring 1231
Internet Key Exchange (IKE) for VPN 1231
IKEv2 1234
Set Up Site-to-Site VPN 1238
Set Up an IKE Gateway 1238
Define Cryptographic Profiles 1245
Set Up an IPSec Tunnel 1249
Set Up Tunnel Monitoring 1252
Enable/Disable, Refresh or Restart an IKE Gateway or IPSec Tunnel 1253
Test VPN Connectivity 1255
Interpret VPN Error Messages 1256
Site-to-Site VPN Quick Configs 1258
Site-to-Site VPN with Static Routing 1258
Site-to-Site VPN with OSPF 1262
Site-to-Site VPN with Static and Dynamic Routing 1268
Large Scale VPN (LSVPN) 1275

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

LSVPN Overview 1276
Create Interfaces and Zones for the LSVPN 1277
Enable SSL Between GlobalProtect LSVPN Components 1279
About Certificate Deployment 1279
Deploy Server Certificates to the GlobalProtect LSVPN Components 1279
Deploy Client Certificates to the GlobalProtect Satellites Using SCEP 1282
Configure the Portal to Authenticate Satellites 1285
Configure GlobalProtect Gateways for LSVPN 1287
Configure the GlobalProtect Portal for LSVPN 1291
GlobalProtect Portal for LSVPN Prerequisite Tasks 1291
Configure the Portal 1291
Define the Satellite Configurations 1292
Prepare the Satellite to Join the LSVPN 1296
Verify the LSVPN Configuration 1299
LSVPN Quick Configs 1300
Basic LSVPN Configuration with Static Routing 1300
Advanced LSVPN Configuration with Dynamic Routing 1303
Advanced LSVPN Configuration with iBGP 1305

Policy 1313

Policy Types 1314
Security Policy 1315
Components of a Security Policy Rule 1315
Security Policy Actions 1318
Create a Security Policy Rule 1319
Policy Objects 1323
Security Profiles 1325
Create a Security Profile Group 1332
Set Up or Override a Default Security Profile Group 1333
Track Rules Within a Rulebase 1336
Rule Numbers 1336
Rule UUIDs 1338
Enforce Policy Rule Description, Tag, and Audit Comment 1343
Move or Clone a Policy Rule or Object to a Different Virtual System 1346
Use an Address Object to Represent IP Addresses 1348
Address Objects 1348
Create an Address Object 1349
Use Tags to Group and Visually Distinguish Objects 1351
Create and Apply Tags 1351
Modify Tags 1352
View Rules by Tag Group 1353

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Use an External Dynamic List in Policy..... 1355
External Dynamic List..... 1355
Formatting Guidelines for an External Dynamic List..... 1358
Built-in External Dynamic Lists..... 1360
Configure the Firewall to Access an External Dynamic List..... 1361
Configure the Firewall to Access an External Dynamic List from the EDL Hosting Service..... 1364
Retrieve an External Dynamic List from the Web Server..... 1371
View External Dynamic List Entries..... 1371
Exclude Entries from an External Dynamic List..... 1372
Enforce Policy on an External Dynamic List..... 1373
Find External Dynamic Lists That Failed Authentication..... 1376
Disable Authentication for an External Dynamic List..... 1377
Register IP Addresses and Tags Dynamically..... 1379
Use Dynamic User Groups in Policy..... 1381
Use Auto-Tagging to Automate Security Actions..... 1384
Monitor Changes in the Virtual Environment..... 1388
Enable VM Monitoring to Track Changes on the Virtual Network..... 1388
Attributes Monitored on Virtual Machines in Cloud Platforms..... 1390
Use Dynamic Address Groups in Policy..... 1395
CLI Commands for Dynamic IP Addresses and Tags..... 1399
Enforce Policy on Endpoints and Users Behind an Upstream Device..... 1402
Use XFF Values for Policy Based on Source Users..... 1402
Use XFF IP Address Values in Security Policy and Logging..... 1403
Use the IP Address in the XFF Header to Troubleshoot Events..... 1406
Policy-Based Forwarding..... 1408
PBF..... 1408
Create a Policy-Based Forwarding Rule..... 1410
Use Case: PBF for Outbound Access with Dual ISPs..... 1413
Test Policy Rules..... 1423

Virtual Systems..... 1425

Virtual Systems Overview..... 1426
Virtual System Components and Segmentation..... 1426
Benefits of Virtual Systems..... 1427
Use Cases for Virtual Systems..... 1427
Platform Support and Licensing for Virtual Systems..... 1428
Administrative Roles for Virtual Systems..... 1428
Shared Objects for Virtual Systems..... 1429
Communication Between Virtual Systems..... 1430
Inter-VSYS Traffic That Must Leave the Firewall..... 1430

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Inter-VSYS Traffic That Remains Within the Firewall....1431
Inter-VSYS Communication Uses Two Sessions....1433
Shared Gateway....1434
External Zones and Shared Gateway....1434
Networking Considerations for a Shared Gateway....1435
Configure Virtual Systems....1436
Configure Inter-Virtual System Communication within the Firewall....1442
Configure a Shared Gateway....1443
Customize Service Routes for a Virtual System....1444
Customize Service Routes to Services for Virtual Systems....1444
Configure a PA-7000 Series Firewall for Logging Per Virtual System....1446
Configure Administrative Access Per Virtual System or Firewall....1448
Virtual System Functionality with Other Features....1450

Zone Protection and DoS Protection....1451

Network Segmentation Using Zones....1452
How Do Zones Protect the Network?....1453
Zone Defense....1454
Zone Defense Tools....1454
How Do the Zone Defense Tools Work?....1456
Firewall Placement for DoS Protection....1457
Baseline CPS Measurements for Setting Flood Thresholds....1457
Zone Protection Profiles....1459
Packet Buffer Protection....1463
DoS Protection Profiles and Policy Rules....1465
Configure Zone Protection to Increase Network Security....1472
Configure Reconnaissance Protection....1472
Configure Packet Based Attack Protection....1473
Configure Protocol Protection....1474
Configure Packet Buffer Protection....1478
Configure Packet Buffer Protection Based on Latency....1479
Configure Ethernet SGT Protection....1480
DoS Protection Against Flooding of New Sessions....1482
Multiple-Session DoS Attack....1482
Single-Session DoS Attack....1486
Configure DoS Protection Against Flooding of New Sessions....1486
End a Single Session DoS Attack....1489
Identify Sessions That Use Too Much of the On-Chip Packet Descriptor... 1490
Discard a Session Without a Commit....1493

Certifications....1495

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Enable FIPS and Common Criteria Support....1496
Access the Maintenance Recovery Tool (MRT)....1496
Change the Operational Mode to FIPS-CC Mode....1498
FIPS-CC Security Functions....1501
Scrub the Swap Memory on Firewalls or Appliances Running in FIPS-CC Mode..1503

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Getting Started

The following topics provide detailed steps to help you deploy a new Palo Alto Networks next-generation firewall. They provide details for integrating a new firewall into your network and how to set up a basic security policy. For guidance on continuing to deploy the security platform features to address your network security needs, review the Best Practices for Completing the Firewall Deployment.

> Integrate the Firewall into Your Management Network
> Register the Firewall
> Segment Your Network Using Interfaces and Zones
> Set Up a Basic Security Policy
> Assess Network Traffic
> Enable Free WildFire Forwarding
> Best Practices for Completing the Firewall Deployment

https://t.me/learningnets

Integrate the Firewall into Your Management Network

All Palo Alto Networks firewalls provide an out-of-band management port (MGT) that you can use to perform the firewall administration functions. By using the MGT port, you separate the management functions of the firewall from the data processing functions, safeguarding access to the firewall and enhancing performance. When using the web interface, you must perform all initial configuration tasks from the MGT port even if you plan to use an in-band data port for managing your firewall going forward.

Some management tasks, such as retrieving licenses and updating the threat and application signatures on the firewall require access to the Internet. If you do not want to enable external access to your MGT port, you will need to either set up an in-band data port to provide access to required external services (using service routes) or plan to manually upload updates regularly.

Do not enable access to your management interface from the internet or from other untrusted zones inside your enterprise security boundary. This applies whether you use the dedicated management port (MGT) or you configured a data port as your management interface. When integrating your firewall into your management network, follow the Adminstrative Access Best Practices to ensure that you are securing administrative access to your firewalls and other security devices in a way that prevents successful attacks.

The following topics describe how to perform the initial configuration steps that are necessary to integrate a new firewall into the management network and deploy it in a basic security configuration.

- Determine Your Management Strategy
- Perform Initial Configuration
- Set Up Network Access for External Services

The following topics describe how to integrate a single Palo Alto Networks next-generation firewall into your network. However, for redundancy, consider deploying a pair of firewalls in a High Availability configuration.

Determine Your Management Strategy

The Palo Alto Networks firewall can be configured and managed locally or it can be managed centrally using Panorama, the Palo Alto Networks centralized security management system. If you have six or more firewalls deployed in your network, use Panorama to achieve the following benefits:

- Reduce the complexity and administrative overhead in managing configuration, policies, software and dynamic content updates. Using device groups and templates on Panorama, you can effectively manage firewall-specific configuration locally on a firewall and enforce shared policies across all firewalls or device groups.
- Aggregate data from all managed firewalls and gain visibility across all the traffic on your network. The Application Command Center (ACC) on Panorama provides a single glass pane for unified reporting across all the firewalls, allowing you to centrally analyze, investigate and report on network traffic, security incidents and administrative modifications.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

The procedures that follow describe how to manage the firewall using the local web interface. If you want to use Panorama for centralized management, first Perform Initial Configuration and verify that the firewall can establish a connection to Panorama. From that point on you can use Panorama to configure your firewall centrally.

Perform Initial Configuration

By default, the PA-Series firewall has an IP address of 192.168.1.1 and a username/password of admin/admin. For security reasons, you must change these settings before continuing with other firewall configuration tasks. You must perform these initial configuration tasks either from the MGT interface, even if you do not plan to use this interface for your firewall management, or using a direct serial connection to the console port on the firewall.

STEP 1 | Install your firewall and connect power to it.



If your firewall model has dual power supplies, connect the second power supply for redundancy. Refer to the hardware reference guide for your model for details.

STEP 2 | Gather the required information from your network administrator.

- IP address for MGT port
- Netmask
- Default gateway
- DNS server address

STEP 3 | Connect your computer to the firewall.

You can connect to the firewall in one of the following ways:

- Connect a serial cable from your computer to the Console port and connect to the firewall using terminal emulation software (9600-8-N-1). Wait a few minutes for the boot-up sequence to complete; when the firewall is ready, the prompt changes to the name of the firewall, for example `PA-220 login`.
- Connect an RJ-45 Ethernet cable from your computer to the MGT port on the firewall. From a browser, go to **`https://192.168.1.1`**.



You may need to change the IP address on your computer to an address in the 192.168.1.0/24 network, such as 192.168.1.2, to access this URL.

STEP 4 | When prompted, log in to the firewall.

You must log in using the default username and password (admin/admin). The firewall will begin to initialize.

 ©2022 Palo Alto Networks, Inc.

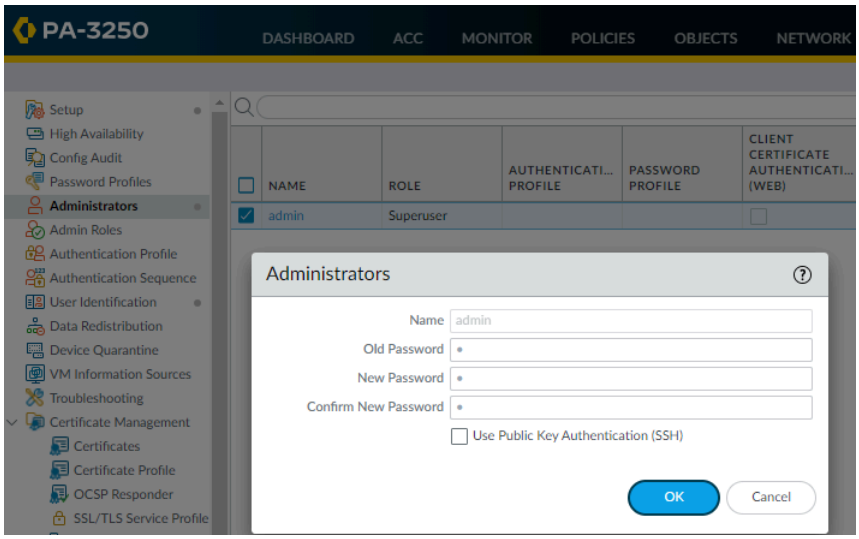
https://t.me/learningnets

STEP 5 | Set a secure password for the admin account.

Starting with PAN-OS 9.0.4, the predefined, default administrator password (admin/admin) must be changed on the first login on a device. The new password must be a minimum of eight characters and include a minimum of one lowercase and one uppercase character, as well as one number or special character.

Be sure to use the best practices for password strength to ensure a strict password and review the password complexity settings.

1. Select **Device** > **Administrators**.
2. Select the **admin** role.
3. Enter the current default password and the new password.



4. Click **OK** to save your settings.

 ©2022 Palo Alto Networks, Inc.

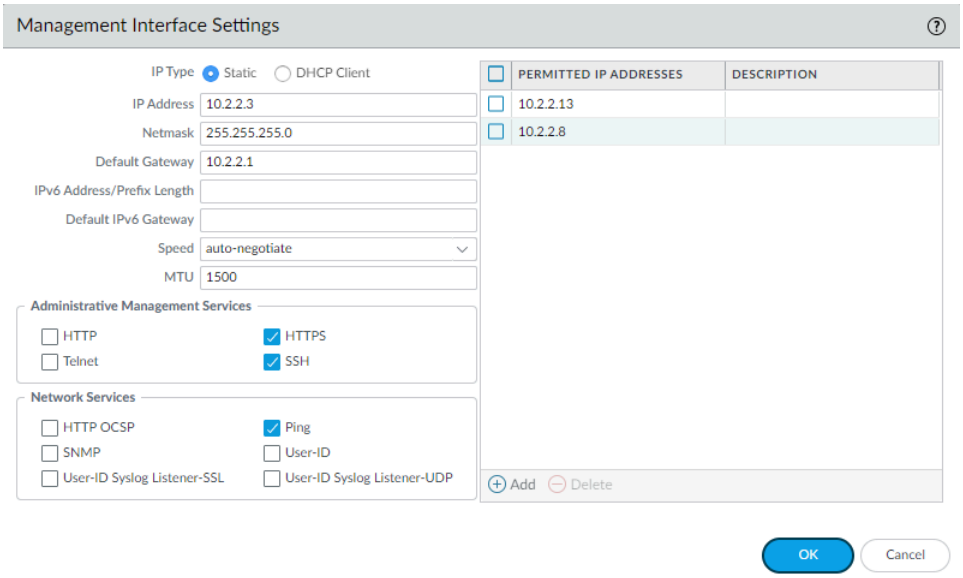
https://t.me/learningnets

STEP 6 | Configure the MGT interface.

1. Select **Device** > **Setup** > **Interfaces** and edit the **Management** interface.
2. Configure the address settings for the MGT interface using one of the following methods:
   - To configure static IP address settings for the MGT interface, set the **IP Type** to **Static** and enter the **IP Address**, **Netmask**, and **Default Gateway**.
   - To dynamically configure the MGT interface address settings, set the **IP Type** to **DHCP Client**. To use this method, you must Configure the Management Interface as a DHCP Client.

   *To prevent unauthorized access to the management interface, it is a an administrative best practice to **Add** the **Permitted IP Addresses** from which an administrator can access the MGT interface.*
3. Set the **Speed** to **auto-negotiate**.
4. Select which management services to allow on the interface.

   *Make sure **Telnet** and **HTTP** are not selected because these services use plaintext and are not as secure as the other services and could compromise administrator credentials.*



5. Click **OK**.

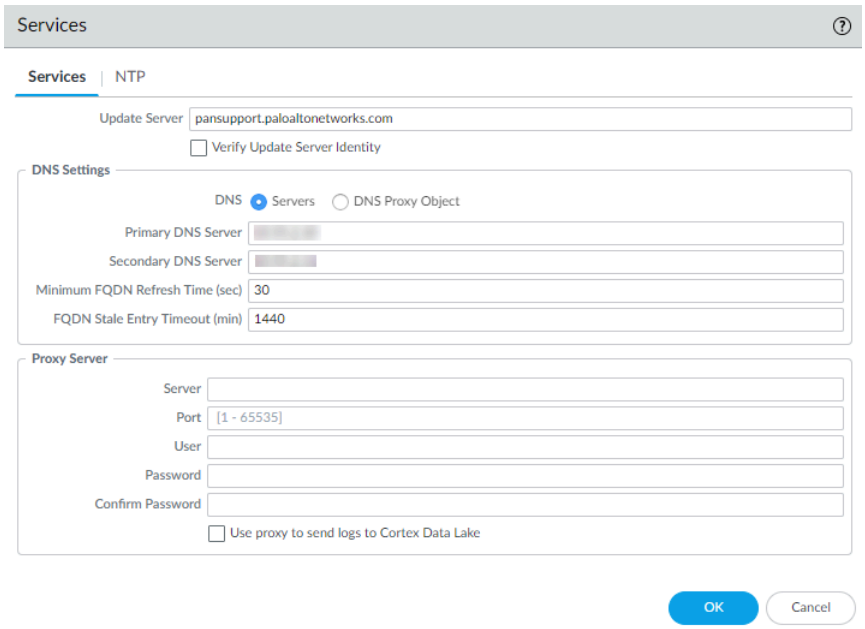
 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 7 | Configure DNS, update server, and proxy server settings.

You must manually configure at least one DNS server on the firewall or it will not be able to resolve hostnames; it will not use DNS server settings from another source, such as an ISP.

1. Select **Device** > **Setup** > **Services**.
   - For multi-virtual system platforms, select **Global** and edit the Services section.
   - For single virtual system platforms, edit the Services section.
2. On the **Services** tab, for **DNS**, select one of the following:
   - **Servers**—Enter the **Primary DNS Server** address and **Secondary DNS Server** address.
   - **DNS Proxy Object**—From the drop-down, select the **DNS Proxy** that you want to use to configure global DNS services, or click **DNS Proxy** to configure a new DNS proxy object.



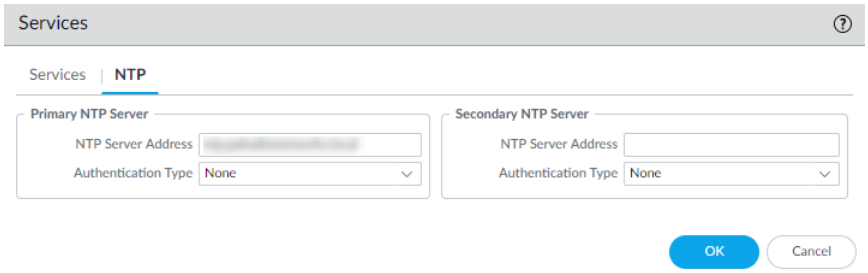
3. Click **OK**.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 8 | Configure date and time (NTP) settings.

1. Select **Device** > **Setup** > **Services**.
   - For multi-virtual system platforms, select **Global** and edit the Services section.
   - For single virtual system platforms, edit the Services section.
2. On the **NTP** tab, to use the virtual cluster of time servers on the Internet, enter the hostname `pool.ntp.org` as the **Primary NTP Server** or enter the IP address of your primary NTP server.



3. (Optional) Enter a **Secondary NTP Server** address.
4. (Optional) To authenticate time updates from the NTP server(s), for **Authentication Type**, select one of the following for each server:
   - **None**—(Default) Disables NTP authentication.
   - **Symmetric Key**—Firewall uses symmetric key exchange (shared secrets) to authenticate time updates.
     - **Key ID**—Enter the Key ID (1-65534).
     - **Algorithm**—Select the algorithm to use in NTP authentication (**MD5** or **SHA1**).
   - **Autokey**—Firewall uses autokey (public key cryptography) to authenticate time updates.
5. Click **OK**.

STEP 9 | (Optional) Configure general firewall settings as needed.

1. Select **Device** > **Setup** > **Management** and edit the General Settings.
2. Enter a **Hostname** for the firewall and enter your network **Domain** name. The domain name is just a label; it will not be used to join the domain.
3. Enter **Login Banner** text that informs users who are about to log in that they require authorization to access the firewall management functions.

   As a best practice, avoid using welcoming verbiage. Additionally, you should ask your legal department to review the banner message to ensure it adequately warns that unauthorized access is prohibited.
4. Enter the **Latitude** and **Longitude** to enable accurate placement of the firewall on the world map.
5. Click **OK**.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 10 | Commit your changes.



When the configuration changes are saved, you lose connectivity to the web interface because the IP address has changed.

Click **Commit** at the top right of the web interface. The firewall can take up to 90 seconds to save your changes.

STEP 11 | Connect the firewall to your network.

1. Disconnect the firewall from your computer.
2. Connect the MGT port to a switch port on your management network using an RJ-45 Ethernet cable. Make sure that the switch port you cable the firewall to is configured for auto-negotiation.

STEP 12 | Open an SSH management session to the firewall.

Using a terminal emulation software, such as PuTTY, launch an SSH session to the firewall using the new IP address you assigned to it.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

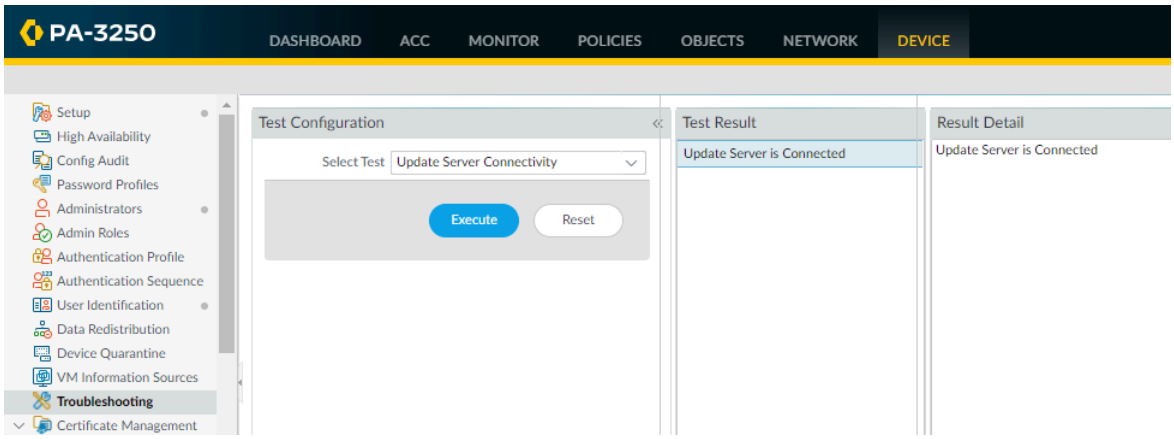
STEP 13 | Verify network access to external services required for firewall management, such as the Palo Alto Networks Update Server.

You can do this in one of the following ways:

- If you do not want to allow external network access to the MGT interface, you will need to set up a data port to retrieve required service updates. Continue to Set Up Network Access for External Services.
- If you do plan to allow external network access to the MGT interface, verify that you have connectivity and then proceed to Register the Firewall and Activate Subscription Licenses.

1. Use update server connectivity test to verify network connectivity to the Palo Alto Networks Update server as shown in the following example:

   1. Select **Device** > **Troubleshooting**, and select **Update Server Connectivity** from the Select Test drop-down.
   2. **Execute** the update server connectivity test.



2. Use the following CLI command to retrieve information on the support entitlement for the firewall from the Palo Alto Networks update server:

```
request support
check
```

If you have connectivity, the update server will respond with the support status for your firewall. If your firewall is not yet registered, the update server returns the following message:

```
Contact Us

https://www.paloaltonetworks.com/company/contact-us.html

Support Home

https://www.paloaltonetworks.com/support/tabs/overview.html

Device not found on this update server
```

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Set Up Network Access for External Services

By default, the firewall uses the MGT interface to access remote services, such as DNS servers, content updates, and license retrieval. If you do not want to enable external network access to your management network, you must set up an in-band data port to provide access to required external services and set up service routes to instruct the firewall what port to use to access the external services.

Do not enable management access from the internet or from other untrusted zones inside your enterprise security boundary. Follow the Adminstrative Access Best Practices to ensure that you are properly securing your firewall.

This task requires familiarity with firewall interfaces, zones, and policies. For more information on these topics, see Configure Interfaces and Zones and Set Up a Basic Security Policy.

STEP 1 | Decide which interface you want to use for access to external services and connect it to your switch or router port.

The interface you use must have a static IP address.

STEP 2 | Log in to the web interface.

Using a secure connection (https) from your web browser, log in using the new IP address and password you assigned during initial configuration (https://<IP address>). You will see a certificate warning; that is okay. Continue to the web page.

STEP 3 | (Optional) The firewall comes preconfigured with a default virtual wire interface between ports Ethernet 1/1 and Ethernet 1/2 (and a corresponding default security policy and zones). If you do not plan to use this virtual wire configuration, you must manually delete the configuration to prevent it from interfering with other interface settings you define.

You must delete the configuration in the following order:

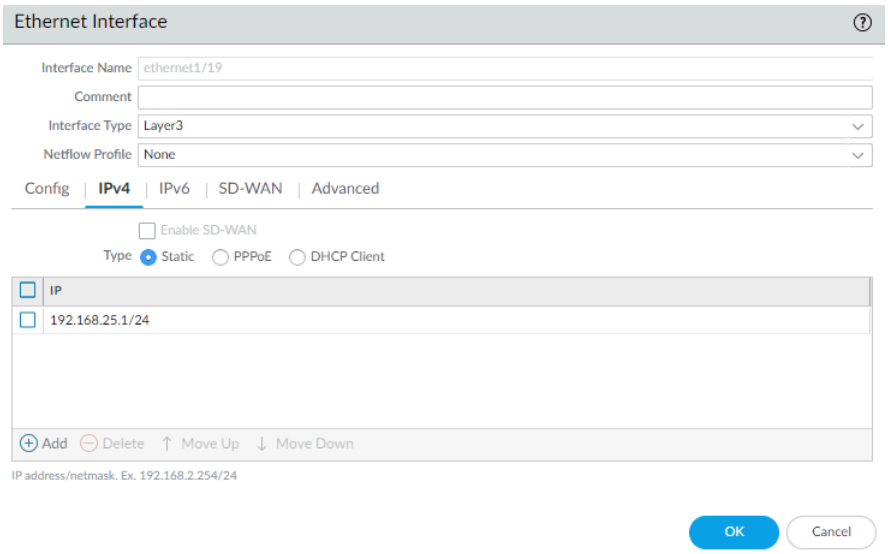
1. To delete the default security policy, select **Policies** > **Security**, select the rule, and click **Delete**.
2. To delete the default virtual wire, select **Network** > **Virtual Wires**, select the virtual wire and click **Delete**.
3. To delete the default trust and untrust zones, select **Network** > **Zones**, select each zone and click **Delete**.
4. To delete the interface configurations, select **Network** > **Interfaces** and then select each interface (ethernet1/1 and ethernet1/2) and click **Delete**.
5. **Commit** the changes.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 4 | Configure the interface you plan to use for external access to management services.

1. Select **Network** > **Interfaces** and select the interface that corresponds to the interface you cabled in Step 1.
2. Select the **Interface Type**. Although your choice here depends on your network topology, this example shows the steps for **Layer3**.
3. On the **Config** tab, expand the **Security Zone** drop-down and select **New Zone**.
4. In the Zone dialog, enter a **Name** for new zone, for example Management, and then click **OK**.
5. Select the **IPv4** tab, select the **Static** radio button, and click **Add** in the IP section, and enter the IP address and network mask to assign to the interface, for example 192.168.1.254/24. You must use a static IP address on this interface.

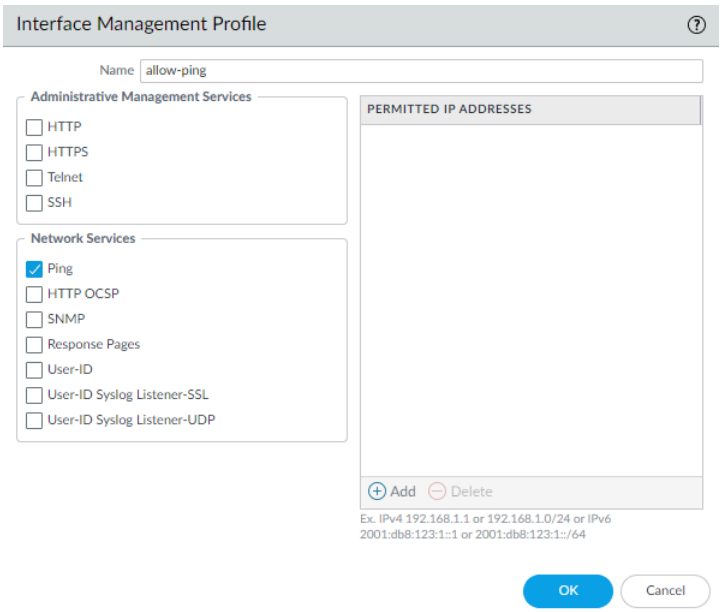


6. Select **Advanced** > **Other Info**, expand the **Management Profile** drop-down, and select **New Management Profile**.
7. Enter a **Name** for the profile, such as allow_ping, and then select the services you want to allow on the interface. For the purposes of allowing access to the external services, you probably only need to enable **Ping** and then click **OK**.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

*These services provide management access to the firewall, so only select the services that correspond to the management activities you want to allow on this interface. For example, don't enable HTTP or Telnet because those protocols transmit in plaintext and therefore aren't secure. Or if you plan to use the MGT interface for firewall configuration tasks through the web interface or CLI, you don't enable HTTP, HTTPS, SSH, or Telnet so that you prevent unauthorized access through the interface (if you must allow HTTPS or SSH in this scenario, limit access to a specific set of **Permitted IP Addresses**). For details, see* Use Interface Management Profiles to Restrict Access.



8. To save the interface configuration, click **OK**.

©2022 Palo Alto Networks, Inc.

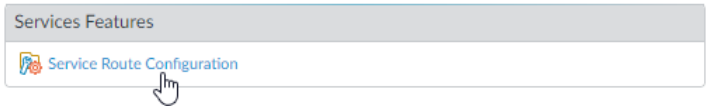
https://t.me/learningnets

STEP 5 | Configure the Service Routes.

By default, the firewall uses the MGT interface to access the external services it requires. To change the interface the firewall uses to send requests to external services, you must edit the service routes.

This example shows how to set up global service routes. For information on setting up network access to external services on a virtual system basis rather than a global basis, see Customize Service Routes to Services for Virtual Systems.

1. Select **Device** > **Setup** > **Services** > **Global** and click **Service Route Configuration**.



*For the purposes of activating your licenses and getting the most recent content and software updates, you will want to change the service route for **DNS**, **Palo Alto Networks Services**, **URL Updates**, and **AutoFocus**.*

2. Click the **Customize** radio button, and select one of the following:

   - For a predefined service, select **IPv4** or **IPv6** and click the link for the service. To limit the drop-down list for Source Address, select **Source Interface** and select the interface you just configured. Then select a Source Address (from that interface) as the service route.

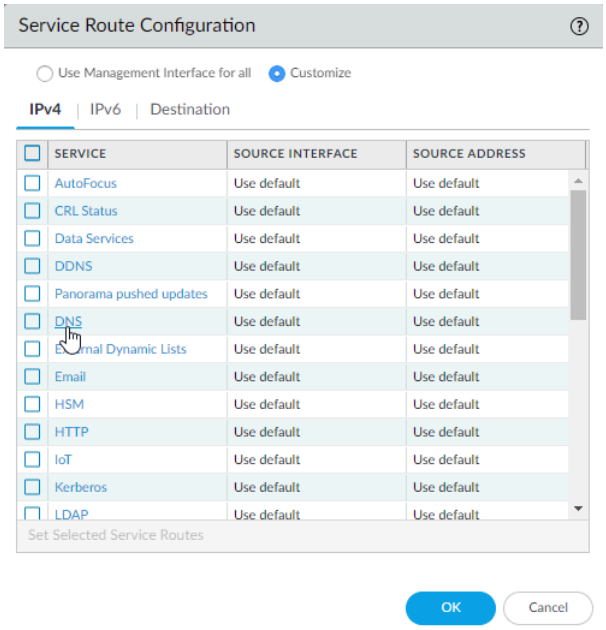
     If more than one IP address is configured for the selected interface, the **Source Address** drop-down allows you to select an IP address.

   - To create a service route for a custom destination, select **Destination**, and click **Add**. Enter a **Destination** IP address. An incoming packet with a destination address that matches this address will use as its source the Source Address you specify for this service route. To limit the drop-down for Source Address, select a **Source Interface**. If

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

more than one IP address is configured for the selected interface, the **Source Address** drop-down allows you to select an IP address.

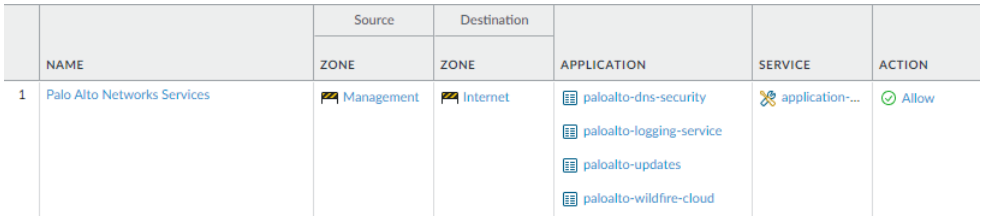


3. Click **OK** to save the settings.
4. Repeat Steps 5.2 - 5.3 above for each service route you want to modify.
5. **Commit** your changes.

STEP 6 | Configure an external-facing interface and an associated zone and then create a security policy rule to allow the firewall to send service requests from the internal zone to the external zone.

1. Select **Network** > **Interfaces** and then select the external-facing interface. Select **Layer3** as the **Interface Type**, **Add** the **IP** address (on the **IPv4** or **IPv6** tab), and create the associated **Security Zone** (on the **Config** tab), such as Internet. This interface must have a static IP address; you do not need to set up management services on this interface.
2. To set up a security rule that allows traffic from your internal network to the Palo Alto Networks update server, select **Policies** > **Security** and click **Add**.

*As a best practice when creating Security policy rules, use application-based rules instead of port-based rules to ensure that you are accurately identifying the underlying application regardless of the port, protocol, evasive tactics, or encryption in use. Always leave the **Service** set to **application-default**. In this case, create a security policy rule that allows access to the update server (and other Palo Alto Networks services).*



 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 7 | Create a NAT policy rule.

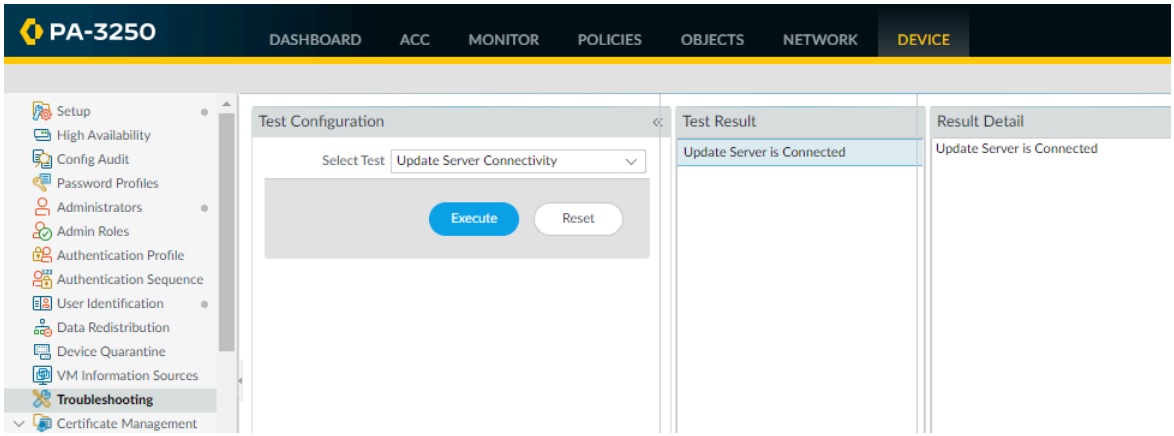
1. If you are using a private IP address on the internal-facing interface, you will need to create a source NAT rule to translate the address to a publicly routable address. Select **Policies** > **NAT** and then click **Add**. At a minimum you must define a name for the rule (**General** tab), specify a source and destination zone, Management to Internet in this case (**Original Packet** tab), and define the source address translation settings (**Translated Packet** tab) and then click **OK**.
2. **Commit** your changes.

		Original Packet			Translated Packet	
	NAME	SOURCE ZONE	DESTINATION ZONE	SERVICE	SOURCE TRANSLATION	DESTINATION TRANSLATION
1	Source NAT	Management	Internet	any	dynamic-ip-and-port	none

STEP 8 | Select **Device** > **Troubleshooting** and verify that you have connectivity from the data port to the external services, including the default gateway, using the **Ping** connectivity test, and the Palo Alto Networks Update Server using the **Update Server Connectivity** test. In this example, the firewall connectivity to the Palo Alto Networks Update Server is tested.

After you verify you have the required network connectivity, continue to Register the Firewall and Activate Subscription Licenses.

1. Select **Update Server** from the Select Test drop-down.
2. **Execute** the Palo Alto Networks Update Server connectivity test.



3. Access the firewall CLI, and use the following command to retrieve information on the support entitlement for the firewall from the Palo Alto Networks update server:

```
request support
check
```

If you have connectivity, the update server will respond with the support status for your firewall. Because your firewall is not registered, the update server will return the following message:

```
Contact Us
https://www.paloaltonetworks.com/company/contact-us.html
Support Home
https://www.paloaltonetworks.com/support/tabs/overview.html
```

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

```
Device not found on this update server
```

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Register the Firewall

Before you can activate support and other licenses and subscriptions, you must first register the firewall. Before you can register a firewall, though, you must first have an active support account. Perform one of the following tasks depending on whether you have an active support account:

- If you don't have an active support account, then Create a New Support Account and Register a Firewall.
- If you already have an active support account, then you are ready to Register a Firewall.
- (Optional) Perform Day 1 Configuration on a registered firewall.
- If your firewall uses line cards such as an NPC (Network Processing Card), then Register the Firewall Line Cards.

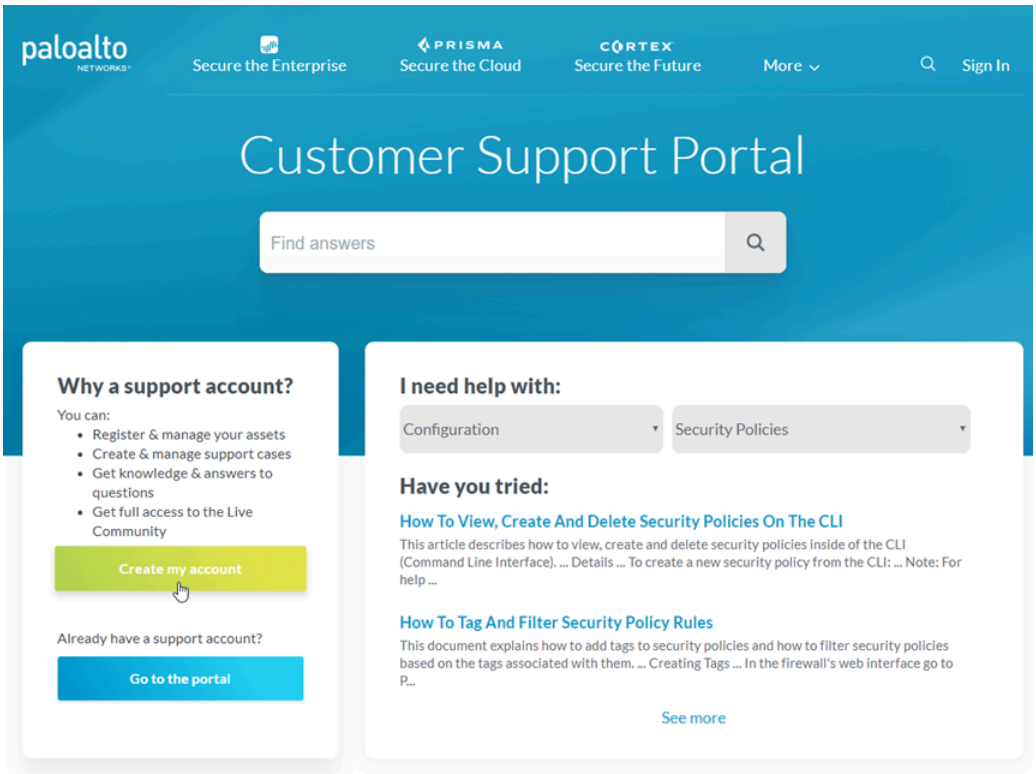
If you are registering a VM-Series firewall, refer to the VM-Series Deployment Guide for instructions.

Create a New Support Account and Register a Firewall

If you do not already have an active Palo Alto Networks support account, then you need to register your firewall when you create your new support account.

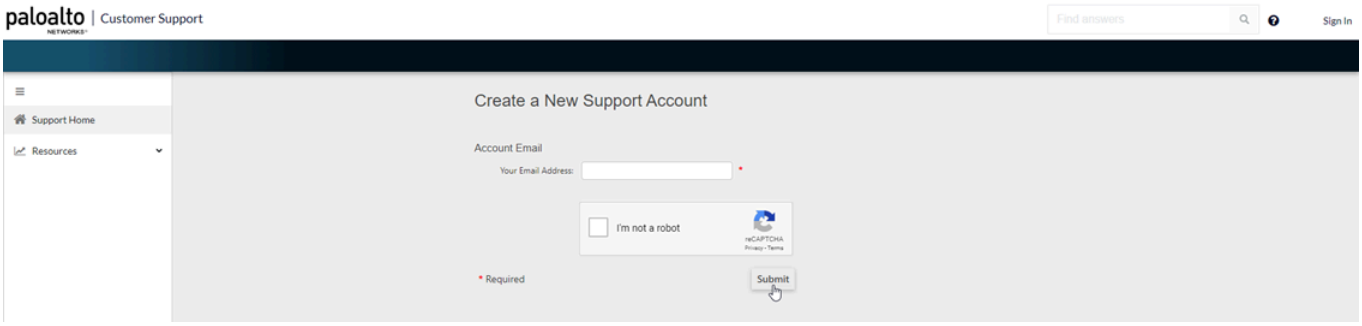
STEP 1 | Go to the Palo Alto Networks Customer Support Portal.

STEP 2 | Click **Create my account**.

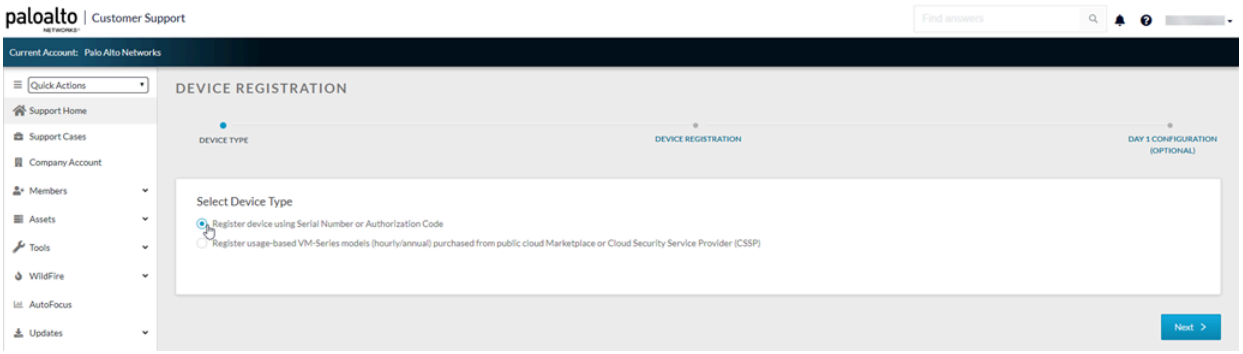


https://t.me/learningnets

STEP 3 | Enter **Your Email Address**, check **I'm not a robot**, and click **Submit**.



STEP 4 | Select **Register device using Serial Number or Authorization Code** and click **Next**.

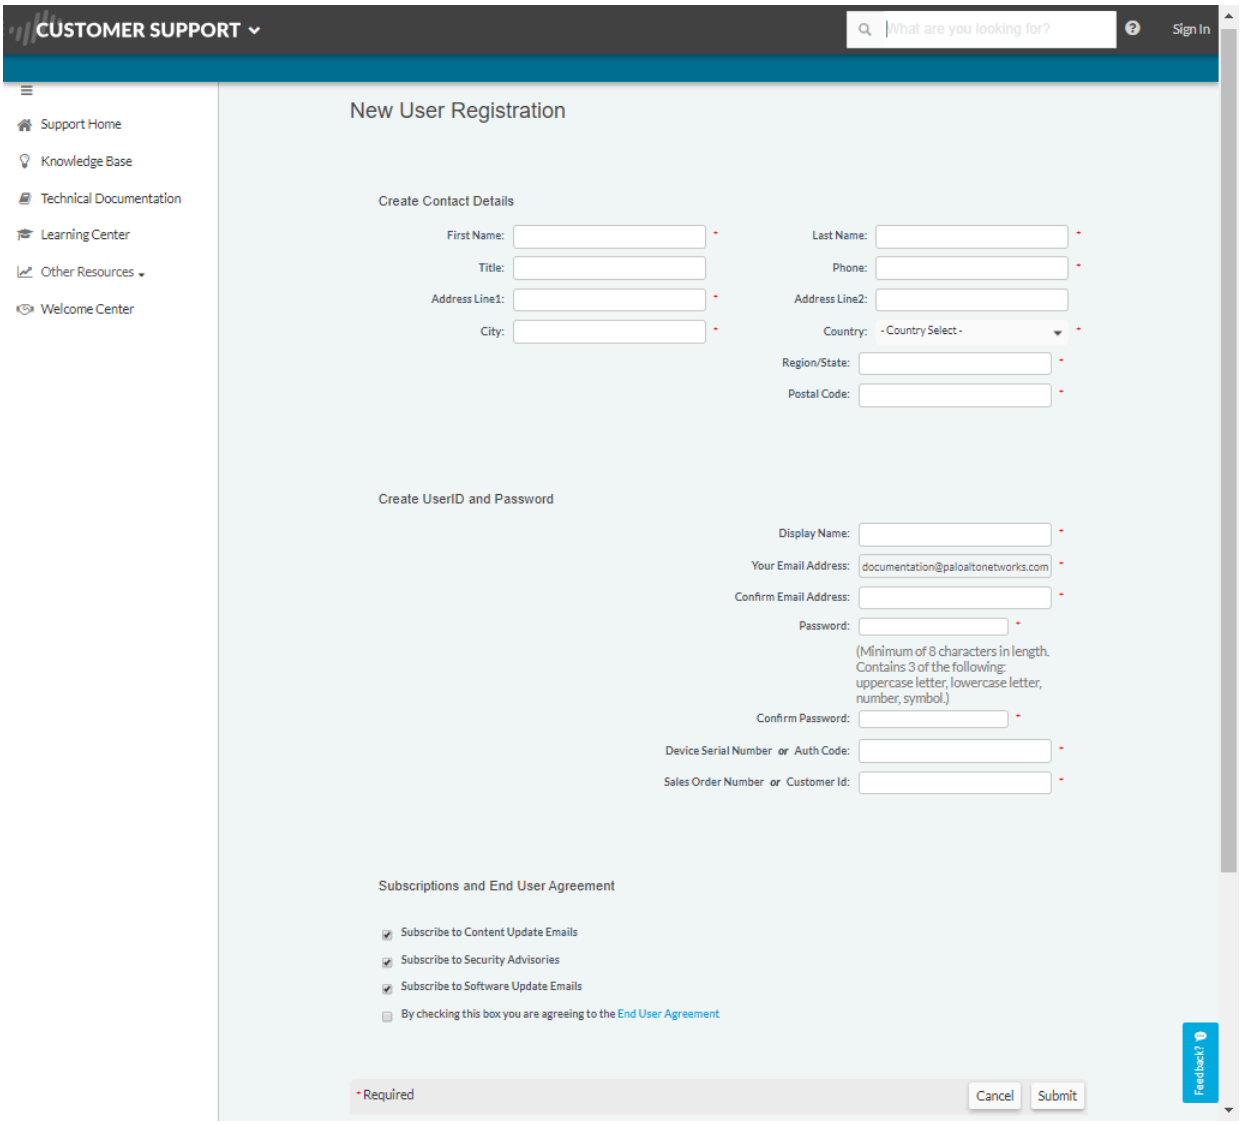


 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 5 | Complete the registration form.

1. Enter the contact details for the person in your organization who will own this account. Required fields are indicated by red asterisks.
2. Create a UserID and Password for the account. Required fields are indicated by red asterisks.
3. Enter the **Device Serial Number** or **Auth Code**.
4. Enter your **Sales Order Number** or **Customer Id**.
5. To ensure that you are always alerted to the latest updates and security advisories, **Subscribe to Content Update Emails**, **Subscribe to Security Advisories**, and **Subscribe to Software Update Emails**.
6. Select the check box to agree to the End User Agreement and **Submit**.



Register a Firewall

If you already have an active Palo Alto Networks Customer Support account, perform the following task to register your firewall.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

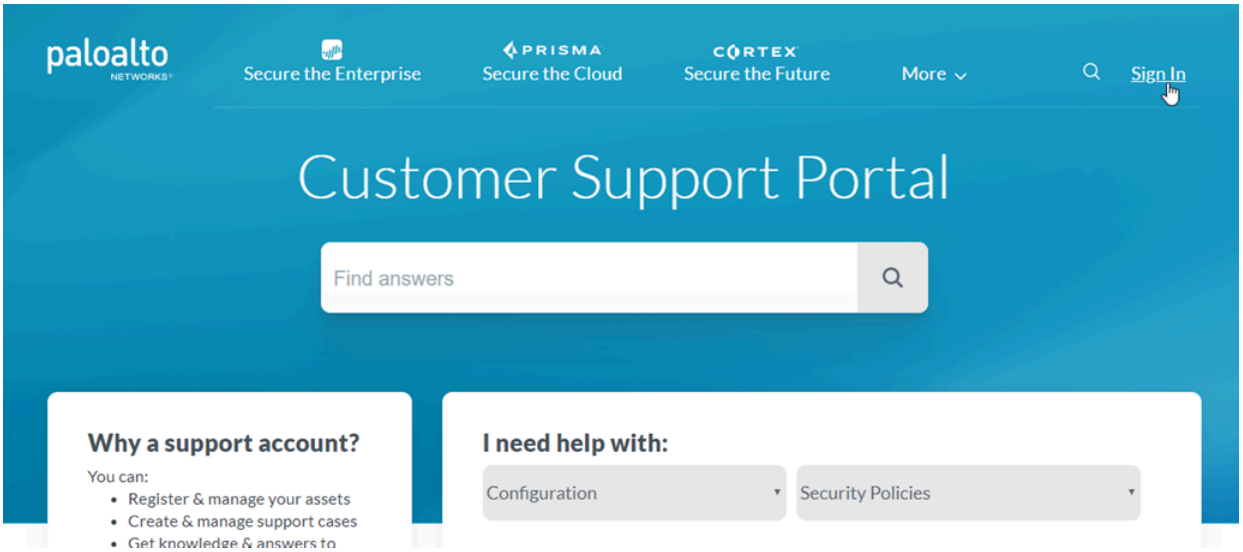
STEP 1 | Log in to the firewall web interface.

Using a secure connection (HTTPS) from your web browser, log in using the new IP address and password you assigned during initial configuration (https://<IP address>).

STEP 2 | Locate your serial number and copy it to the clipboard.

On the **Dashboard**, locate your **Serial Number** in the General Information section of the screen.

STEP 3 | Go to the Palo Alto Networks Customer Support Portal and, if not already logged in, **Sign In** now.

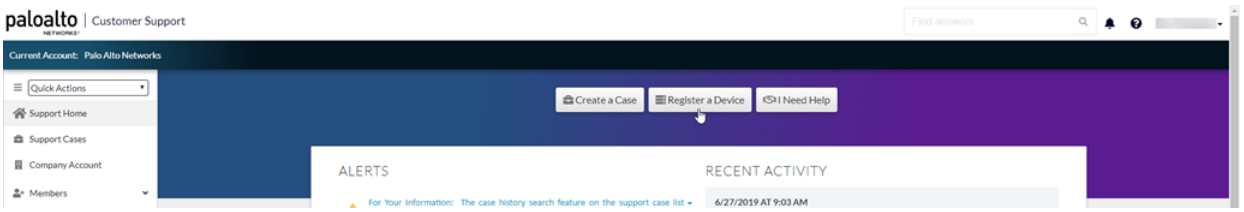


 ©2022 Palo Alto Networks, Inc.

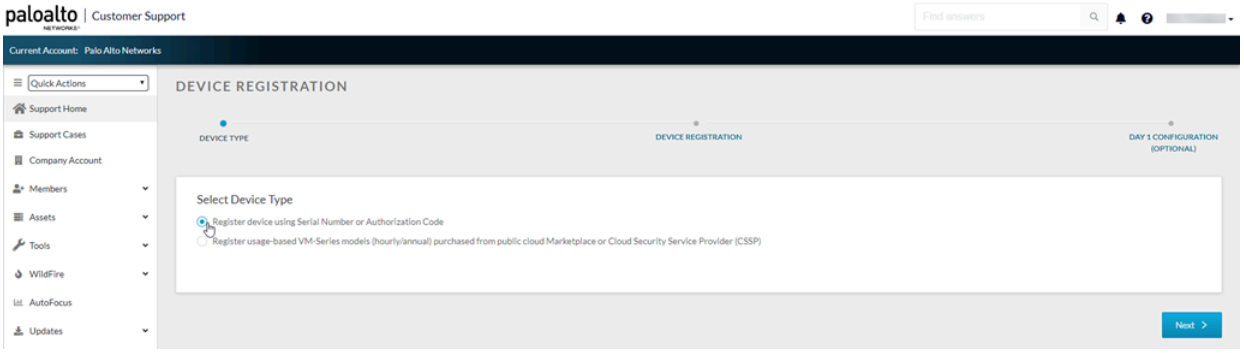
https://t.me/learningnets

STEP 4 | Register the firewall.

1. On the Support Home page, click **Register a Device**.



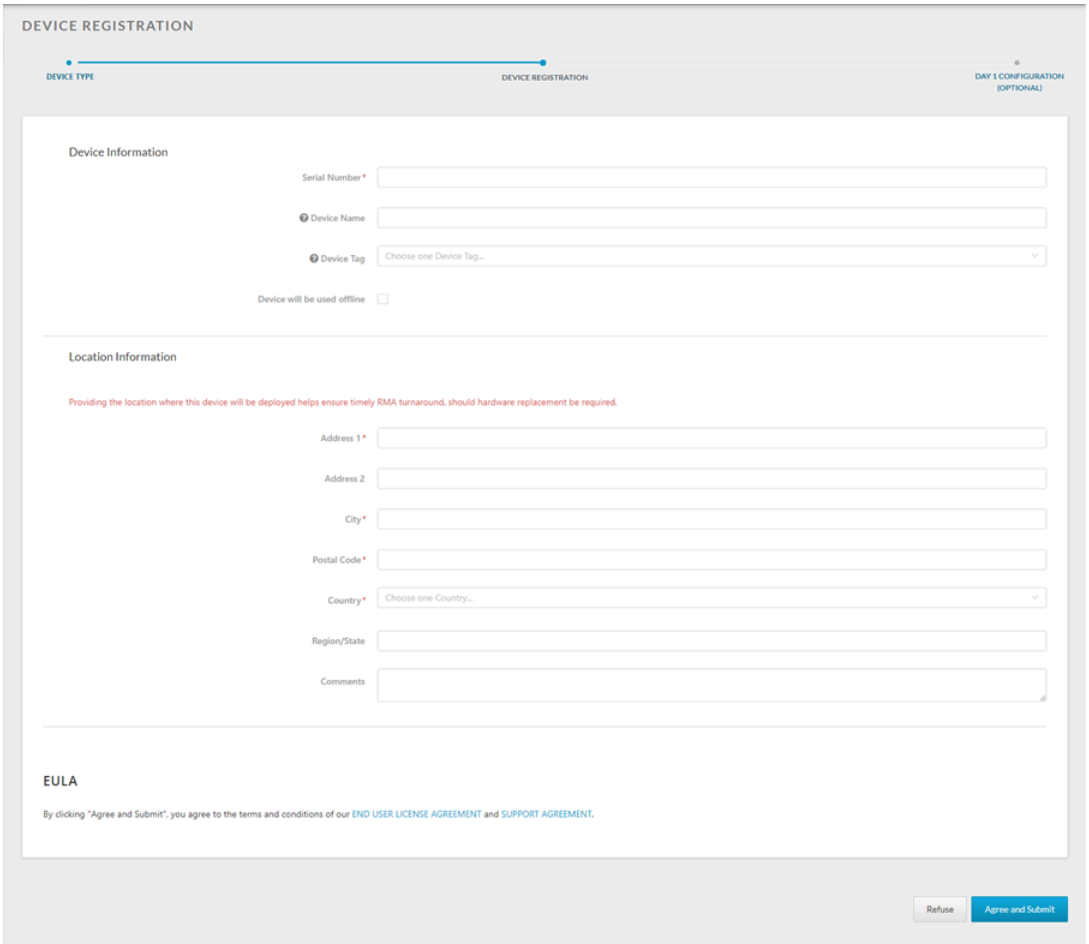
2. Select **Register device using Serial Number or Authorization Code**, and then click **Next**.



3. Enter the firewall **Serial Number** (you can copy and paste it from the firewall Dashboard).
4. (Optional) Enter the **Device Name** and **Device Tag**.
5. (Optional) If the device will not have a connection to the internet, select the **Device will be used offline** check box and then, from the drop-down, select the **OS Release** you plan to use.
6. Provide information about where you plan to deploy the firewall including the **Address**, **City**, **Postal Code**, and **Country**.
7. Read the End User License Agreement (EULA) and the Support Agreement, then **Agree and Submit**.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets



You can view the entry for the firewall you just registered under **Devices**.

STEP 5 | (Firewalls with line cards) To ensure that you receive support for your firewall's line cards, make sure to Register the Firewall Line Cards.

(Optional) Perform Day 1 Configuration

After you register your firewall, you have the option of running Day 1 Configuration. The Day 1 Configuration tool provides configuration templates informed by Palo Alto Networks best practices, which you can use as a starting point to build the rest of your configuration.

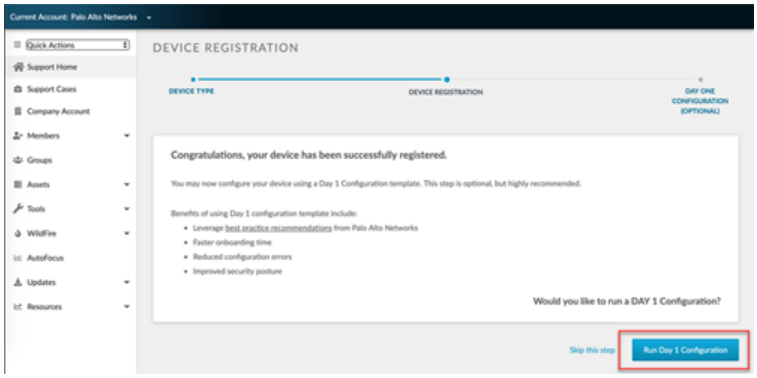
The benefits of Day 1 Configuration templates include:

- Faster implementation time
- Reduced configuration errors
- Improved security posture

Perform Day 1 Configuration by following these steps:

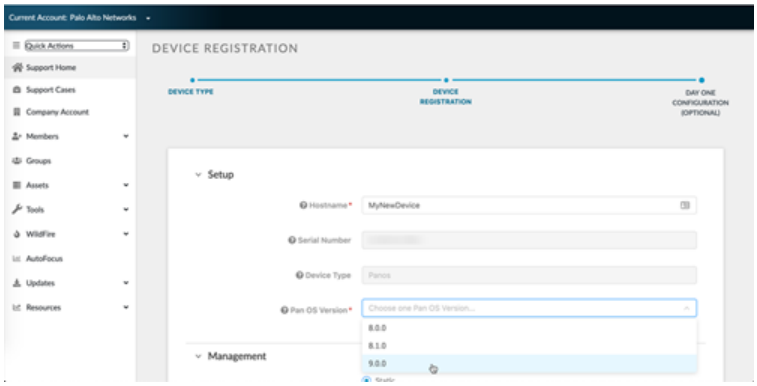
https://t.me/learningnets

STEP 1 | From the page that displays after you have registered your firewall, select **Run Day 1 Configuration**.



*If you've already registered your firewall but haven't run Day 1 Configuration, you can also run it from the Customer Support Portal home page by selecting **Tools** > **Run Day 1 Configuration.***

STEP 2 | Enter the **Hostname** and **Pan OS Version** for your new device, and optionally, the **Serial Number** and **Device Type**.

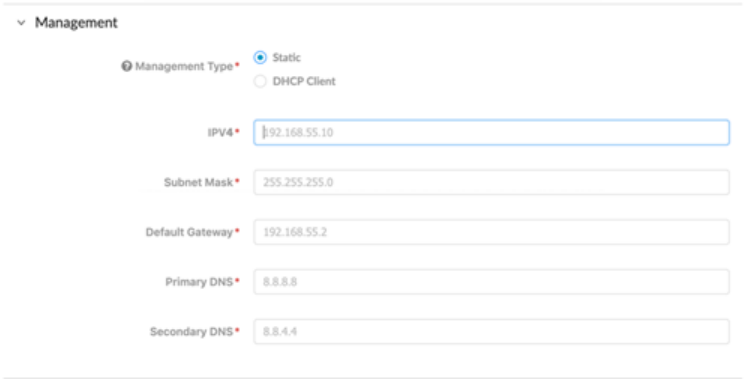


 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 3 | Under **Management**, select either **Static** or **DHCP Client** for your **Management Type**.

Selecting **Static** will require you fill out the **IPV4**, **Subnet Mask**, and **Default Gateway** fields.

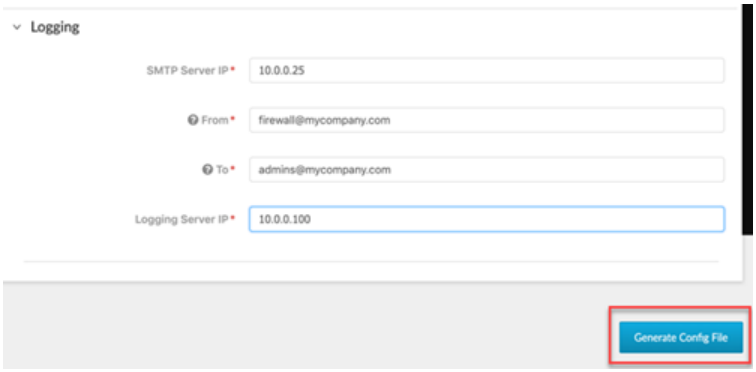


Selecting **DHCP Client** only requires that you enter the **Primary DNS** and **Secondary DNS**. A device configured in DHCP client mode will ensure the management interface receives an IP address from the local DHCP server, or it will fill out all the parameters if they are known.



STEP 4 | Fill out all fields under **Logging**.

STEP 5 | Click **Generate Config File**.

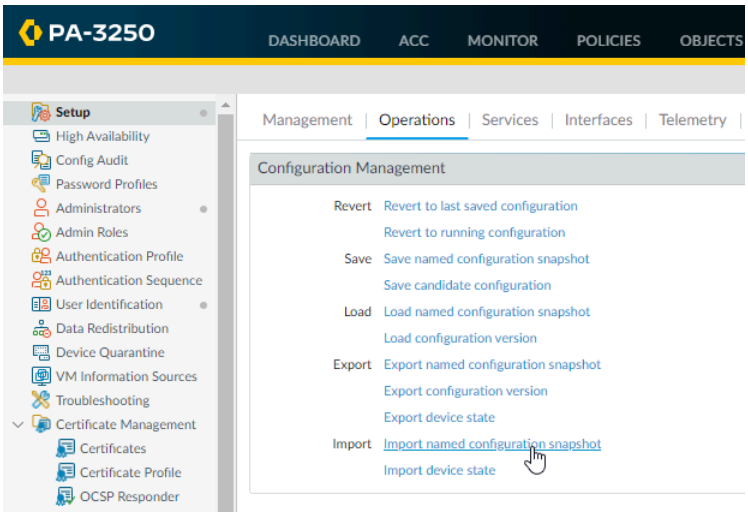


 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 6 | To import and load the Day 1 Configuration file you just downloaded to your firewall:

1. Log into your firewall web interface.
2. Select **Device** > **Setup** > **Operations**.
3. Click **Import named configuration snapshot**.
4. Select the file.



Register the Firewall Line Cards

The following firewalls use line cards that must be registered to receive support with troubleshooting and returns:

- PA-7000 Series firewalls
- PA-5450 firewall

If you do not have a Palo Alto Networks Customer Support account, create one by following the steps at Create a New Support Account and Register a Firewall. Return to these instructions after creating your Customer Support account and registering your firewall.

STEP 1 | Go to the Palo Alto Networks Customer Support Portal and, if not already logged in, **Sign In** now.

STEP 2 | Select **Assets** > **Line Cards/Optics/FRUs**.

STEP 3 | **Register Components**.

STEP 4 | Enter the Palo Alto Networks Sales Order Number of the line cards into the **Sales Order Number** field to display the line cards eligible for registration.

STEP 5 | Register the line cards to your firewall by entering its chassis serial number in the **Serial Number** field. The **Location Information** below auto-populates based on the registration information of your firewall.

STEP 6 | Click **Agree and Submit** to accept the legal terms. The system updates to display the registered line cards under **Assets** > **Line Cards/Optics/FRUs**.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

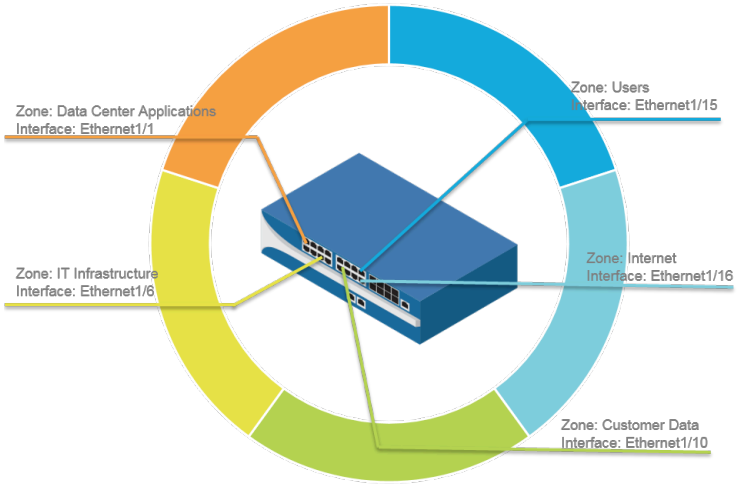
Segment Your Network Using Interfaces and Zones

Traffic must pass through the firewall in order for the firewall to manage and control it. Physically, traffic enters and exits the firewall through *interfaces*. The firewall determines how to act on a packet based on whether the packet matches a *Security policy rule*. At the most basic level, each Security policy rule must identify where the traffic came from and where it is going. On a Palo Alto Networks next-generation firewall, Security policy rules are applied between zones. A *zone* is a grouping of interfaces (physical or virtual) that represents a segment of your network that is connected to, and controlled by, the firewall. Because traffic can only flow between zones if there is a Security policy rule to allow it, this is your first line of defense. The more granular the zones you create, the greater control you have over access to sensitive applications and data and the more protection you have against malware moving laterally throughout your network. For example, you might want to segment access to the database servers that store your customer data into a zone called Customer Data. You can then define security policies that only permit certain users or groups of users to access the Customer Data zone, thereby preventing unauthorized internal or external access to the data stored in that segment.

- Network Segmentation for a Reduced Attack Surface
- Configure Interfaces and Zones

Network Segmentation for a Reduced Attack Surface

The following diagram shows a very basic example of Network Segmentation Using Zones. The more granular you make your zones (and the corresponding security policy rules that allows traffic between zones), the more you reduce the attack surface on your network. This is because traffic can flow freely within a zone (intra-zone traffic), but traffic cannot flow between zones (inter-zone traffic) until you define a Security policy rule that allows it. Additionally, an interface cannot process traffic until you have assigned it to a zone. Therefore, by segmenting your network into granular zones you have more control over access to sensitive applications or data and you can prevent malicious traffic from establishing a communication channel within your network, thereby reducing the likelihood of a successful attack on your network.



 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

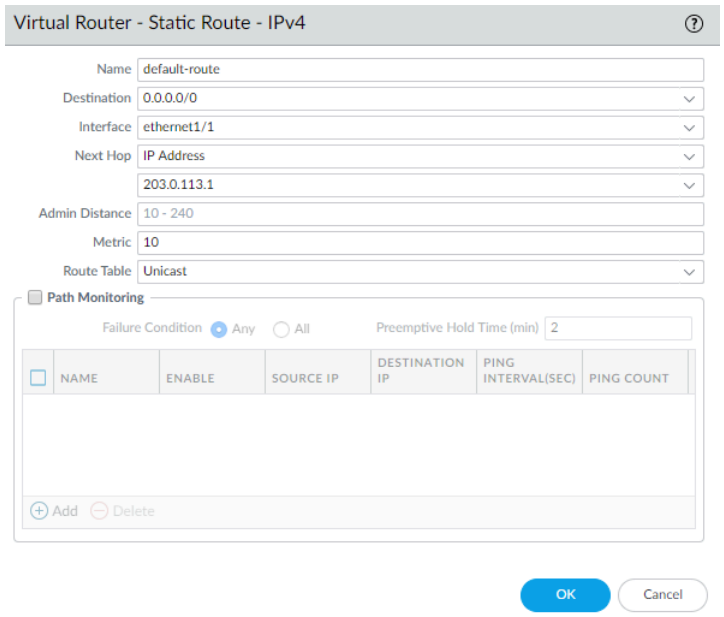
Configure Interfaces and Zones

After you identify how you want to segment your network and the zones you will need to create to achieve the segmentation (as well as the interfaces to map to each zone), you can begin configuring the interfaces and zones on the firewall. Configure interfaces on the firewall the to support the topology of each part of the network you are connecting to. The following workflow shows how to configure Layer 3 interfaces and assign them to zones. For details on integrating the firewall using a different type of interface deployments (for example as virtual wire interfaces or as Layer 2 interfaces), see the PAN-OS Networking Adminstrator's Guide.

The firewall comes preconfigured with a default virtual wire interface between ports Ethernet 1/1 and Ethernet 1/2 (and a corresponding default security policy and virtual router). If you do not plan to use the default virtual wire, you must manually delete the configuration and commit the change before proceeding to prevent it from interfering with other settings you define. For instructions on how to delete the default virtual wire and its associated security policy and zones, see Step 3 in Set Up Network Access for External Services.

STEP 1 | Configure a default route to your Internet router.

1. Select **Network** > **Virtual Router** and then select the **default** link to open the Virtual Router dialog.
2. Select the **Static Routes** tab and click **Add**. Enter a **Name** for the route and enter the route in the **Destination** field (for example, 0.0.0.0/0).
3. Select the **IP Address** radio button in the **Next Hop** field and then enter the IP address and netmask for your Internet gateway (for example, 203.0.113.1).



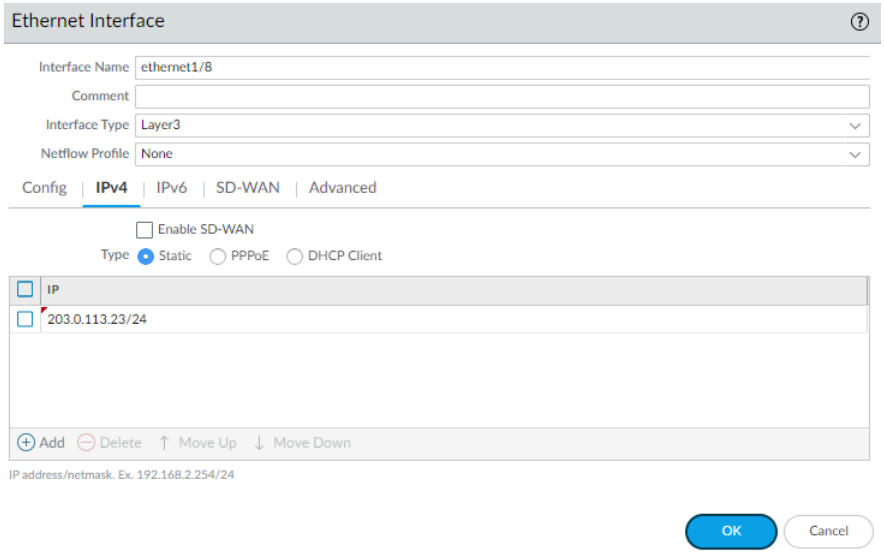
4. Click **OK** twice to save the virtual router configuration.

©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 2 | Configure the external interface (the interface that connects to the Internet).

1. Select **Network** > **Interfaces** and then select the interface you want to configure. In this example, we are configuring Ethernet1/8 as the external interface.
2. Select the **Interface Type**. Although your choice here depends on interface topology, this example shows the steps for **Layer3**.
3. On the **Config** tab, select **New Zone** from the **Security Zone** drop-down. In the Zone dialog, define a **Name** for new zone, for example Internet, and then click **OK**.
4. In the **Virtual Router** drop-down, select **default**.
5. To assign an IP address to the interface, select the **IPv4** tab, click **Add** in the IP section, and enter the IP address and network mask to assign to the interface, for example 203.0.113.23/24.



6. To enable you to ping the interface, select **Advanced** > **Other Info**, expand the **Management Profile** drop-down, and select **New Management Profile**. Enter a **Name** for the profile, select **Ping** and then click **OK**.
7. To save the interface configuration, click **OK**.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 3 | Configure the interface that connects to your internal network.

In this example, the interface connects to a network segment that uses private IP addresses. Because private IP addresses cannot be routed externally, you have to configure NAT.

1. Select **Network** > **Interfaces** and select the interface you want to configure. In this example, we are configuring Ethernet1/15 as the internal interface our users connect to.
2. Select **Layer3** as the **Interface Type**.
3. On the **Config** tab, expand the **Security Zone** drop-down and select **New Zone**. In the Zone dialog, define a **Name** for new zone, for example Users, and then click **OK**.
4. Select the same Virtual Router you used previously, default in this example.
5. To assign an IP address to the interface, select the **IPv4** tab, click **Add** in the IP section, and enter the IP address and network mask to assign to the interface, for example 192.168.1.4/24.
6. To enable you to ping the interface, select the management profile that you just created.
7. To save the interface configuration, click **OK**.

STEP 4 | Configure the interface that connects to your data center applications.

Make sure you define granular zones to prevent unauthorized access to sensitive applications or data and eliminate the possibility of malware moving laterally within your data center.

1. Select the interface you want to configure.
2. Select **Layer3** from the **Interface Type** drop-down. In this example, we are configuring Ethernet1/1 as the interface that provides access to your data center applications.
3. On the **Config** tab, expand the **Security Zone** drop-down and select **New Zone**. In the Zone dialog, define a **Name** for new zone, for example Data Center Applications, and then click **OK**.
4. Select the same Virtual Router you used previously, default in this example.
5. To assign an IP address to the interface, select the **IPv4** tab, click **Add** in the IP section, and enter the IP address and network mask to assign to the interface, for example 10.1.1.1/24.
6. To enable you to ping the interface, select the management profile that you created.
7. To save the interface configuration, click **OK**.

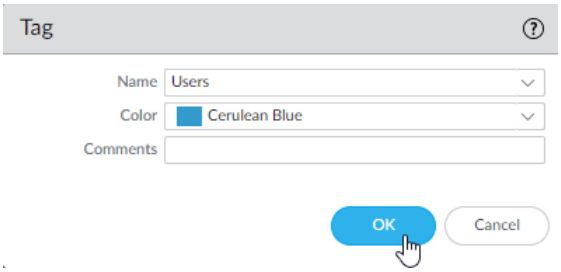
 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 5 | (Optional) Create tags for each zone.

Tags allow you to visually scan policy rules.

1. Select **Objects** > **Tags** and **Add**.
2. Select a zone **Name**.
3. Select a tag **Color** and click **OK**.



STEP 6 | Save the interface configuration.

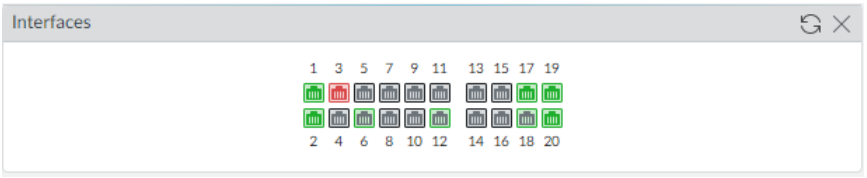
Click **Commit**.

STEP 7 | Cable the firewall.

Attach straight through cables from the interfaces you configured to the corresponding switch or router on each network segment.

STEP 8 | Verify that the interfaces are active.

Select **Dashboard** and verify that the interfaces you configured show as green in the Interfaces widget.



 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Set Up a Basic Security Policy

Now that you defined some zones and attached them to interfaces, you are ready to begin creating your Security Policy. The firewall will not allow any traffic to flow from one zone to another unless there is a Security policy rule that allows it. When a packet enters a firewall interface, the firewall matches the attributes in the packet against the Security policy rules to determine whether to block or allow the session based on attributes such as the source and destination security zone, the source and destination IP address, the application, user, and the service. The firewall evaluates incoming traffic against the Security policy rulebase from left to right and from top to bottom and then takes the action specified in the first Security rule that matches (for example, whether to allow, deny, or drop the packet). This means that you must order the rules in your Security policy rulebase so that more specific rules are at the top of the rulebase and more general rules are at the bottom to ensure that the firewall is enforcing policy as expected.

Even though a Security policy rule allows a packet, this does not mean that the traffic is free of threats. To enable the firewall to scan the traffic that it allows based on a Security policy rule, you must also attach Security Profiles—including URL Filtering, Antivirus, Anti-Spyware, File Blocking, and WildFire Analysis—to each rule (the profiles you can use depend on which Subscriptions you purchased). When creating your basic Security policy, use the predefined security profiles to ensure that the traffic you allow into your network is being scanned for threats. You can customize these profiles later as needed for your environment.

Use the following workflow set up a very basic Security policy that enables access to the network infrastructure, to data center applications, and to the internet. This enables you to get the firewall up and running so that you can verify that you have successfully configured the firewall. However, this initial policy is not comprehensive enough to protect your network. After you verify that you successfully configured the firewall and integrated it into your network, proceed with creating a Best Practice Internet Gateway Security Policy that safely enables application access while protecting your network from attack.

STEP 1 | (Optional) Delete the default Security policy rule.

By default, the firewall includes a Security policy rule named *rule1* that allows all traffic from Trust zone to Untrust zone. You can either delete the rule or modify the rule to reflect your zone-naming conventions.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 2 | Allow access to your network infrastructure resources.

1. Select **Policies** > **Security** and click **Add**.
2. In the **General** tab, enter a descriptive **Name** for the rule.
3. In the **Source** tab, set the **Source Zone** to **Users**.
4. In the **Destination** tab, set the **Destination Zone** to **IT Infrastructure**.

   *As a best practice, use address objects in the **Destination Address** field to enable access to specific servers or groups of servers only, particularly for services such as DNS and SMTP that are commonly exploited. By restricting users to specific destination server addresses, you can prevent data exfiltration and command and control traffic from establishing communication through techniques such as DNS tunneling.*
5. In the **Applications** tab, **Add** the applications that correspond to the network services you want to safely enable. For example, select **dns**, **ntp**, **ocsp**, **ping**, and **smtp**.
6. In the **Service/URL Category** tab, keep the **Service** set to **application-default**.
7. In the **Actions** tab, set the **Action Setting** to **Allow**.
8. Set **Profile Type** to **Profiles** and select the following security profiles to attach to the policy rule:
   - For **Antivirus**, select **default**
   - For **Vulnerability Protection**, select **strict**
   - For **Anti-Spyware**, select **strict**
   - For **URL Filtering**, select **default**
   - For **File Blocking**, select **basic file blocking**
   - For **WildFire Analysis**, select **default**
9. Verify that **Log at Session End** is enabled. Only traffic that matches a Security policy rule will be logged.
10. Click **OK**.

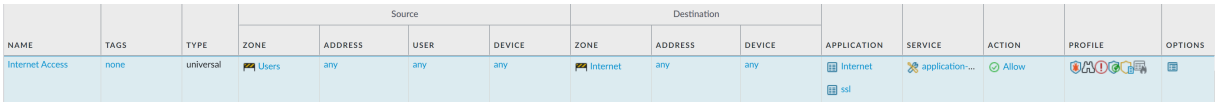
			Source				Destination							
NAME	TAGS	TYPE	ZONE	ADDRESS	USER	DEVICE	ZONE	ADDRESS	DEVICE	APPLICATION	SERVICE	ACTION	PROFILE	OPTIONS
Network Infrastructu...	none	universal	Users	any	any	any	IT Infrastruct...	any	any	dns ntp ocsp ping smtp	application-...	Allow		

https://t.me/learningnets

STEP 3 | Enable access to general internet applications.

This is a temporary rule that allows you to gather information about the traffic on your network. After you have more insight into which applications your users need to access, you can make informed decisions about which applications to allow and create more granular application-based rules for each user group.

1. Select **Policies** > **Security** and **Add** a rule.
2. In the **General** tab, enter a descriptive **Name** for the rule.
3. In the **Source** tab, set the **Source Zone** to **Users**.
4. In the **Destination** tab, set the **Destination Zone** to **Internet**.
5. In the **Applications** tab, **Add** an **Application Filter** and enter a **Name**. To safely enable access to legitimate web-based applications, set the **Category** in the application filter to **general-internet** and then click **OK**. To enable access to encrypted sites, **Add** the **ssl** application.
6. In the **Service/URL Category** tab, keep the **Service** set to **application-default**.
7. In the **Actions** tab, set the **Action Setting** to **Allow**.
8. Set **Profile Type** to **Profiles** and select the following security profiles to attach to the policy rule:
   - For **Antivirus**, select **default**
   - For **Vulnerability Protection**, select **strict**
   - For **Anti-Spyware**, select **strict**
   - For **URL Filtering**, select **default**
   - For **File Blocking**, select **strict file blocking**
   - For **WildFire Analysis**, select **default**
9. Verify that **Log at Session End** is enabled. Only traffic that matches a security rule will be logged.
10. Click **OK**.

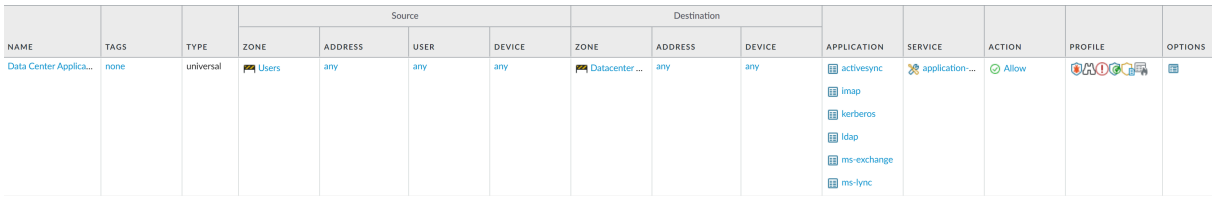


 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 4 | Enable access to data center applications.

1. Select **Policies** > **Security** and **Add** a rule.
2. In the **General** tab, Enter a descriptive **Name** for the rule.
3. In the **Source** tab, set the **Source Zone** to **Users**.
4. In the **Destination** tab, set the **Destination Zone** to **Data Center Applications**.
5. In the **Applications** tab, **Add** the applications that correspond to the network services you want to safely enable. For example, select **activesync**, **imap**, **kerberos**, **ldap**, **ms-exchange**, and **ms-lync**.
6. In the **Service/URL Category** tab, keep the **Service** set to **application-default**.
7. In the **Actions** tab, set the **Action Setting** to **Allow**.
8. Set **Profile Type** to **Profiles** and select the following security profiles to attach to the policy rule:
   - For **Antivirus**, select **default**
   - For **Vulnerability Protection** select **strict**
   - For **Anti-Spyware** select **strict**
   - For **URL Filtering** select **default**
   - For **File Blocking** select **basic file blocking**
   - For **WildFire Analysis** select **default**
9. Verify that **Log at Session End** is enabled. Only traffic that matches a security rule will be logged.
10. Click **OK**.



STEP 5 | Save your policy rules to the running configuration on the firewall.

Click **Commit**.

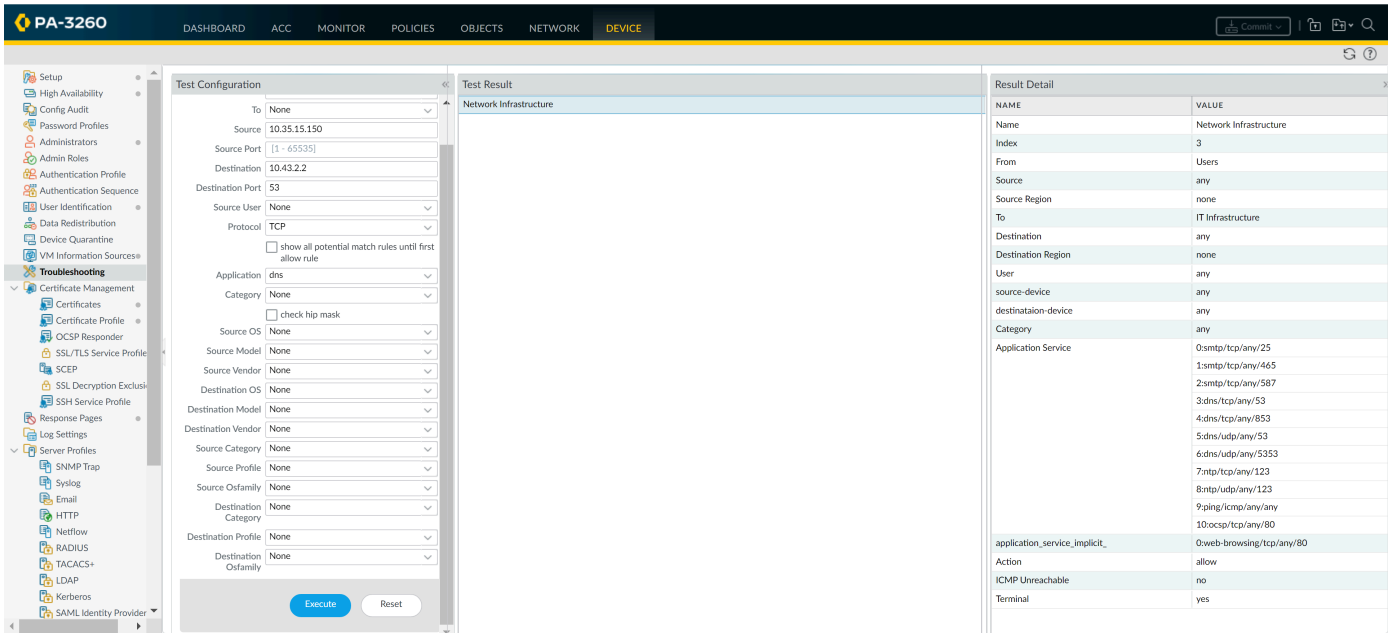
 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 6 | To verify that you have set up your basic policies effectively, test whether your Security policy rules are being evaluated and determine which Security policy rule applies to a traffic flow.

For example, to verify the policy rule that will be applied for a client in the user zone with the IP address 10.35.14.150 when it sends a DNS query to the DNS server in the data center:

1. Select **Device** > **Troubleshooting** and select **Security Policy Match** (**Select Test**).
2. Enter the **Source** and **Destination** IP addresses.
3. Enter the **Protocol**.
4. Select **dns** (**Application**)
5. **Execute** the Security policy match test.



 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Assess Network Traffic

Now that you have a basic security policy, you can review the statistics and data in the Application Command Center (ACC), traffic logs, and the threat logs to observe trends on your network. Use this information to identify where you need to create more granular security policy rules.

- Use the Application Command Center and Use the Automated Correlation Engine.

  In the ACC, review the most used applications and the high-risk applications on your network. The ACC graphically summarizes the log information to highlight the applications traversing the network, who is using them (with User-ID enabled), and the potential security impact of the content to help you identify what is happening on the network in real time. You can then use this information to create appropriate security policy rules that block unwanted applications, while allowing and enabling applications in a secure manner.

  The Compromised Hosts widget in **ACC** > **Threat Activity** displays potentially compromised hosts on your network and the logs and match evidence that corroborates the events.

- Determine what updates/modifications are required for your network security policy rules and implement the changes.

  For example:

  - Evaluate whether to allow web content based on schedule, users, or groups.
  - Allow or control certain applications or functions within an application.
  - Decrypt and inspect content.
  - Allow but scan for threats and exploits.

  For information on refining your security policies and for attaching custom security profiles, see how to Create a Security Policy Rule and Security Profiles.

- View Logs.

  Specifically, view the traffic and threat logs (**Monitor** > **Logs**).

  Traffic logs are dependent on how your security policies are defined and set up to log traffic. The Application Usage widget in the ***ACC****, however, records applications and statistics regardless of policy configuration; it shows all traffic that is allowed on your network, therefore it includes the inter-zone traffic that is allowed by policy and the same zone traffic that is allowed implicitly.*

- Configure Log Storage Quotas and Expiration Periods.

  Review the AutoFocus intelligence summary for artifacts in your logs. An *artifact* is an item, property, activity, or behavior associated with logged events on the firewall. The intelligence summary reveals the number of sessions and samples in which WildFire detected the artifact.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Use WildFire verdict information (benign, grayware, malware) and AutoFocus matching tags to look for potential risks in your network.



AutoFocus tags created by Unit 42, the Palo Alto Networks threat intelligence team, call attention to advanced, targeted campaigns and threats in your network.

From the AutoFocus intelligence summary, you can start an AutoFocus search for artifacts and assess their pervasiveness within global, industry, and network contexts.

- Monitor Web Activity of Network Users.

  Review the URL filtering logs to scan through alerts, denied categories/URLs. URL logs are generated when a traffic matches a security rule that has a URL filtering profile attached with an action of alert, continue, override or block.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Enable Free WildFire Forwarding

WildFire is a cloud-based virtual environment that analyzes and executes unknown samples (files and email links) and determines the samples to be malicious, phishing, grayware, or benign. With WildFire enabled, a Palo Alto Networks firewall can forward unknown samples to WildFire for analysis. For newly-discovered malware, WildFire generates a signature to detect the malware, which is made available for retrieval in real-time for all firewalls with an active WildFire subscription. This enables all Palo Alto next-generation firewalls worldwide to detect and prevent malware found by a single firewall. Malware signatures often match multiple variants of the same malware family, and as such, block new malware variants that the firewall has never seen before. The Palo Alto Networks threat research team uses the threat intelligence gathered from malware variants to block malicious IP addresses, domains, and URLs.

A basic WildFire service is included as part of the Palo Alto Networks next-generation firewall and does not require a WildFire subscription. With the basic WildFire service, you can enable the firewall to forward portable executable (PE) files. Additionally, if you do not have a WildFire subscription, but you do have a Threat Prevention subscription, you can receive signatures for malware WildFire identifies every 24- 48 hours (as part of the Antivirus updates).

Beyond the basic WildFire service, a WildFire subscription is required for the firewall to:

- Get the latest WildFire signatures in real-time.
- Prevent malicious PE (portable executables), ELF and MS Office files, and PowerShell and shell scripts from entering your network in real-time using WildFire Inline ML.
- Forward advanced file types and email links for analysis.
- Use the WildFire API.
- Use a WildFire appliance to host a WildFire private cloud or a WildFire hybrid cloud.

If you have a WildFire subscription, go ahead and get started with WildFire to get the most out of your subscription. Otherwise, take the following steps to enable basic WildFire forwarding:

STEP 1 | Confirm that your firewall is registered and that you have a valid support account as well as any subscriptions you require.

1. Log in to the Palo Alto Networks Customer Support Portal(CSP) and on the left-hand side navigation pane, select **Assets** > **Devices**.
2. Verify that the firewall is listed. If it is not listed, select **Register New Device** and continue to Register the Firewall.
3. (Optional) If you have a Threat Prevention subscription, be sure to Activate Subscription Licenses.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 2 | Log in to the firewall and configure WildFire forwarding settings.

1. Select **Device** > **Setup** > **WildFire** and edit the General Settings.
2. Set the **WildFire Public Cloud** field to forward files to the WildFire global cloud (U.S.) at: **`wildfire.paloaltonetworks.com`**.

   You can also forward files to a WildFire regional cloud or a private cloud based on your location and your organizational requirements.
3. Review the **File Size Limits** for PEs the firewall forwards for WildFire analysis. set the **Size Limit** for PEs that the firewall can forward to the maximum available limit of 10 MB.

   *As a WildFire best practice, set the **Size Limit** for PEs to the maximum available limit of 10 MB.*
4. Click **OK** to save your changes.

STEP 3 | Enable the firewall to forward PEs for analysis.

1. Select **Objects** > **Security Profiles** > **WildFire Analysis** and **Add** a new profile rule.
2. **Name** the new profile rule.
3. **Add** a forwarding rule and enter a **Name** for it.
4. In the **File Types** column, add **pe** files to the forwarding rule.
5. In the **Analysis** column, select **public-cloud** to forward PEs to the WildFire public cloud.
6. Click **OK**.

STEP 4 | Apply the new WildFire Analysis profile to traffic that the firewall allows.

1. Select **Policies** > **Security** and either select an existing policy rule or create a new policy rule as described in Set Up a Basic Security Policy.
2. Select **Actions** and in the Profile Settings section, set the **Profile Type** to **Profiles**.
3. Select the **WildFire Analysis** profile you just created to apply that profile rule to all traffic this policy rule allows.
4. Click **OK**.

STEP 5 | Enable the firewall to forward decrypted SSL traffic for WildFire analysis.

STEP 6 | Review and implement WildFire best practices to ensure that you are getting the most of WildFire detection and prevention capabilities.

STEP 7 | **Commit** your configuration updates.

STEP 8 | Verify that the firewall is forwarding PE files to the WildFire public cloud.

Select **Monitor** > **Logs** > **WildFire Submissions** to view log entries for PEs the firewall successfully submitted for WildFire analysis. The Verdict column displays whether WildFire found the PE to be malicious, grayware, or benign. (WildFire only assigns the phishing verdict to email links). The Action column indicates whether the firewall allowed or blocked the sample. The Severity column indicates how much of a threat a sample poses to an organization using the following values: critical, high, medium, low, information.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 9 | (Threat Prevention subscription only) If you have a Threat Prevention subscription, but do not have a WildFire subscription, you can still receive WildFire signature updates every 24-48 hours.

1. Select **Device** > **Dynamic Updates**.
2. Check that the firewall is scheduled to download, and install Antivirus updates.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Best Practices for Completing the Firewall Deployment

Now that you have integrated the firewall into your network and enabled the basic security features, you can begin configuring more advanced features. Here are some things to consider next:

- Follow the Adminstrative Access Best Practices to make sure you are properly securing the management interfaces.
- Configure a best-practice security policy rulebase to safely enable applications and protect your network from attack. Go to the Best Practices page and select security policy best practice for your firewall deployment.
- Set up High Availability—High availability (HA) is a configuration in which two firewalls are placed in a group and their configuration and session tables are synchronized to prevent a single point to failure on your network. A heartbeat connection between the firewall peers ensures seamless failover in the event that a peer goes down. Setting up a two-firewall cluster provides redundancy and allows you to ensure business continuity.
- Enable User Identification (User-ID)—User-ID is a Palo Alto Networks next-generation firewall feature that allows you to create policies and perform reporting based on users and groups rather than individual IP addresses.
- Enable Decryption—Palo Alto Networks firewalls provide the capability to decrypt and inspect traffic for visibility, control, and granular security. Use decryption on a firewall to prevent malicious content from entering your network or sensitive content from leaving your network concealed as encrypted or tunneled traffic.
- Follow the Best Practices for Securing Your Network from Layer 4 and Layer 7 Evasions.
- Share Threat Intelligence with Palo Alto Networks—Permit the firewall to periodically collect and send information about applications, threats, and device health to Palo Alto Networks. Telemetry includes options to enable passive DNS monitoring and to allow experimental test signatures to run in the background with no impact to your security policy rules, firewall logs, or firewall performance. All Palo Alto Networks customers benefit from the intelligence gathered from telemetry, which Palo Alto Networks uses to improve the threat prevention capabilities of the firewall.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Subscriptions

Learn about all the subscriptions and services that work with the firewall, and get started by activating subscription licenses:

- Subscriptions You Can Use With the Firewall
- Activate Subscription Licenses
- What Happens When Licenses Expire?
- Enhanced Application Logs for Palo Alto Networks Cloud Services

Certain cloud services, like Cortex XDR™, do not integrate with the firewall directly, but rely on data stored in Cortex Data Lake for visibility into network activity. Enhanced application logging is a feature that comes with a Cortex Data Lake subscription—it allows the firewall to collect data specifically for Cortex XDR to use to detect anomalous network activity. Turning on enhanced application logging is a Cortex XDR best practice.

https://t.me/learningnets

Subscriptions You Can Use With the Firewall

The following Palo Alto Networks subscriptions unlock certain firewall features or enable the firewall to leverage a Palo Alto Networks cloud-delivered service (or both). Here you can read more about each service or feature that requires a subscription to work with the firewall. To enable a subscription, you must first Activate Subscription Licenses; once active, most subscription services can use Dynamic Content Updates to provide new and updated functionality to the firewall.

Subscriptions You Can Use With the Firewall	
IoT Security	The IoT Security solution works with next-generation firewalls to dynamically discover and maintain a real-time inventory of the IoT devices on your network. Through AI and machine-learning algorithms, the IoT Security solution achieves a high level of accuracy, even classifying IoT device types encountered for the first time. And because it's dynamic, your IoT device inventory is always up to date. IoT Security also provides the automatic generation of policy recommendations to control IoT device traffic, as well as the automatic creation of IoT device attributes for use in firewall policies. • Get Started with IoT Security.
SD-WAN	Provides intelligent and dynamic path selection on top of the industry-leading security that PAN-OS software already delivers. Managed by Panorama, the SD-WAN implementation includes: • Centralized configuration management • Automatic VPN topology creation • Traffic distribution • Monitoring and troubleshooting • Get Started with SD-WAN
Threat Prevention	Threat Prevention provides: • Antivirus, anti-spyware (command-and-control), and vulnerability protection. • Built-in external dynamic lists that you can use to secure your network against malicious hosts. • Ability to identify infected hosts that try to connect to malicious domains. • Get Started with Threat Prevention
DNS Security	Provides enhanced DNS sinkholing capabilities by querying DNS Security, an extensible cloud-based service capable of

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Subscriptions You Can Use With the Firewall	
	generating DNS signatures using advanced predictive analytics and machine learning. This service provides full access to the continuously expanding DNS-based threat intelligence produced by Palo Alto Networks. To set up DNS Security, you must first purchase and install a Threat Prevention license. • Get Started with DNS Security
URL Filtering	Provides the ability to not only control web-access, but how users interact with online content based on dynamic URL categories. You can also prevent credential theft by controlling the sites to which users can submit their corporate credentials. To set up URL Filtering, you must purchase and install a subscription for the supported URL filtering database, PAN-DB. With PAN-DB, you can set up access to the PAN-DB public cloud or to the PAN-DB private cloud. *URL filtering is no longer available as a standalone subscription. All features contained in URL filtering are included with the Advanced URL filtering subscription.* • Get Started with URL Filtering
Advanced URL Filtering	Advanced URL Filtering uses a cloud-based ML-powered web security engine to perform ML-based inspection of web traffic in real-time. This reduces reliance on URL databases and out-of-band web crawling to detect and prevent advanced, file-less web-based attacks including targeted phishing, web-delivered malware and exploits, command-and-control, social engineering, and other types of web attacks. • Get Started with Advanced URL Filtering
WildFire	Although basic WildFire® support is included as part of the Threat Prevention license, the WildFire subscription service provides enhanced services for organizations that require immediate coverage for threats, frequent WildFire signature updates, advanced file type forwarding (APK, PDF, Microsoft Office, and Java Applet), as well as the ability to upload files using the WildFire API. A WildFire subscription is also required if your firewalls will be forwarding files to an on-premise WF-500 appliance. • Get Started with WildFire
AutoFocus	Provides a graphical analysis of firewall traffic logs and identifies potential risks to your network using threat

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Subscriptions You Can Use With the Firewall	
	intelligence from the AutoFocus portal. With an active license, you can also open an AutoFocus search based on logs recorded on the firewall. • Get Started with AutoFocus
Cortex Data Lake	Provides cloud-based, centralized log storage and aggregation. The Cortex Data Lake is required or highly-recommended to support several other cloud-delivered services, including Cortex XDR, IoT Security, and Prisma Access, and Traps management service. • Get Started with Cortex Data Lake
GlobalProtect Gateway	Provides mobility solutions and/or large-scale VPN capabilities. By default, you can deploy GlobalProtect portals and gateways (without HIP checks) without a license. If you want to use advanced GlobalProtect features (HIP checks and related content updates, the GlobalProtect Mobile App, IPv6 connections, or a GlobalProtect Clientless VPN) you will need a GlobalProtect Gateway license for each gateway. • Get Started with GlobalProtect
Virtual Systems	This is a perpetual license, and is required to enable support for multiple virtual systems on PA-3200 Series firewalls. In addition, you must purchase a Virtual Systems license if you want to increase the number of virtual systems beyond the base number provided by default on PA-5200 Series, and PA-7000 Series firewalls (the base number varies by platform). The PA-800 Series, PA-220, and VM-Series firewalls do not support virtual systems. • Get Started with Virtual Systems
Enterprise Data Loss Prevention (DLP)	Provides cloud-based protection against unauthorized access, misuse, extraction, and sharing of sensitive information. Enterprise DLP provides a single engine for accurate detection and consistent policy enforcement for sensitive data at rest and in motion using machine learning-based data classification, hundreds of data patterns using regular expressions or keywords, and data profiles using Boolean logic to scan for collective types of data. • Get Started with Enterprise Data loss Prevention
SaaS Security Inline	The SaaS Security solution works with Cortex Data Lake to discover all of the SaaS applications in use on your network. SaaS Security Inline can discover thousands of Shadow IT applications and their users and usage details. SaaS Security Inline also enforces SaaS policy rule recommendations

©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Subscriptions You Can Use With the Firewall	
	seamlessly across your existing Palo Alto Networks firewalls. App-ID Cloud Engine (ACE) also requires SaaS Security Inline. • Get Started with Prisma SaaS

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Activate Subscription Licenses

Follow these steps to activate a new license on the firewall.

The Decryption Mirroring feature requires you to activate a free license to unlock feature functionality. For those features, you should instead follow the steps to Activate Free Licenses for Decryption Features.

STEP 1 | Locate the activation codes for the licenses you purchased.

When you purchased your subscriptions you should have received an email from Palo Alto Networks customer service listing the activation code associated with each subscription. If you cannot locate this email, contact Customer Support to obtain your activation codes before you proceed.

STEP 2 | Activate your Support license.

You will not be able to update your PAN-OS software if you do not have a valid Support license.

1. Log in to the web interface and then select **Device** > **Support**.
2. Click **Activate support using authorization code**.
3. Enter your **Authorization Code** and then click **OK**.

STEP 3 | Activate each license you purchased.

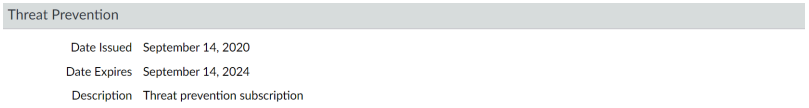
Select **Device** > **Licenses** and then activate your licenses and subscriptions in one of the following ways:

- **Retrieve license keys from license server**—Use this option if you activated your license on the Customer Support portal.
- **Activate feature using authorization code**—Use this option to enable purchased subscriptions using an authorization code for licenses that have not been previously activated on the support portal. When prompted, enter the **Authorization Code** and then click **OK**.
- **Manually upload license key**—Use this option if your firewall does not have connectivity to the Palo Alto Networks Customer Support Portal. In this case, you must download a license key file from the support site on an internet-connected computer and then upload to the firewall.

To automate activation using the Customer Support Portal API, see the process to Activate Licenses. This process works for both the hardware and VM-Series firewalls.

STEP 4 | Verify that the license is successfully activated

On the **Device** > **Licenses** page, verify that the license is successfully activated. For example, after activating the WildFire license, you should see that the license is valid:



 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 5 | (WildFire subscriptions only) Perform a commit to complete WildFire subscription activation.

After activating a WildFire subscription, a commit is required for the firewall to begin forwarding advanced file types. You should either:

- Commit any pending changes.
- Check that the WildFire Analysis profile rules include the advanced file types that are now supported with the WildFire subscription. If no change to any of the rules is required, make a minor edit to a rule description and perform a commit.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

What Happens When Licenses Expire?

Palo Alto Networks subscriptions provide the firewall with added functionality and/or access to a Palo Alto Networks cloud-delivered service. When a license is within 30 days of expiration, a warning message displays in the system log daily until the subscription is renewed or expires. Upon license expiration, some subscriptions continue to function in a limited capacity, and others stop operating completely. Here you can find out what happens when each subscription expires.



The precise moment of license expiry is at the beginning of the following day at 12:00 AM (GMT). For example, if your license is scheduled to end on 1/20 you will have functionality for the remainder of that day. At the start of the new day on 1/21 at 12:00 AM (GMT), the license will expire. All license-related functions operate on Greenwich Mean Time (GMT), regardless of the configured time zone on the firewall.

Subscription	Expiry Behavior
Threat Prevention	Alerts appear in the System Log indicating that the license has expired. **You can still:** • Use signatures that were installed at the time the license expired, unless you install a new Applications-only content update either manually or as part of an automatic schedule. If you do, the update will delete your existing threat signatures and you will no longer receive protection against them. • Use and modify Custom App-ID™ and threat signatures. **You can no longer:** • Install new signatures. • Roll signatures back to previous versions.
DNS Security	**You can still:** • Use local DNS signatures if you have an active Threat Prevention license. **You can no longer:** • Get new DNS signatures.
Advanced URL Filtering / URL Filtering	**You can still:** • Enforce policy using custom URL categories. **You can no longer:** • Get updates to cached PAN-DB categories. • Connect to the PAN-DB URL filtering database. • Get PAN-DB URL categories.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Subscription	Expiry Behavior
	• Analyze URL requests in real-time using advanced URL filtering.
WildFire	**You can still:** • Forward PEs for analysis. • Get signature updates every 24-48 hours if you have an active Threat Prevention subscription. **You can no longer:** • Get five-minute updates through the WildFire public and private clouds. • Forward advanced file types such as APKs, Flash files, PDFs, Microsoft Office files, Java Applets, Java files (.jar and .class), and HTTP/HTTPS email links contained in SMTP and POP3 email messages. • Use the WildFire API. • Use the WildFire appliance to host a WildFire private cloud or a WildFire hybrid cloud.
AutoFocus	**You can still:** • Use an external dynamic list with AutoFocus data for a grace period of three months. **You can no longer:** • Access the AutoFocus portal. • View the AutoFocus Intelligence Summary for Monitor log or ACC artifacts.
Cortex Data Lake	**You can still:** • Store log data for a 30-day grace period, after which it is deleted. • Forward logs to Cortex Data Lake until the end of the 30-day grace period.
GlobalProtect	**You can still:** • Use the app for endpoints running Windows and macOS. • Configure single or multiple internal/external gateways. **You can no longer:** • Access the Linux OS app and mobile app for iOS, Android, Chrome OS, and Windows 10 UWP. • Use IPv6 for external gateways. • Run HIP checks.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Subscription	Expiry Behavior
	• Use Clientless VPN. • Enforce split tunneling based on destination domain, client process, and video streaming application.
VM-Series	**You can still:** • Configure and use the firewall(s) you had deployed when the license expired.
Support	**You can no longer:** • Receive software updates. • Download VM images. • Benefit from technical support.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Enhanced Application Logs for Palo Alto Networks Cloud Services

The firewall can collect data that increases visibility into network activity for Palo Alto Networks apps and services, like Cortex XDR. These enhanced application logs are designed strictly for Palo Alto Networks apps and services to consume and process; you cannot view enhanced application logs on the firewall or Panorama. Only firewalls sending logs to Cortex Data Lake can generate enhanced application logs.

Examples of the types of data that enhanced application logs gather includes records of DNS queries, the HTTP header User Agent field that specifies the web browser or tool used to access a URL, and information about DHCP automatic IP address assignment. With DHCP information, for example, Cortex XDR™ can alert on unusual activity based on hostname instead of IP address. This allows the security analyst using Cortex XDR to meaningfully assess whether the user's activity is within the scope of his or her role, and if not, to more quickly take action to stop the activity.

To benefit from the most comprehensive set of enhanced application logs, you should enable User-ID; deployments for the Windows-based User-ID agent and the PAN-OS integrated User-ID agent both collect some data that is not reflected in the firewall User-ID logs but that is useful towards associating network activity with specific users.

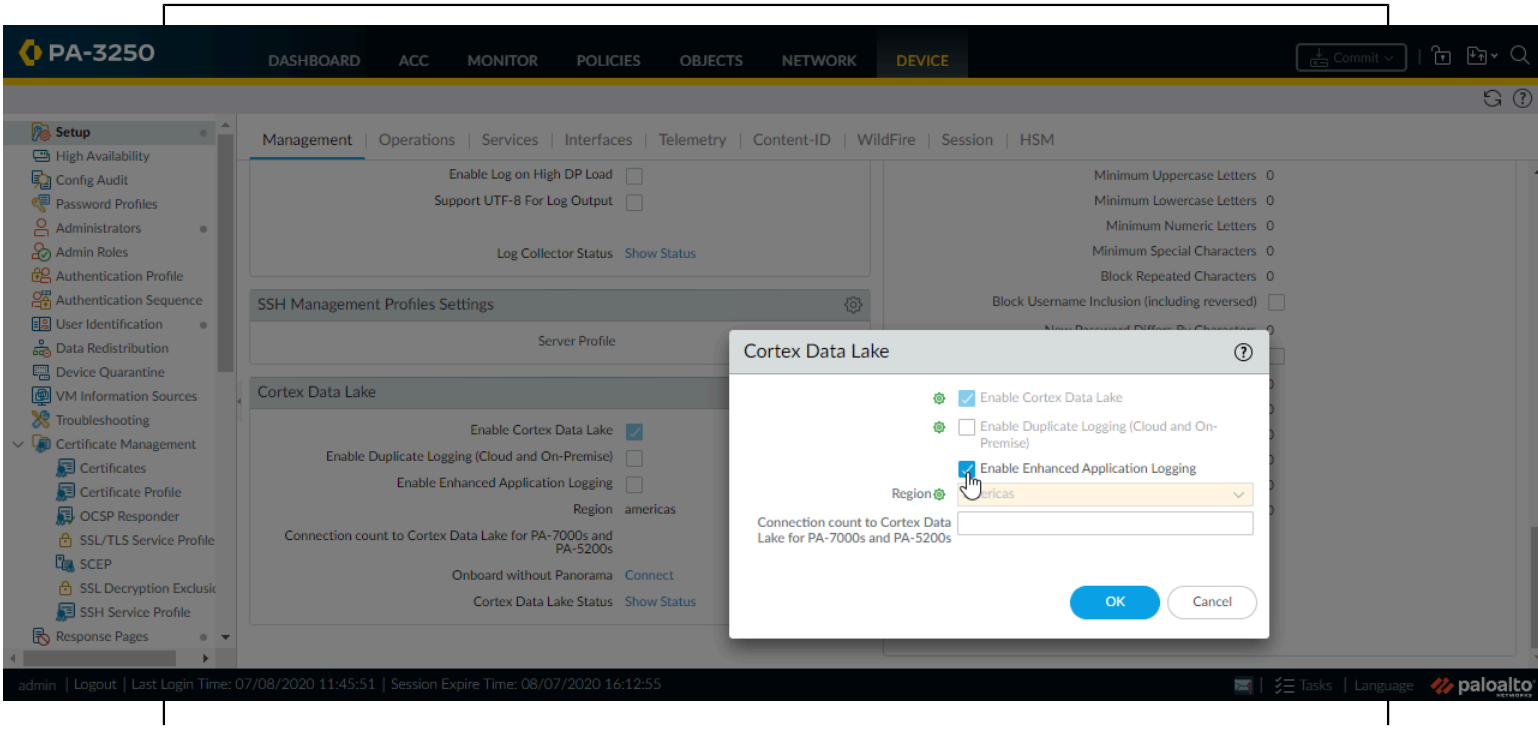
To start forwarding enhanced application logs to Cortex Data Lake, turn on enhanced application logging globally, and then enable it on a per-security rule basis (using a Log Forwarding profile). The global setting is required and captures data for traffic that is not session-based (ARP requests, for example). The per-security policy rule setting is strongly recommended; the majority of enhanced application logs are gathered from the session-based traffic that your security policy rules enforce.

STEP 1 | Enhanced application logging requires a Cortex Data Lake subscription and User-ID is also recommended. Here are steps to get started with Cortex Data Lake and enable User-ID.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 2 | To **Enable Enhanced Application Logging** on the firewall, select **Device** > **Setup** > **Management** > **Cortex Data Lake** and edit Cortex Data Lake Settings.

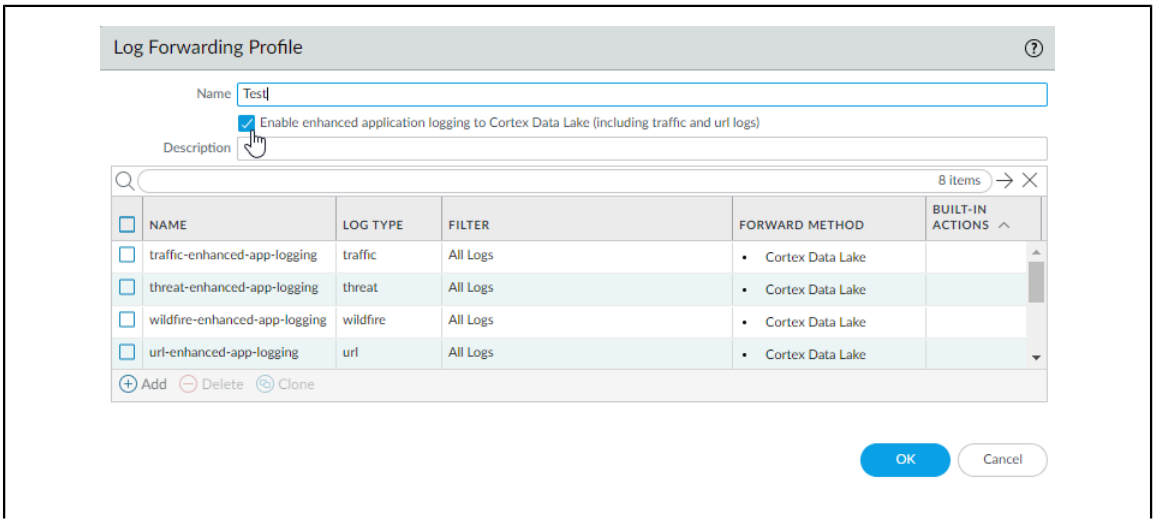


 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 3 | Continue to enable enhanced application logging for the security policy rules that control the traffic into which you want extended visibility.

1. Select **Objects** > **Log Forwarding** and **Add** or modify a log forwarding profile.
2. Update the profile to **Enable enhanced application logging to Cortex Data Lake (including traffic and url logs)**.



Notice that when you enable enhanced application logging in a Log Forwarding profile, match lists that specify the log types required for enhanced application logging are automatically added to the profile.

3. Click **OK** to save the profile and continue to update as many profiles as needed.
4. Ensure that the Log Forwarding profile that you've updated is attached to a security policy rule, to trigger log generation and forwarding for the traffic matched to the rule.
   1. Select **Policies** > **Security** to view the profiles attached to each security policy rule.
   2. To update the log forwarding profile attached to a rule, **Add** or edit a rule and select **Policies** > **Security** > **Actions** > **Log Forwarding** and select the Log Forwarding profile enabled with enhanced application logging.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Firewall Administration

Administrators can configure, manage, and monitor Palo Alto Networks firewalls using the web interface, CLI, and API management interface. You can customize role-based administrative access to the management interfaces to delegate specific tasks or permissions to certain administrators.

See Administrative Access Best Practices for how to safeguard your management network and the firewall and Panorama management interfaces.

- Management Interfaces
- Use the Web Interface
- Manage Configuration Backups
- Manage Firewall Administrators
- Reference: Web Interface Administrator Access
- Reference: Port Number Usage
- Reset the Firewall to Factory Default Settings
- Bootstrap the Firewall

https://t.me/learningnets

Management Interfaces

You can use the following user interfaces to manage the Palo Alto Networks firewall:



Do not enable management access from the internet or from other untrusted zones inside your enterprise security boundary. Follow the Adminstrative Access Best Practices to ensure that you are properly securing your firewall.

- Use the Web Interface to perform configuration and monitoring tasks with relative ease. This graphical interface allows you to access the firewall using HTTPS (recommended) or HTTP and it is the best way to perform administrative tasks.
- Use the Command Line Interface (CLI) to perform a series of tasks by entering commands in rapid succession over SSH (recommended), Telnet, or the console port. The CLI is a no-frills interface that supports two command modes, operational and configure, each with a distinct hierarchy of commands and statements. When you become familiar with the nesting structure and syntax of the commands, the CLI provides quick response times and administrative efficiency.
- Use the XML API to streamline your operations and integrate with existing, internally developed applications and repositories. The XML API is a web service implemented using HTTP/HTTPS requests and responses.
- Use Panorama to perform web-based management, reporting, and log collection for multiple firewalls. The Panorama web interface resembles the firewall web interface but with additional functions for centralized management.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Use the Web Interface

The following topics describe how to use the firewall web interface. For detailed information about specific tabs and fields in the web interface, refer to the Web Interface Reference Guide.

- Launch the Web Interface
- Configure Banners, Message of the Day, and Logos
- Use the Administrator Login Activity Indicators to Detect Account Misuse
- Manage and Monitor Administrative Tasks
- Commit, Validate, and Preview Firewall Configuration Changes
- Export Configuration Table Data
- Use Global Find to Search the Firewall or Panorama Management Server
- Manage Locks for Restricting Configuration Changes

Launch the Web Interface

The following web browsers are supported for access to the web interface:

- Internet Explorer 11+
- Firefox 3.6+
- Safari 5+
- Chrome 11+

Perform the following tasks to launch the web interface.

STEP 1 | Launch an Internet browser and enter the IP address of the firewall in the URL field (https://<IP address>).



*By default, the management (MGT) interface allows only HTTPS access to the web interface. To enable other protocols, select **Device** > **Setup** > **Interfaces** and edit the **Management** interface.*

STEP 2 | Log in to the firewall according to the type of authentication used for your account. If logging in to the firewall for the first time, use the default value **admin** for your username and password.

- **SAML**—Click **Use Single Sign-On** (SSO). If the firewall performs authorization (role assignment) for administrators, enter your **Username** and **Continue**. If the SAML identity provider (IdP) performs authorization, **Continue** without entering a **Username**. In both cases, the firewall redirects you to the IdP, which prompts you to enter a username and password. After you authenticate to the IdP, the firewall web interface displays.
- **Any other type of authentication**—Enter your user **Name** and **Password**. Read the login banner and select **I Accept and Acknowledge the Statement Below** if the login page has the banner and check box. Then click **Login**.

STEP 3 | Read and **Close** the messages of the day.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Configure Banners, Message of the Day, and Logos

A *login banner* is optional text that you can add to the login page so that administrators will see information they must know before they log in. For example, you could add a message to notify users of restrictions on unauthorized use of the firewall.

You can add colored bands that highlight overlaid text across the top (*header banner*) and bottom (*footer banner*) of the web interface to ensure administrators see critical information, such as the classification level for firewall administration.

A *message of the day* dialog automatically displays after you log in. The dialog displays messages that Palo Alto Networks embeds to highlight important information associated with a software or content release. You can also add one custom message to ensure administrators see information, such as an impending system restart, that might affect their tasks.

You can replace the default logos that appear on the login page and in the header of the web interface with the logos of your organization.

STEP 1 | Configure the login banner.

1. Select **Device** > **Setup** > **Management** and edit the General Settings.
2. Enter the **Login Banner** (up to 3,200 characters).
3. (Optional) Select **Force Admins to Acknowledge Login Banner** to force administrators to select an **I Accept and Acknowledge the Statement Below** check box above the banner text to activate the **Login** button.
4. Click **OK**.

STEP 2 | Set the message of the day.

1. Select **Device** > **Setup** > **Management** and edit the Banners and Messages settings.
2. Enable the **Message of the Day**.
3. Enter the **Message of the Day** (up to 3,200 characters).

   *After you enter the message and click **OK**, administrators who subsequently log in, and active administrators who refresh their browsers, see the new or updated message immediately; a commit isn't necessary. This enables you to inform other administrators of an impending commit that might affect their configuration changes. Based on the commit time that your message specifies, the administrators can then decide whether to complete, save, or undo their changes.*
4. (Optional) Select **Allow Do Not Display Again** (default is disabled) to give administrators the option to suppress a message of the day after the first login session. Each administrator can suppress messages only for his or her own login sessions. In the message of the day dialog, each message will have its own suppression option.
5. (Optional) Enter a header **Title** for the message of the day dialog (default is `Messageof the Day`).

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 3 | Configure the header and footer banners.

A bright background color and contrasting text color can increase the likelihood that administrators will notice and read a banner. You can also use colors that correspond to classification levels in your organization.

1. Enter the **Header Banner** (up to 3,200 characters).
2. (Optional) Clear **Same Banner Header and Footer** (enabled by default) to use different header and footer banners.
3. Enter the **Footer Banner** (up to 3,200 characters) if the header and footer banners differ.
4. Click **OK**.

STEP 4 | Replace the logos on the login page and in the header.

The maximum size for any logo image is 128KB. The supported file types are png and jpg. The firewall does not support image files that are interlaced, images that contain alpha channels, and gif file types because such files interfere with PDF generation.

1. Select **Device** > **Setup** > **Operations** and click **Custom Logos** in the Miscellaneous section.
2. Perform the following steps for both the **Login Screen** logo and the **Main UI** (header) logo:
   1. Click upload .
   2. Select a logo image and click **Open**.

      You can preview the image to see how PAN-OS will crop it to fit by clicking the magnifying glass icon.
   3. Click **Close**.
3. **Commit** your changes.

STEP 5 | Verify that the banners, message of the day, and logos display as expected.

1. Log out to return to the login page, which displays the new logos you selected.
2. Enter your login credentials, review the banner, select **I Accept and Acknowledge the Statement Below** to enable the **Login** button, and then **Login**.

   A dialog displays the message of the day. Messages that Palo Alto Networks embedded display on separate pages in the same dialog. To navigate the pages, click the right or left arrows along the sides of the dialog or click a page selector at the bottom of the dialog.
3. (Optional) You can select **Do not show again** for the message you configured and for any messages that Palo Alto Networks embedded.
4. **Close** the message of the day dialog to access the web interface.

   Header and footer banners display in every web interface page with the text and colors that you configured. The new logo you selected for the web interface displays below the header banner.

 ©2022 Palo Alto Networks, Inc.

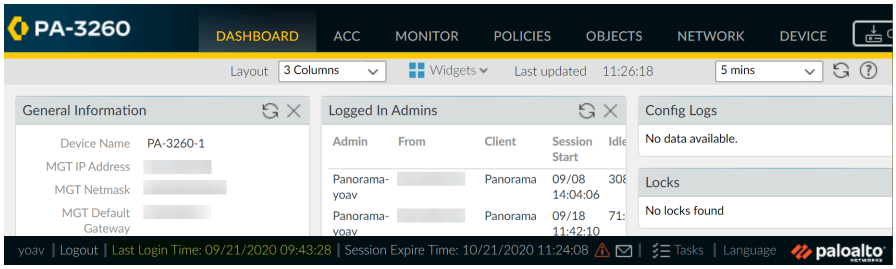
https://t.me/learningnets

Use the Administrator Login Activity Indicators to Detect Account Misuse

The last login time and failed login attempts indicators provide a visual way to detect misuse of your administrator account on a Palo Alto Networks firewall or Panorama management server. Use the last login information to determine if someone else logged in using your credentials and use the failed login attempts indicator to determine if your account is being targeted in a brute-force attack.

STEP 1 | View the login activity indicators to monitor recent activity on your account.

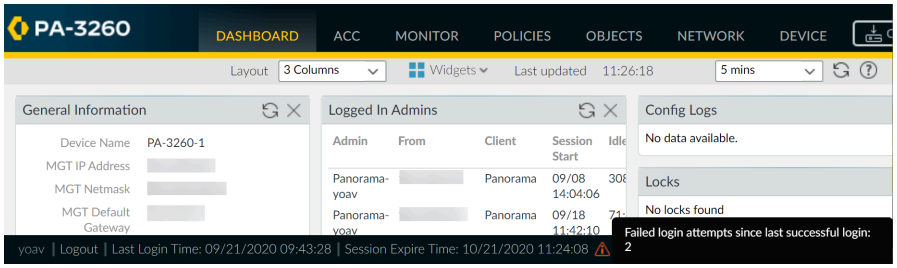
1. Log in to the web interface on your firewall or Panorama management server.
2. View the last login details located at the bottom left of the window and verify that the timestamp corresponds to your last login.



3. Look for a caution symbol to the right of the last login time information for failed login attempts.

   The failed login indicator appears if one or more failed login attempts occurred using your account since the last successful login.

   1. If you see the caution symbol, hover over it to display the number of failed login attempts.



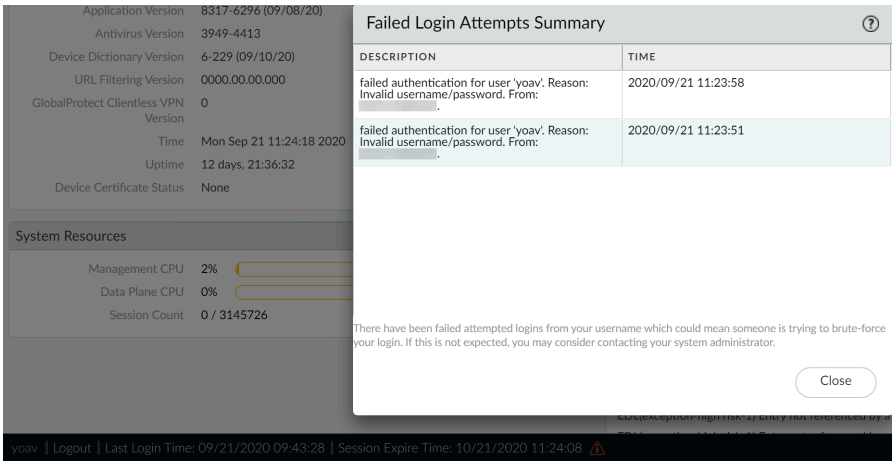
   2. Click the caution symbol to view the failed login attempts summary. Details include the admin account name, the reason for the login failure, the source IP address, and the date and time.

> *After you successfully log in and then log out, the failed login counter resets to zero so you will see new failed login details, if any, the next time you log in.*

https://t.me/learningnets

STEP 2 | Locate hosts that are continually attempting to log in to your firewall or Panorama management server.

1. Click the failed login caution symbol to view the failed login attempts summary.
2. Locate and record the source IP address of the host that attempted to log in. For example, the following figure shows multiple failed login attempts.



3. Work with your network administrator to locate the user and host that is using the IP address that you identified.

   If you cannot locate the system that is performing the brute-force attack, consider renaming the account to prevent future attacks.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 3 | Take the following actions if you detect an account compromise.

1. Select **Monitor** > **Logs** > **Configuration** and view the configuration changes and commit history to determine if your account was used to make changes without your knowledge.
2. Select **Device** > **Config Audit** to compare the current configuration and the configuration that was running just prior to the configuration you suspect was changed using your credentials. You can also do this using Panorama.

   If your administrator account was used to create a new account, performing a configuration audit helps you detect changes that are associated with any unauthorized accounts, as well.
3. Revert the configuration to a known good configuration if you see that logs were deleted or if you have difficulty determining if improper changes were made using your account.

   Before you commit to a previous configuration, review it to ensure that it contains the correct settings. For example, the configuration that you revert to may not contain recent changes, so apply those changes after you commit the backup configuration.

   Use the following best practices to help prevent brute-force attacks on privileged accounts.

   - *Limit the number of failed attempts allowed before the firewall locks a privileged account by setting the number of Failed Attempts and the Lockout Time (min) in the authentication profile or in the Authentication Settings for the Management interface (**Device** > **Setup** > **Management** > **Authentication Settings**).*
   - Use Interface Management Profiles to Restrict Access.
   - *Enforce* complex passwords *for privileged accounts.*

Manage and Monitor Administrative Tasks

The Task Manager displays details about all the operations that you and other administrators initiated (such as manual commits) or that the firewall initiated (such as scheduled report generation) since the last firewall reboot. You can use the Task Manager to troubleshoot failed operations, investigate warnings associated with completed commits, view details about queued commits, or cancel pending commits.

You can also view System Logs *to monitor system events on the firewall or view* Config Logs *to monitor firewall configuration changes.*

STEP 1 | Click **Tasks** at the bottom of the web interface.

STEP 2 | **Show** only **Running** tasks (in progress) or **All** tasks (default). Optionally, filter the tasks by type:

- **Jobs**—Administrator-initiated commits, firewall-initiated commits, and software or content downloads and installations.
- **Reports**—Scheduled reports.
- **Log Requests**—Log queries that you trigger by accessing the **Dashboard** or a **Monitor** page.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 3 | Perform any of the following actions:

- **Display or hide task details**—By default, the Task Manager displays the Type, Status, Start Time, and Messages for each task. To see the End Time and Job ID for a task, you must manually configure the display to expose those columns. To display or hide a column, open the drop-down in any column header, select **Columns**, and select or deselect the column names as needed.
- **Investigate warnings or failures**—Read the entries in the Messages column for task details. If the column says `Too many messages`, click the corresponding entry in the Type column to see more information.
- **Display a commit description**—If an administrator entered a description when configuring a commit, you can click **Commit Description** in the Messages column to display the description.
- **Check the position of a commit in the queue**—The Messages column indicates the queue position of commits that are in progress.
- **Cancel pending commits**—Click **Clear Commit Queue** to cancel all pending commits (available only to predefined administrative roles). To cancel an individual commit, click **x** in the Action column for that commit (the commit remains in the queue until the firewall dequeues it). You cannot cancel commits that are in progress.

Commit, Validate, and Preview Firewall Configuration Changes

A commit is the process of activating pending changes to the firewall configuration. You can filter pending changes by administrator or *location* and then preview, validate, or commit only those changes. The locations can be specific virtual systems, shared policies and objects, or shared device and network settings.

The firewall queues commit requests so that you can initiate a new commit while a previous commit is in progress. The firewall performs the commits in the order they are initiated but prioritizes auto-commits that are initiated by the firewall (such as FQDN refreshes). However, if the queue already has the maximum number of administrator-initiated commits, you must wait for the firewall to finish processing a pending commit before initiating a new one. To cancel pending commits or view details about commits of any status, see Manage and Monitor Administrative Tasks.

When you initiate a commit, the firewall checks the validity of the changes before activating them. The validation output displays conditions that either block the commit (errors) or that are important to know (warnings). For example, validation could indicate an invalid route destination that you need to fix for the commit to succeed. The validation process enables you to find and fix errors before you commit (it makes no changes to the running configuration). This is useful if you have a fixed commit window and want to be sure the commit will succeed without errors.

When enabled and managed by a Panorama™ management server, managed firewalls locally test the configuration committed locally or pushed from Panorama to verify that the new changes do not break the connection between Panorama and the managed firewall. If the committed configuration breaks the connection between Panorama and a managed firewall, then the firewall automatically fails the commit and the configuration is reverted to the previous running configuration. Additionally, firewalls managed by a Panorama management server test their connection to Panorama every 60 minutes and if a managed firewalls detects that it can no

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

longer successfully connect to Panorama, then it reverts its configuration to the previous running configuration.

The commit, validate, preview, save, and revert operations apply only to changes made after the last commit. To restore configurations to the state they were in before the last commit, you must load a previously backed up configuration.

To prevent multiple administrators from making configuration changes during concurrent sessions, see Manage Locks for Restricting Configuration Changes.

STEP 1 | Configure the scope of configuration changes that you will commit, validate, or preview.

1. Click **Commit** at the top of the web interface.
2. Select one of the following options:
   - **Commit All Changes** (default)—Applies the commit to all changes for which you have administrative privileges. You cannot manually filter the commit scope when you select this option. Instead, the administrator role assigned to the account you used to log in determines the commit scope.
   - **Commit Changes Made By**—Enables you to filter the commit scope by administrator or location. The administrative role assigned to the account you used to log in determines which changes you can filter.

   To commit the changes of other administrators, the account you used to log in must be assigned the Superuser role or an Admin Role profile *with the* ***Commit For Other Admins*** *privilege enabled.*
3. (Optional) To filter the commit scope by administrator, select **Commit Changes Made By**, click the adjacent link, select the administrators, and click **OK**.
4. (Optional) To filter by location, select **Commit Changes Made By** and clear any changes that you want to exclude from the Commit Scope.

   If dependencies between the configuration changes you included and excluded cause a validation error, perform the commit with all the changes included. For example, when you commit changes to a virtual system, you must include the changes of all administrators who added, deleted, or repositioned rules for the same rulebase in that virtual system.

STEP 2 | Preview the changes that the commit will activate.

This can be useful if, for example, you don't remember all your changes and you're not sure you want to activate all of them.

The firewall compares the configurations you selected in the Commit Scope to the running configuration. The preview window displays the configurations side-by-side and uses color coding to indicate which changes are additions (green), modifications (yellow), or deletions (red).

Preview Changes and select the **Lines of Context**, which is the number of lines from the compared configuration files to display before and after each highlighted difference. These

https://t.me/learningnets

additional lines help you correlate the preview output to settings in the web interface. Close the preview window when you finish reviewing the changes.

Because the preview results display in a new browser window, your browser must allow pop-ups. If the preview window does not open, refer to your browser documentation for the steps to allow pop-ups.

STEP 3 | Preview the individual settings for which you are committing changes.

This can be useful if you want to know details about the changes, such as the types of settings and who changed them.

1. Click **Change Summary**.
2. (Optional) **Group By** a column name (such as the **Type** of setting).
3. **Close** the Change Summary dialog when you finish reviewing the changes.

STEP 4 | Validate the changes before you commit to ensure the commit will succeed.

1. **Validate Changes**.

   The results display all the errors and warnings that an actual commit would display.
2. Resolve any errors that the validation results identify.

STEP 5 | Commit your configuration changes.

Commit your changes to validate and activate them.

To view details about commits that are pending (which you can still cancel), in progress, completed, or failed, see Manage and Monitor Administrative Tasks.

Export Configuration Table Data

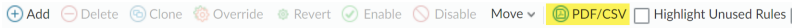
Export policy rules, configuration objects, and IPS signatures from Panorama™ and firewalls to demonstrate regulatory compliance to external auditors, to conduct periodic reviews of the firewall configuration, and to generate reports on firewall policies. This prevents you from having to give auditors direct access to your firewalls and appliances, to take screen shots or to access the XML API to generate configuration reports. From the web interface, you can export the configuration table data for policies, objects, network, firewall, and Panorama configurations, as well as Signature exceptions in the Antivirus, Anti-Spyware, and Vulnerability Protection Security profiles, in either a PDF or CSV file.

Configuration table export works like a print function—you cannot import generated files back into Panorama or the firewall. When you export data as a PDF file and the table data exceeds 50,000 rows, the data is split in to multiple PDF files (for example, <report-name>_part1.pdf and <report-name>_part2.pdf) When you export data as a CSV file, the data is exported as a single file. These export formats allow you to apply filters that match your report criteria and search within PDF reports to quickly find specific data. Additionally, when you export the configuration table data, a system log is generated to record the event.

STEP 1 | Launch the Web Interface and identify the configuration data you need to export.

https://t.me/learningnets

STEP 2 | Apply filters as needed to produce the configuration data you need to export and click **PDF/CSV**.



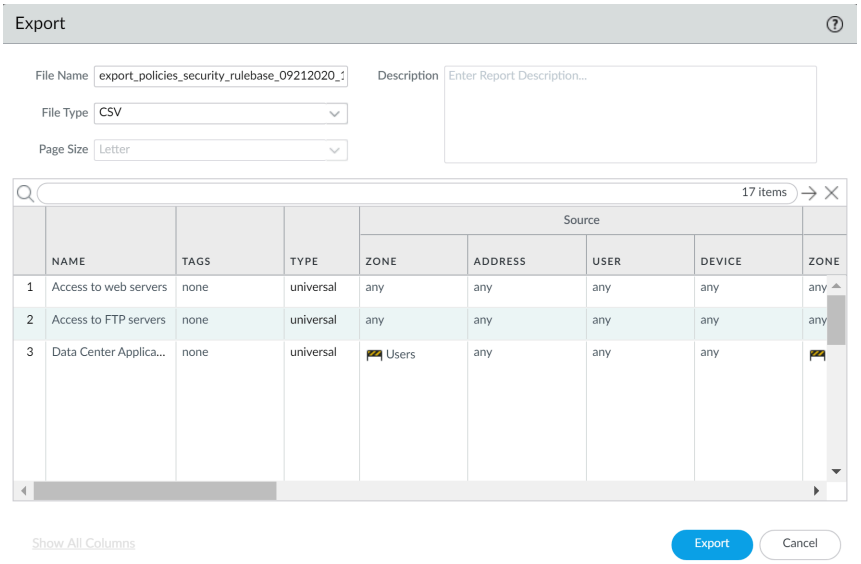
STEP 3 | Configure the Configuration Table Export report:

1. Enter a **File Name**.
2. Select the **File Type**.
3. (Optional) Enter a report Description.
4. Confirm the configuration table data matches the filters you applied.

*Select **Show All Columns** to show all filters applied.*

STEP 4 | **Export** the configuration table data.

Configuration table export works like a print function—you cannot import generated files back in to Panorama or the firewall.



STEP 5 | Select a location to save the exported file.

Use Global Find to Search the Firewall or Panorama Management Server

Global Find enables you to search the candidate configuration on a firewall or on Panorama for a particular string, such as an IP address, object name, policy rule name, threat ID, UUID, or application name. In addition to searching for configuration objects and settings, you can search by job ID or job type for manual commits that administrators performed or auto-commits that the firewall or Panorama performed. The search results are grouped by category and provide links to the configuration location in the web interface, so that you can easily find all of the places where the string is referenced. The search results also help you identify other objects that depend on or make reference to the search term or string. For example, when deprecating a security profile enter the profile name in Global Find to locate all instances of the profile and then click each

https://t.me/learningnets

instance to navigate to the configuration page and make the necessary change. After all references are removed, you can then delete the profile. You can do this for any configuration item that has dependencies.

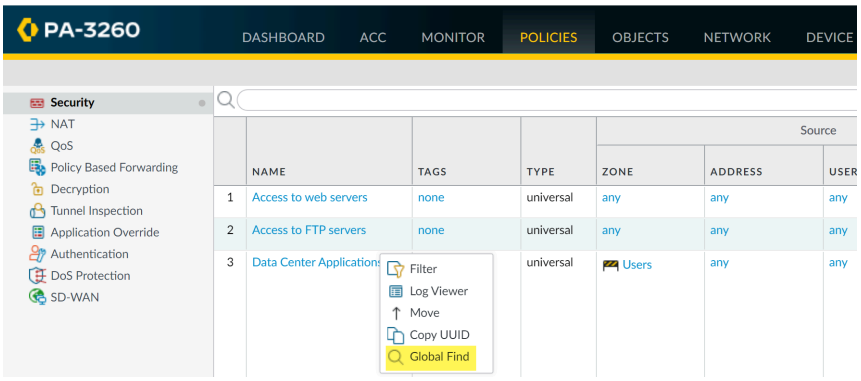
Watch the video.

Global Find does not search dynamic content (such as logs, address ranges, or allocated DHCP addresses). In the case of DHCP, you can search on a DHCP server attribute, such as the DNS entry, but you cannot search for individual addresses allocated to users. Global Find also does not search for individual user or group names identified by User-ID unless the user/group is defined in a policy. In general, you can only search content that the firewall writes to the configuration.

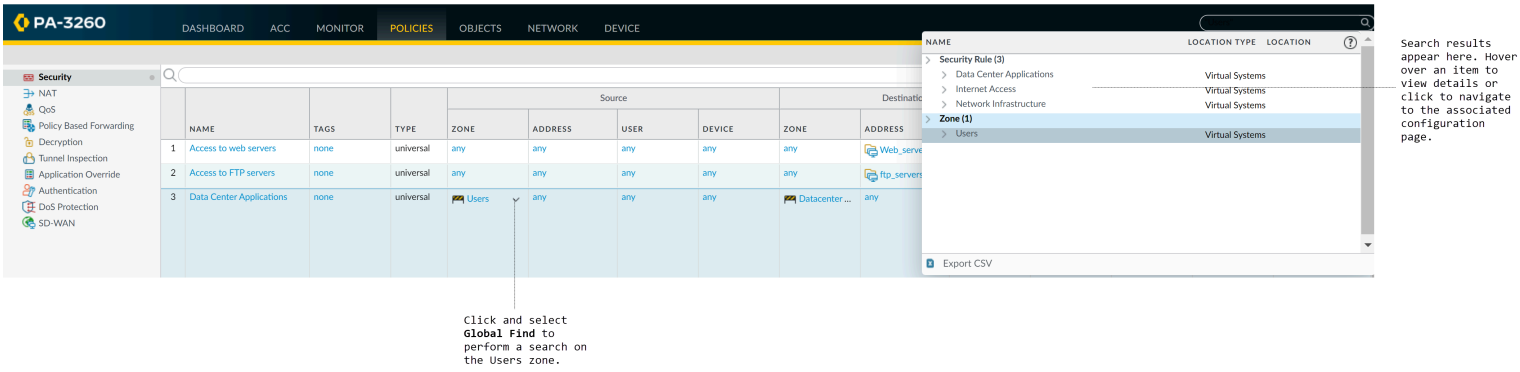
- Launch Global Find by clicking the **Search** icon located on the upper right of the web interface.



- To access the Global Find from within a configuration area, click the drop-down next to an item and select **Global Find**:



For example, click **Global Find** on a zone named **Users** to search the candidate configuration for each location where the zone is referenced. The following screen capture shows the search results for the zone Users:



Search tips:

- If you initiate a search on a firewall that has multiple virtual systems enabled or if custom Administrative Role Types are defined, Global Find will only return results for areas of the

https://t.me/learningnets

firewall in which the administrator has permissions. The same applies to Panorama device groups.

- Spaces in search terms are handled as AND operations. For example, if you search on **`corp policy`**, the search results include instances where corp and policy exist in the configuration.
- To find an exact phrase, enclose the phrase in quotation marks.
- Enter no more than five keywords or use an exact phrase match with quotation marks.
- To rerun a previous search, click Search (located on the upper right of the web interface) to see a list of the last 20 searches. Click an item in the list to rerun that search. Search history is unique to each administrator account.
- To search for a UUID, you must copy and paste the UUID.

Manage Locks for Restricting Configuration Changes

You can use configuration locks to prevent other administrators from changing the candidate configuration or from committing configuration changes until you manually remove the lock or the firewall automatically removes it (after a commit). Locks ensure that administrators don't make conflicting changes to the same settings or interdependent settings during concurrent login sessions.

The firewall queues commit requests and performs them in the order that administrators initiate the commits. For details, see Commit, Validate, and Preview Firewall Configuration Changes. *To view the status of queued commits, see* Manage and Monitor Administrative Tasks.

- View details about current locks.

  For example, you can check whether other administrators have set locks and read comments they entered to explain the locks.

  Click the lock at the top of the web interface. An adjacent number indicates the number of current locks.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

- Lock a configuration.
  1. Click the lock at the top of the web interface.

     The lock image varies based on whether existing locks are or are not set.
  2. **Take a Lock** and select the lock **Type**:
     - **Config**—Blocks other administrators from changing the candidate configuration.
     - **Commit**—Blocks other administrators from committing changes made to the candidate configuration.
  3. (Firewall with multiple virtual systems only) Select a **Location** to lock the configuration for a specific virtual system or the **Shared** location.
  4. (Optional) As a best practice, enter a **Comment** so that other administrators will understand the reason for the lock.
  5. Click **OK** and **Close**.

- Unlock a configuration.

  Only a superuser or the administrator who locked the configuration can manually unlock it. However, the firewall automatically removes a lock after completing the commit operation.
  1. Click the lock at the top of the web interface.
  2. Select the lock entry in the list.
  3. Click **Remove Lock**, **OK**, and **Close**.

- Configure the firewall to automatically apply a commit lock when you change the candidate configuration. This setting applies to all administrators.
  1. Select **Device** > **Setup** > **Management** and edit the General Settings.
  2. Select **Automatically Acquire Commit Lock** and then click **OK** and **Commit**.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Manage Configuration Backups

The running configuration on the firewall comprises all settings you have committed and that are therefore active, such as policy rules that currently block or allow various types of traffic in your network. The candidate configuration is a copy of the running configuration plus any inactive changes that you made after the last commit. Saving backup versions of the running or candidate configuration enables you to later restore those versions. For example, if a commit validation shows that the current candidate configuration has more errors than you want to fix, you can restore a previous candidate configuration. You can also revert to the current running configuration without saving a backup first. If you need to export specific parts of the configuration for internal review or audit, you can Export Configuration Table Data.



See Commit, Validate, and Preview Firewall Configuration Changes for details about commit operations.

- Save and Export Firewall Configurations
- Revert Firewall Configuration Changes

Save and Export Firewall Configurations

Saving a backup of the candidate configuration to persistent storage on the firewall enables you to later revert to that backup (see Revert Firewall Configuration Changes). This is useful for preserving changes that would otherwise be lost if a system event or administrator action causes the firewall to reboot. After rebooting, PAN-OS automatically reverts to the current version of the running configuration, which the firewall stores in a file named running-config.xml. Saving backups is also useful if you want to revert to a firewall configuration that is earlier than the current running configuration. The firewall does not automatically save the candidate configuration to persistent storage. You must manually save the candidate configuration as a default snapshot file (.snapshot.xml) or as a custom-named snapshot file. The firewall stores the snapshot file locally but you can export it to an external host.



*You don't have to save a configuration backup to revert the changes made since the last commit or reboot; just select **Config** > **Revert Changes** (see Revert Firewall Configuration Changes).*

*When you edit a setting and click **OK**, the firewall updates the candidate configuration but does not save a backup snapshot.*

Additionally, saving changes does not activate them. To activate changes, perform a commit (see Commit, Validate, and Preview Firewall Configuration Changes).

Palo Alto Networks recommends that you back up any important configuration to a host external to the firewall.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 1 | Save a local backup snapshot of the candidate configuration if it contains changes that you want to preserve in the event the firewall reboots.

These are changes you are not ready to commit—for example, changes you cannot finish in the current login session.

To overwrite the default snapshot file (.snapshot.xml) with all the changes that all administrators made, perform one of the following steps:

- Select **Device** > **Setup** > **Operations** and **Save candidate configuration**.
- Log in to the firewall with an administrative account that is assigned the Superuser role or an Admin Role profile with the **Save For Other Admins** privilege enabled. Then select **Config** > **Save Changes** at the top of the web interface, select **Save All Changes** and **Save**.

To create a snapshot that includes all the changes that all administrators made but without overwriting the default snapshot file:

1. Select **Device** > **Setup** > **Operations** and **Save named configuration snapshot**.
2. Specify the **Name** of a new or existing configuration file.
3. Click **OK** and **Close**.

To save only specific changes to the candidate configuration without overwriting any part of the default snapshot file:

1. Log in to the firewall with an administrative account that has the role privileges required to save the desired changes.
2. Select **Config** > **Save Changes** at the top of the web interface.
3. Select **Save Changes Made By**.
4. To filter the Save Scope by administrator, click **<administrator-name>**, select the administrators, and click **OK**.
5. To filter the Save Scope by location, clear any locations that you want to exclude. The locations can be specific virtual systems, shared policies and objects, or shared device and network settings.
6. Click **Save**, specify the **Name** of a new or existing configuration file, and click **OK**.

STEP 2 | Export a candidate configuration, a running configuration, or the firewall state information to a host external to the firewall.

Select **Device** > **Setup** > **Operations** and click an export option:

- **Export named configuration snapshot**—Export the current running configuration, a named candidate configuration snapshot, or a previously imported configuration (candidate or running). The firewall exports the configuration as an XML file with the **Name** you specify.
- **Export configuration version**—Select a **Version** of the running configuration to export as an XML file. The firewall creates a version whenever you commit configuration changes.
- **Export device state**—Export the firewall state information as a bundle. Besides the running configuration, the state information includes device group and template settings pushed from Panorama. If the firewall is a GlobalProtect portal, the information also includes certificate information, a list of satellites, and satellite authentication information. If you replace a firewall or portal, you can restore the exported information on the replacement by importing the state bundle.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Revert Firewall Configuration Changes

Revert operations replace settings in the current candidate configuration with settings from another configuration. Reverting changes is useful when you want to undo changes to multiple settings as a single operation instead of manually reconfiguring each setting.

You can revert pending changes that were made to the firewall configuration since the last commit. The firewall provides the option to filter the pending changes by administrator or *location*. The locations can be specific virtual systems, shared policies and objects, or shared device and network settings. If you saved a snapshot file for a candidate configuration that is earlier than the current running configuration (see Save and Export Firewall Configurations), you can also revert to that snapshot. Reverting to a snapshot enables you to restore a candidate configuration that existed before the last commit. The firewall automatically saves a new version of the running configuration whenever you commit changes, and you can restore any of those versions.

- Revert to the current running configuration (file named running-config.xml).

  This operation undoes changes you made to the candidate configuration since the last commit.

  To revert all the changes that all administrators made, perform one of the following steps:

  - Select **Device** > **Setup** > **Operations, Revert to running configuration**, and click **Yes** to confirm the operation.
  - Log in to the firewall with an administrative account that is assigned the Superuser role or an Admin Role profile with the **Commit For Other Admins** privilege enabled. Then select **Config** > **Revert Changes** at the top of the web interface, select **Revert All Changes** and **Revert**.

  To revert only specific changes to the candidate configuration:

  1. Log in to the firewall with an administrative account that has the role privileges required to revert the desired changes.

     The privileges that control commit operations also control revert operations.

  2. Select **Config** > **Revert Changes** at the top of the web interface.
  3. Select **Revert Changes Made By**.
  4. To filter the Revert Scope by administrator, click **<administrator-name>**, select the administrators, and click **OK**.
  5. To filter the Revert Scope by location, clear any locations that you want to exclude.
  6. **Revert** the changes.

- Revert to the default snapshot of the candidate configuration.

  This is the snapshot that you create or overwrite when you click **Config** > **Save Changes** at the top of the web interface.

  1. Select **Device** > **Setup** > **Operations** and **Revert to last saved configuration**.
  2. Click **Yes** to confirm the operation.
  3. (Optional) Click **Commit** to overwrite the running configuration with the snapshot.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

- Revert to a previous version of the running configuration that is stored on the firewall.

  The firewall creates a version whenever you commit configuration changes.

  1. Select **Device** > **Setup** > **Operations** and **Load configuration version**.
  2. Select a configuration **Version** and click **OK**.
  3. (Optional) Click **Commit** to overwrite the running configuration with the version you just restored.

- Revert to one of the following:

  - Custom-named version of the running configuration that you previously imported
  - Custom-named candidate configuration snapshot (instead of the default snapshot)

  1. Select **Device** > **Setup** > **Operations** and click **Load named configuration snapshot**.
  2. Select the snapshot **Name** and click **OK**.
  3. (Optional) Click **Commit** to overwrite the running configuration with the snapshot.

- Revert to a running or candidate configuration that you previously exported to an external host.

  1. Select **Device** > **Setup** > **Operations**, click **Import named configuration snapshot**, **Browse** to the configuration file on the external host, and click **OK**.
  2. Click **Load named configuration snapshot**, select the **Name** of the configuration file you just imported, and click **OK**.
  3. (Optional) Click **Commit** to overwrite the running configuration with the snapshot you just imported.

- Restore state information that you exported from a firewall.

  Besides the running configuration, the state information includes device group and template settings pushed from Panorama. If the firewall is a GlobalProtect portal, the information also includes certificate information, a list of satellites, and satellite authentication information. If you replace a firewall or portal, can you can restore the information on the replacement by importing the state bundle.

  Import state information:

  1. Select **Device** > **Setup** > **Operations**, click **Import device state**, **Browse** to the state bundle, and click **OK**.
  2. (Optional) Click **Commit** to apply the imported state information to the running configuration.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Manage Firewall Administrators

Administrative accounts specify roles and authentication methods for the administrators of Palo Alto Networks firewalls. Every Palo Alto Networks firewall has a predefined default administrative account (admin) that provides full read-write access (also known as superuser access) to the firewall.

As a best practice, create a separate administrative account for each person who needs access to the administrative or reporting functions of the firewall. This enables you to better protect the firewall from unauthorized configuration and enables logging of the actions of individual administrators. Make sure you are following the Adminstrative Access Best Practices to ensure that you are securing administrative access to your firewalls and other security devices in a way that prevents successful attacks.

- Administrative Role Types
- Configure an Admin Role Profile
- Administrative Authentication
- Configure Administrative Accounts and Authentication
- Configure Tracking of Administrator Activity

Administrative Role Types

A *role* defines the type of access that an administrator has to the firewall. The Administrator Types are:

- **Role Based**—Custom roles you can configure for more granular access control over the functional areas of the web interface, CLI, and XML API. For example, you can create an Admin Role profile for your operations staff that provides access to the firewall and network configuration areas of the web interface and a separate profile for your security administrators that provides access to security policy definitions, logs, and reports. On a firewall with multiple virtual systems, you can select whether the role defines access for all virtual systems or specific virtual systems. When new features are added to the product, you must update the roles with corresponding access privileges: the firewall does not automatically add new features to custom role definitions. For details on the privileges you can configure for custom administrator roles, see Reference: Web Interface Administrator Access.
- **Dynamic**—Built-in roles that provide access to the firewall. When new features are added, the firewall automatically updates the definitions of dynamic roles; you never need to manually update them. The following table lists the access privileges associated with dynamic roles.

Dynamic Role	Privileges
Superuser	Full access to the firewall, including defining new administrator accounts and virtual systems. You must have Superuser privileges to create an administrative user with Superuser privileges.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Dynamic Role	Privileges
Superuser (read-only)	Read-only access to the firewall.
Device administrator	Full access to all firewall settings except for defining new accounts or virtual systems.
Device administrator (read-only)	Read-only access to all firewall settings except password profiles (no access) and administrator accounts (only the logged in account is visible).
Virtual system administrator	Access to selected virtual systems on the firewall to create and manage specific aspects of virtual systems. A virtual system administrator doesn't have access to network interfaces, VLANs, virtual wires, virtual routers, IPSec tunnels, GRE tunnels, DHCP, DNS Proxy, QoS, LLDP, or network profiles.
Virtual system administrator (read-only)	Read-only access to selected virtual systems on the firewall and specific aspects of virtual systems. A virtual system administrator with read-only access doesn't have access to network interfaces, VLANs, virtual wires, virtual routers, IPSec tunnels, GRE tunnels, DHCP, DNS Proxy, QoS, LLDP, or network profiles.

Configure an Admin Role Profile

Admin Role profiles enable you to define granular administrative access privileges to ensure protection for sensitive company information and privacy for end users.

Follow the principle of least privilege access to create Admin Role profiles that enable administrators to access only the areas of the management interface that they need to access to perform their jobs and follow Administrative Access Best Practices.

STEP 1 | Select **Device** > **Admin Roles** and click **Add**.

STEP 2 | Enter a **Name** to identify the role.

STEP 3 | For the scope of the **Role**, select **Device** or **Virtual System**.

STEP 4 | In the **Web UI** and **REST API** tabs, click the icon for each functional area to toggle it to the desired setting: Enable, Read Only or Disable. For the **XML API** tab select, Enable or Disable. For details on the **Web UI** options, see Web Interface Access Privileges.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 5 | Select the **Command Line** tab and select a CLI access option. The **Role** scope controls the available options:

- **Device** role:
  - **None**—CLI access is not permitted (default).
  - **superuser**—Full access. Can define new administrator accounts and virtual systems. Only a superuser can create administrator users with superuser privileges.
  - **superreader**—Full read-only access.
  - **deviceadmin**—Full access to all settings except defining new accounts or virtual systems.
  - **devicereader**—Read-only access to all settings except password profiles (no access) and administrator accounts (only the logged in account is visible).
- **Virtual System** role:
  - **None**—Access is not permitted (default).
  - **vsysadmin**—Access to specific virtual systems to create and manage specific aspects of virtual systems. Does not enable access to firewall-level or network-level functions including static and dynamic routing, interface IP addresses, IPSec tunnels, VLANs, virtual wires, virtual routers, GRE tunnels, DCHP, DNS Proxy, QoS, LLDP, or network profiles.
  - **vsysreader**—Read-only access to specific virtual systems to specific aspects of virtual systems. Does not enable access to firewall-level or network-level functions including static and dynamic routing, interface IP addresses, IPSec tunnels, VLANs, virtual wires, virtual routers, GRE tunnels, DCHP, DNS Proxy, QoS, LLDP, or network profiles.

STEP 6 | Click **OK** to save the profile.

STEP 7 | Assign the role to an administrator. See Configure a Firewall Administrator Account.

Example Admin Role Profile Construction

This example shows an Admin Role profile for a Security Operations Center (SOC) manager who needs access to investigate potential issues. The SOC Manager needs read access to many areas of the firewall, but generally doesn't need write access. The example covers all four of the Admin Role Profile's tabs and each step describes why the profile enables or disables a particular area of access to the SOC manager.

This is an example profile for a fictional SOC manager. Configure Admin Role profiles for your administrators based on the functions they manage and the access required to do their job. Do not enable unnecessary access. Create separate profiles for each administrative group that shares the same duties and for administrators who have unique duties. Each administrator should have the exact level of access required to perform their duties and no access beyond that.

STEP 1 | Configure Web UI access permissions. Each snip of the Web UI screen shows a different area of Web UI permissions. Permissions are listed by firewall tab, in the order you see the tabs in the Web UI, followed by permissions for other actions.

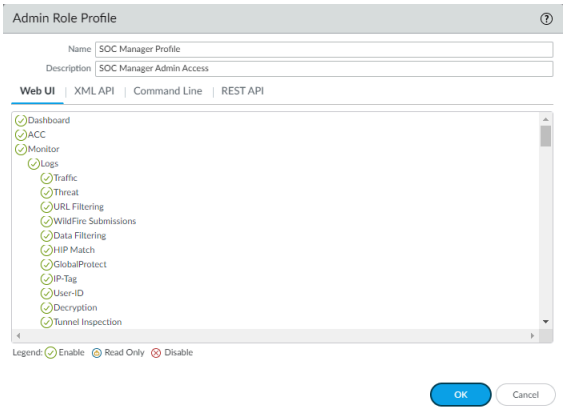
The **Dashboard**, **ACC**, and **Monitor** > **Logs** areas of the firewall don't contain configuration elements—all of the objects are informational (you can only toggle them between enable and

 ©2022 Palo Alto Networks, Inc.

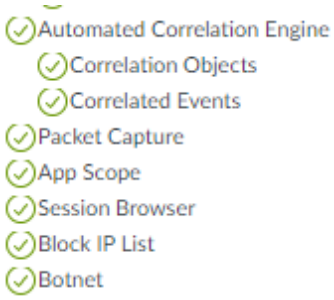
https://t.me/learningnets

disable because they are already read only). Because the SOC Manager needs to investigate potential issues, the SOC Manager needs access to the information on these tabs.

The profile name and description make it easy to understand the profile's objective. This snip doesn't show all of the **Logs** permissions, but all of them are enabled for this profile.



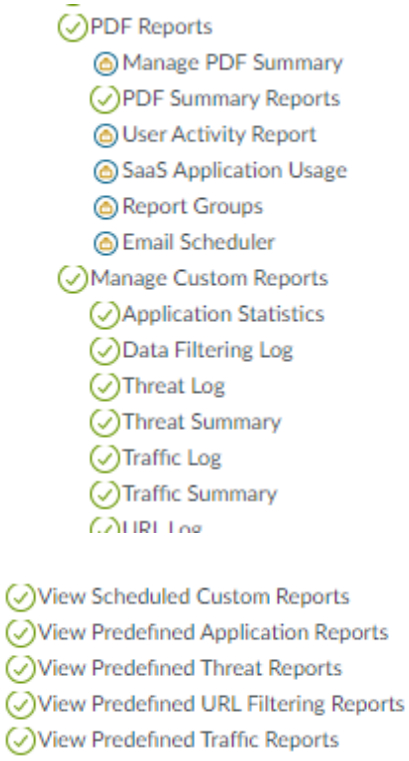
The next snip shows permissions for more informational objects on the **Monitor** tab. The SOC Manager uses these tools to investigate potential issues and therefore requires access.



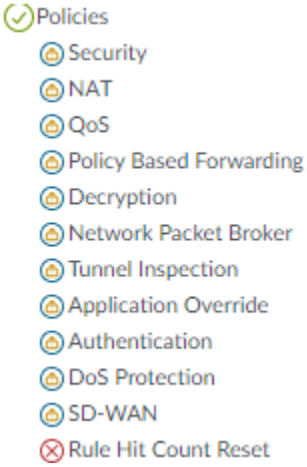
The next two snips show permissions for PDF Reports, Custom Reports, and predefined reports on the **Monitor** tab. While the SOC Manager needs access to PDF reports to gather information, in this example, the SOC Manager does not need to configure reports, so access is set to read-only (summary reports are not configurable). However, the SOC Manager needs to manage custom reports to investigate specific potential issues, so full access permissions are

https://t.me/learningnets

granted for all custom reports (including those not shown in the snip). Finally the SOC Manager requires access to predefined reports for investigating potential issues.



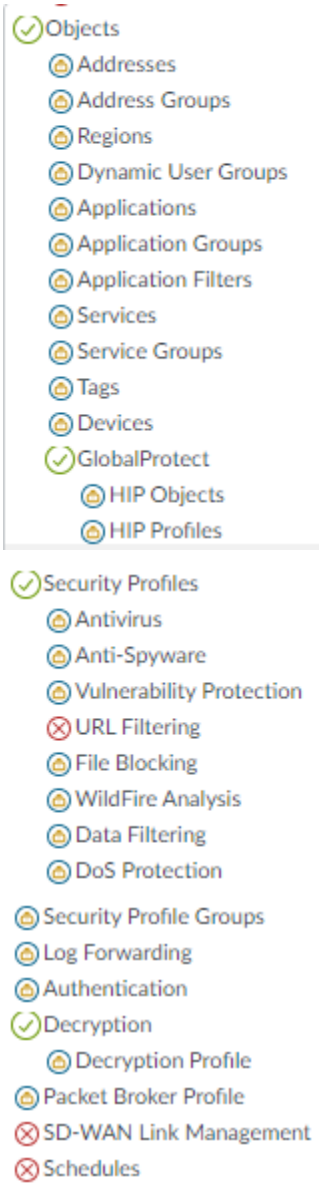
Because the SOC Manager is an investigator and not an administrator who configures the firewall, permissions for the **Policies** tab are read-only, with the exception of resetting the rule hit count. Resetting the rule hit count is not one of the SOC Manager's duties (and changing the hit count could adversely affect or confuse other administrators), so access is disabled. Read access enables the SOC Manager to investigate the construction of a policy that the SOC Manager suspects may have caused an issue.



Permissions for the **Objects** tab are also read-only for the same reason—the SOC Manager's job doesn't require configuration, so no configuration permissions are assigned. For areas that aren't included in the SOC Manager's duties, access is disabled. In this example, the SOC Manager has read-only access to investigate objects configurations for all objects except **URL**

https://t.me/learningnets

Filtering, **SD-WAN Link Management** and **Schedules**, which are under the control of different administrators in this example.

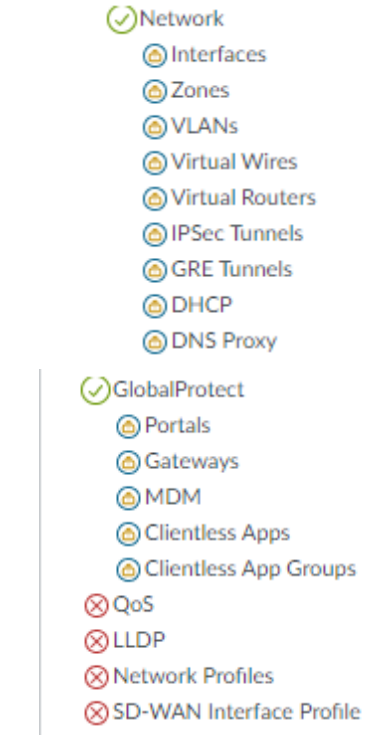


For **Network** tab permissions, the scenario is similar: the SOC Manager doesn't need to configure any of the objects, but may need information to investigate issues, so read-only access is assigned to the areas that the SOC Manager may need to investigate. In this example,

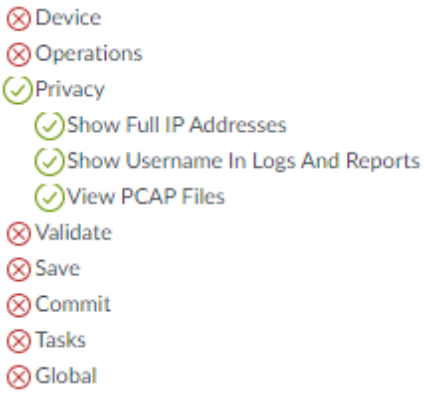
 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

access is disabled for QoS, LLDP, Network Profiles, or SD-WAN Interface profiles because these items are not part of the SOC Manager's duties.



In this example, the SOC Manager needs no access to the **Device** tab capabilities for investigative purposes, so all **Device** tab permissions are blocked. In addition, investigation doesn't require commit actions or access to any of the remaining actions, so those permissions are also blocked.

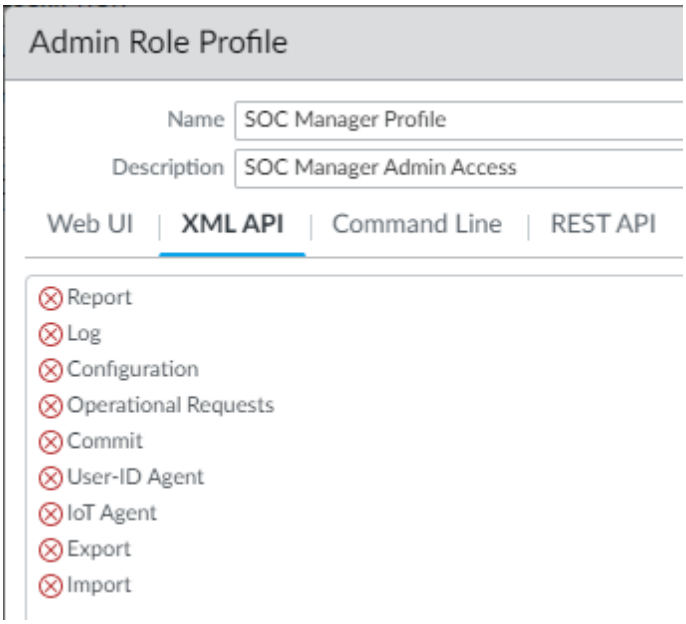


 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

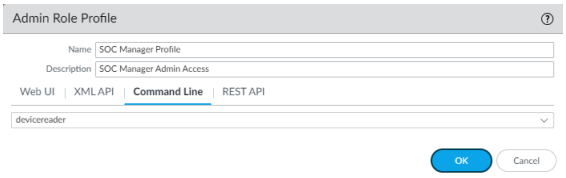
STEP 2 | Configure XML API access permissions.

The following snip shows that all XML API permissions are disabled for the SOC Manager because the SOC Manager doesn't access the firewall using XML API commands.



STEP 3 | Configure Command Line (CLI) access permissions.

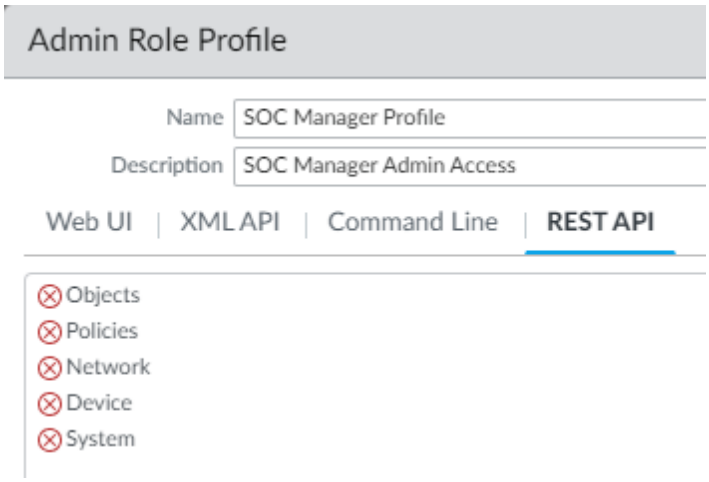
CLI access permissions are read-only for the SOC Manager because the SOC Manager needs access to logs and other monitoring tools and also needs to be able to see certain configurations in order to investigate potential issues. However, the SOC Manager doesn't configure the firewall, so no configuration permissions are assigned. The access level is set to **devicereader** instead of to **superreader** because the SOC Manager doesn't need access to password profiles or to other administrative accounts.



https://t.me/learningnets

STEP 4 | Configure REST API access permissions.

The SOC Manager doesn't access the firewall using REST API commands, so all REST API access is disabled.



Administrative Authentication

You can configure the following types of authentication and authorization (role and access domain assignment) for firewall administrators:

Authentication Method	Authorization Method	Description
Local	Local	The administrative account credentials and authentication mechanisms are local to the firewall. You can define the accounts with or without a user database that is local to the firewall—see Local Authentication for the advantages and disadvantages of using a local database. You use the firewall to manage role assignments but access domains are not supported. For details, see Configure Local or External Authentication for Firewall Administrators.
SSH Keys	Local	The administrative accounts are local to the firewall, but authentication to the CLI is based on SSH keys. You use the firewall to manage role assignments but access domains are not supported. For details, see Configure SSH Key-Based Administrator Authentication to the CLI.
Certificates	Local	The administrative accounts are local to the firewall, but authentication to the web interface is based on client certificates. You use the firewall to manage role assignments but access domains are not supported. For details, see Configure Certificate-Based Administrator Authentication to the Web Interface.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Authentication Method	Authorization Method	Description
External service	Local	The administrative accounts you define locally on the firewall serve as references to the accounts defined on an external Multi-Factor Authentication, SAML, Kerberos, TACACS+, RADIUS, or LDAP server. The external server performs authentication. You use the firewall to manage role assignments but access domains are not supported. For details, see Configure Local or External Authentication for Firewall Administrators.
External service	External service	The administrative accounts are defined on an external SAML, TACACS+, or RADIUS server. The server performs both authentication and authorization. For authorization, you define Vendor-Specific Attributes (VSAs) on the TACACS+ or RADIUS server, or SAML attributes on the SAML server. PAN-OS maps the attributes to administrator roles, access domains, user groups, and virtual systems that you define on the firewall. For details, see: • Configure SAML Authentication • Configure TACACS+ Authentication • Configure RADIUS Authentication

Configure Administrative Accounts and Authentication

If you have already configured an authentication profile (see Configure an Authentication Profile and Sequence) or you don't require one to authenticate administrators, you are ready to Configure a Firewall Administrator Account. Otherwise, perform one of the other procedures listed below to configure administrative accounts for specific types of authentication.

- Configure a Firewall Administrator Account
- Configure Local or External Authentication for Firewall Administrators
- Configure Certificate-Based Administrator Authentication to the Web Interface
- Configure SSH Key-Based Administrator Authentication to the CLI
- Configure API Key Lifetime

Configure a Firewall Administrator Account

Administrative accounts specify roles and authentication methods for firewall administrators. The service that you use to assign roles and perform authentication determines whether you add the accounts on the firewall, on an external server, or both (see Administrative Authentication). If the authentication method relies on a local firewall database or an external service, you must configure an authentication profile before adding an administrative account (see Configure Administrative Accounts and Authentication). If you already configured the authentication profile or you will use Local Authentication without a firewall database, perform the following steps to add an administrative account on the firewall.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Create a separate administrative account for each person who needs access to the administrative or reporting functions of the firewall. This enables you to better protect the firewall from unauthorized configuration and enables logging of the actions of individual administrators.

Make sure you are following the Adminstrative Access Best Practices to ensure that you are securing administrative access to your firewalls and other security devices in a way that prevents successful attacks.

STEP 1 | Modify the number of supported administrator accounts.

Configure the total number of supported concurrent administrative accounts sessions for a firewall in the normal operational mode or in FIPS-CC mode. You can allow up to four concurrent administrative account sessions or configure the firewall to support an unlimited number of concurrent administrative account sessions.

1. Select **Device** > **Setup** > **Management** and edit the Authentication Settings.
2. Edit the **Max Session Count** to specify the number of supported concurrent sessions (range is **0** to **4**) allowed for all administrator and user accounts.

   Enter **0** to configure the firewall to support an unlimited number of administrative accounts.
3. Edit the **Max Session Time** in minutes for an administrative account. Default is 720 minutes.
4. Click **OK**.
5. **Commit**.

You can also configure the total number of supported concurrent sessions by logging in to the firewall CLI.

```
admin> configure
```

```
admin# set deviceconfig setting management admin-session
 max-session-count <0-4>
```

```
admin# set deviceconfig setting management admin-session
 max-session-time <0, 60-1499>
```

```
admin# commit
```

STEP 2 | Select **Device** > **Administrators** and **Add** an account.

STEP 3 | Enter a user **Name**.

If the firewall uses a local user database to authenticate the account, enter the name that you specified for the account in the database (see Add the user group to the local database.)

STEP 4 | Select an **Authentication Profile** or sequence if you configured either for the administrator.

If the firewall uses Local Authentication without a local user database for the account, select **None** (default) and enter a **Password**.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 5 | Select the **Administrator Type**.

If you configured a custom role for the user, select **Role Based** and select the Admin Role **Profile**. Otherwise, select **Dynamic** (default) and select a dynamic role. If the dynamic role is **virtual system administrator**, add one or more virtual systems that the virtual system administrator is allowed to manage.

STEP 6 | (Optional) Select a **Password Profile** for administrators that the firewall authenticates locally without a local user database. For details, see Define a Password Profile.

STEP 7 | Click **OK** and **Commit**.

Configure Local or External Authentication for Firewall Administrators

You can use Local Authentication and External Authentication Services to authenticate administrators who access the firewall. These authentication methods prompt administrators to respond to one or more authentication challenges, such as a login page for entering a username and password.

If you use an external service to manage both authentication and authorization (role and access domain assignments), see:

- Configure SAML Authentication
- Configure TACACS+ Authentication
- Configure RADIUS Authentication

To authenticate administrators without a challenge-response mechanism, you can Configure Certificate-Based Administrator Authentication to the Web Interface *and* Configure SSH Key-Based Administrator Authentication to the CLI.

STEP 1 | (External authentication only) Enable the firewall to connect to an external server for authenticating administrators.

Configure a server profile:

- Add a RADIUS server profile.

  If the firewall integrates with a Multi-Factor Authentication (MFA) service through RADIUS, you must add a RADIUS server profile. In this case, the MFA service provides all the authentication factors (challenges). If the firewall integrates with an MFA service through a vendor API, you can still use a RADIUS server profile for the first factor but MFA server profiles are required for additional factors.
- Add an MFA server profile.
- Add a TACACS+ server profile.
- Add a SAML IdP server profile. You cannot combine Kerberos single sign-on (SSO) with SAML SSO; you can use only one type of SSO service.
- Add a Kerberos server profile.
- Add an LDAP server profile.

https://t.me/learningnets

STEP 2 | (Local database authentication only) Configure a user database that is local to the firewall.

1. Add the user account to the local database.
2. (Optional) Add the user group to the local database.

STEP 3 | (Local authentication only) Define password complexity and expiration settings.

These settings help protect the firewall against unauthorized access by making it harder for attackers to guess passwords.

1. Define global password complexity and expiration settings for all local administrators. The settings don't apply to local database accounts for which you specified a password hash instead of a password (see Local Authentication).
   1. Select **Device** > **Setup** > **Management** and edit the Minimum Password Complexity settings.
   2. Select **Enabled**.
   3. Define the password settings and click **OK**.
2. Define a Password Profile.

   You assign the profile to administrator accounts for which you want to override the global password expiration settings. The profiles are available only to accounts that are not associated with a local database (see Local Authentication).
   1. Select **Device** > **Password Profiles** and **Add** a profile.
   2. Enter a **Name** to identify the profile.
   3. Define the password expiration settings and click **OK**.

STEP 4 | (Kerberos SSO only) Create a Kerberos keytab.

A keytab is a file that contains Kerberos account information for the firewall. To support Kerberos SSO, your network must have a Kerberos infrastructure.

STEP 5 | Configure an authentication profile.



If your administrative accounts are stored across multiple types of servers, you can create an authentication profile for each type and add all the profiles to an authentication sequence.

Configure an Authentication Profile and Sequence. In the authentication profile, specify the **Type** of authentication service and related settings:

- **External service**—Select the **Type** of external service and select the **Server Profile** you created for it.
- **Local database authentication**—Set the **Type** to **Local Database**.
- **Local authentication without a database**—Set the **Type** to **None**.
- **Kerberos SSO**—Specify the **Kerberos Realm** and **Import** the **Kerberos Keytab**.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 6 | Assign the authentication profile or sequence to an administrator account.

1. Configure a Firewall Administrator Account.
   - Assign the **Authentication Profile** or sequence that you configured.
   - (Local database authentication only) Specify the **Name** of the user account you added to the local database.
2. **Commit** your changes.
3. (Optional) Test Authentication Server Connectivity to verify that the firewall can use the authentication profile to authenticate administrators.

Configure Certificate-Based Administrator Authentication to the Web Interface

As a more secure alternative to password-based authentication to the firewall web interface, you can configure certificate-based authentication for administrator accounts that are local to the firewall. Certificate-based authentication involves the exchange and verification of a digital signature instead of a password.

Configuring certificate-based authentication for any administrator disables the username/password logins for all administrators on the firewall; administrators thereafter require the certificate to log in.

STEP 1 | Generate a certificate authority (CA) certificate on the firewall.

You will use this CA certificate to sign the client certificate of each administrator.

Create a Self-Signed Root CA Certificate.

Alternatively, Import a Certificate and Private Key *from your enterprise CA or a third-party CA.*

STEP 2 | Configure a certificate profile for securing access to the web interface.

Configure a Certificate Profile.

- Set the **Username Field** to **Subject**.
- In the CA Certificates section, **Add** the **CA Certificate** you just created or imported.

STEP 3 | Configure the firewall to use the certificate profile for authenticating administrators.

1. Select **Device** > **Setup** > **Management** and edit the Authentication Settings.
2. Select the **Certificate Profile** you created for authenticating administrators and click **OK**.

STEP 4 | Configure the administrator accounts to use client certificate authentication.

For each administrator who will access the firewall web interface, Configure a Firewall Administrator Account and select **Use only client certificate authentication**.

If you have already deployed client certificates that your enterprise CA generated, skip to Step 8. Otherwise, go to Step 5.

STEP 5 | Generate a client certificate for each administrator.

Generate a Certificate. In the **Signed By** drop-down, select a self-signed root CA certificate.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 6 | Export the client certificate.

1. Export a Certificate and Private Key.
2. **Commit** your changes. The firewall restarts and terminates your login session. Thereafter, administrators can access the web interface only from client systems that have the client certificate you generated.

STEP 7 | Import the client certificate into the client system of each administrator who will access the web interface.

Refer to your web browser documentation.

STEP 8 | Verify that administrators can access the web interface.

1. Open the firewall IP address in a browser on the computer that has the client certificate.
2. When prompted, select the certificate you imported and click **OK**. The browser displays a certificate warning.
3. Add the certificate to the browser exception list.
4. Click **Login**. The web interface should appear without prompting you for a username or password.

Configure SSH Key-Based Administrator Authentication to the CLI

For administrators who use Secure Shell (SSH) to access the CLI of a Palo Alto Networks firewall, SSH keys provide a more secure authentication method than passwords. SSH keys almost eliminate the risk of brute-force attacks, provide the option for two-factor authentication (key and passphrase), and don't send passwords over the network. SSH keys also enable automated scripts to access the CLI.

STEP 1 | Use an SSH key generation tool to create an asymmetric keypair on the client system of the administrator.

The supported key formats are IETF SECSH and Open SSH. The supported algorithms are DSA (1,024 bits) and RSA (768-4,096 bits).

For the commands to generate the keypair, refer to your SSH client documentation.

The public key and private key are separate files. Save both to a location that the firewall can access. For added security, enter a passphrase to encrypt the private key. The firewall prompts the administrator for this passphrase during login.

STEP 2 | Configure the administrator account to use public key authentication.

1. Configure a Firewall Administrator Account.
   - Configure the authentication method to use as a fallback if SSH key authentication fails. If you configured an **Authentication Profile** for the administrator, select it in the drop-down. If you select **None**, you must enter a **Password** and **Confirm Password**.
   - Select **Use Public Key Authentication (SSH)**, then **Import Key**, **Browse** to the public key you just generated, and click **OK**.
2. **Commit** your changes.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 3 | Configure the SSH client to use the private key to authenticate to the firewall.

Perform this task on the client system of the administrator. For the steps, refer to your SSH client documentation.

STEP 4 | Verify that the administrator can access the firewall CLI using SSH key authentication.

1. Use a browser on the client system of the administrator to go to the firewall IP address.
2. Log in to the firewall CLI as the administrator. After entering a username, you will see the following output (the key value is an example):

```
Authenticating with public key “dsa-key-20130415”
```

3. If prompted, enter the passphrase you defined when creating the keys.

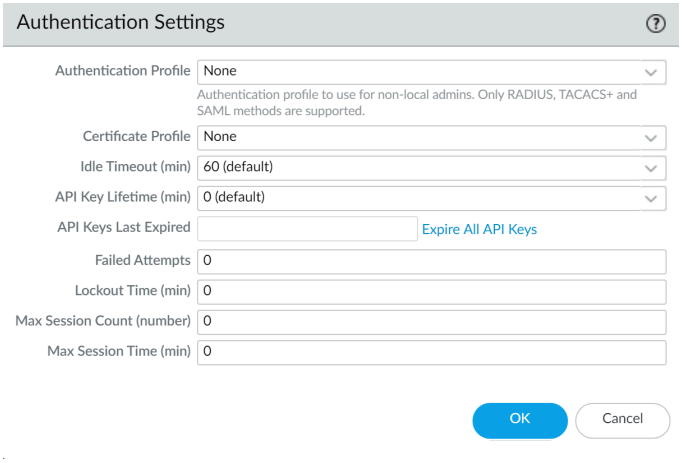
Configure API Key Lifetime

The API keys on the firewall and Panorama enable you to authenticate API calls to the XML API and REST API. Because these keys grant access to the firewall and Panorama that are critical elements of your security posture, as a best practice, specify an API key lifetime to enforce regular key rotation. After you specify the key lifetime, when you regenerate an API key, each key is unique.

In addition to setting a key lifetime that prompts you to regenerate new keys periodically, you can also revoke all currently valid API keys in the event one or more keys are compromised. Revoking keys is a way to expire all currently valid keys.

STEP 1 | Select **Device** > **Setup** > **Management**.

STEP 2 | Edit Authentication Settings to specify the **API Key Lifetime (min)**.



Set the API key lifetime to protect against compromise and to reduce the effects of an accidental exposure. By default, the API key lifetime is set to 0, which means that the keys will never expire. To ensure that your keys are frequently rotated and each key is unique when regenerated, you must specify a validity period that ranges between 1—525600 minutes. Refer to the audit and compliance policies for your enterprise to determine how you should specify the lifetime for which your API keys are valid.

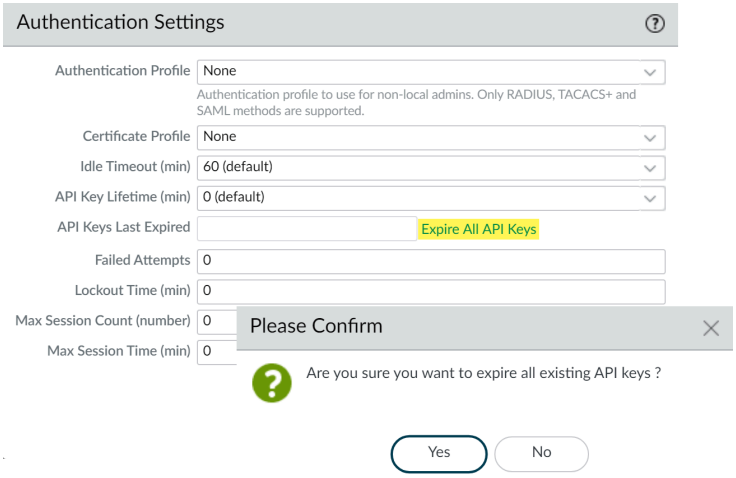
STEP 3 | **Commit** the changes.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 4 | (To revoke all API keys) Select **Expire all API Keys** to reset currently valid API keys.

If you have just set a key lifetime and want to reset all API keys to adhere to the new term, you can expire all existing keys.



On confirmation, the keys are revoked and you can view the timestamp for when the **API Keys Last Expired**.

Configure Tracking of Administrator Activity

Track administrator activity on the firewall web interface and CLI to achieve real time reporting of activity across your firewall. If you have reason to believe an administrator account is compromised, you have a full history of where this administrator account navigated throughout the web interface or what operational commands they executed so you can analyze in detail and respond to all actions the compromised administrator took.

When an event occurs, an audit log is generated and forwarded to the specified syslog server each time an administrator navigates through the web interface or when an operational command is executed in the CLI. An audit log is generated for each navigation or commend executed. Take for example if you want to create a new address object. An audit log is generated when you click on **Objects**, and a second audit log is generated when you then click on Addresses.

Audit logs are only visible as syslogs forwarded to your syslog server and cannot be viewed in the firewall web interface. Audit logs can only be forwarded to a syslog server, cannot be forwarded to Cortex Data Lake (CDL), and are not stored locally on the firewall.

STEP 1 | Configure a syslog server profile to forward audit logs of administrator activity on the firewall.

This step is required to successfully store audit logs for tracking administrator activity on the firewall.

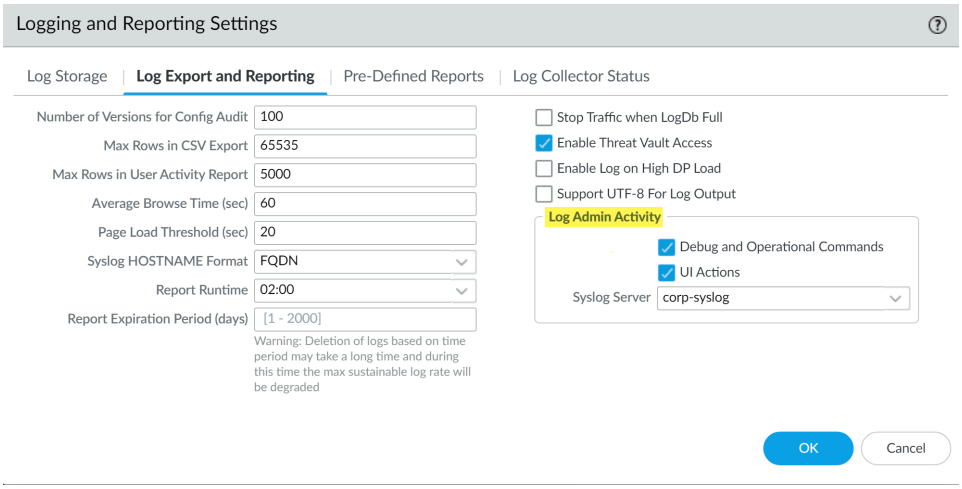
1. Log in to the firewall web interface.
2. Configure a syslog server profile.

https://t.me/learningnets

STEP 2 | Configure tracking of administrator activity.

1. Select **Device** > **Setup** > **Management** and edit the Logging and Reporting Settings.
2. Select **Log Export and Reporting**.
3. In the Log Admin Activity section, configure what administrator activity to track.
   - **Operational Commands**—Generate an audit log when an administrator executes an operational or debug command in the CLI or an operational command triggered from the web interface. See the CLI Operational Command Hierarchy for a full list of PAN-OS operational and debug commands.
   - **UI Actions**—Generate an audit log when an administrator navigates throughout the web interface. This includes navigation between configuration tabs, as well as individual objects within a tab.

     For example, an audit log is generated when an administrator navigates from the **ACC** to the **Policies** tab. Additionally, an audit log is generated when an administrator navigates from **Objects** > **Addresses** to **Objects** > **Tags**.
   - **Syslog Server**—Select a target syslog server profile to forward audit logs.
4. Click **OK**
5. Select **Commit**.



 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Reference: Web Interface Administrator Access

You can configure privileges for an entire firewall or for one or more virtual systems (on platforms that support multiple virtual systems). Within that **Device** or **Virtual System** designation, you can configure privileges for custom administrator roles, which are more granular than the fixed privileges associated with a dynamic administrator role.

Configuring privileges at a granular level ensures that lower level administrators cannot access certain information. You can create custom roles for firewall administrators (see Configure a Firewall Administrator Account), Panorama administrators, or Device Group and Template administrators (refer to the Panorama Administrator's Guide). You apply the admin role to a custom role-based administrator account where you can assign one or more virtual systems. The following topics describe the privileges you can configure for custom administrator roles.

- Web Interface Access Privileges
- Panorama Web Interface Access Privileges

Web Interface Access Privileges

If you want to prevent a role-based administrator from accessing specific tabs on the web interface, you can disable the tab and the administrator will not even see it when logging in using the associated role-based administrative account. For example, you could create an Admin Role Profile for your operations staff that provides access to the **Device** and **Network** tabs only and a separate profile for your security administrators that provides access to the **Object**, **Policy**, and **Monitor** tabs.

An admin role can apply at the **Device** level or **Virtual System** level as defined by the **Device** or **Virtual System** radio button. If you select **Virtual System**, the admin assigned this profile is restricted to the virtual system(s) he or she is assigned to. Furthermore, only the **Device** > **Setup** > **Services** > **Virtual Systems** tab is available to that admin, not the **Global** tab.

The following topics describe how to set admin role privileges to the different parts of the web interface:

- Define Access to the Web Interface Tabs
- Provide Granular Access to the Monitor Tab
- Provide Granular Access to the Policy Tab
- Provide Granular Access to the Objects Tab
- Provide Granular Access to the Network Tab
- Provide Granular Access to the Device Tab
- Define User Privacy Settings in the Admin Role Profile
- Restrict Administrator Access to Commit and Validate Functions
- Provide Granular Access to Global Settings
- Provide Granular Access to the Panorama Tab
- Provide Granular Access to Operations Settings

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Define Access to the Web Interface Tabs

The following table describes the top-level access privileges you can assign to an admin role profile (**Device** > **Admin Roles**). You can enable, disable, or define read-only access privileges at the top-level tabs in the web interface.

Access Level	Description	Enable	Read Only	Disable
Dashboard	Controls access to the **Dashboard** tab. If you disable this privilege, the administrator will not see the tab and will not have access to any of the Dashboard widgets.	Yes	No	Yes
ACC	Controls access to the Application Command Center (ACC). If you disable this privilege, the **ACC** tab will not display in the web interface. Keep in mind that if you want to protect the privacy of your users while still providing access to the ACC, you can disable the **Privacy** > **Show Full IP Addresses** option and/or the **Show User Names In Logs And Reports** option.	Yes	No	Yes
Monitor	Controls access to the **Monitor** tab. If you disable this privilege, the administrator will not see the **Monitor** tab and will not have access to any of the logs, packet captures, session information, reports or to App Scope. For more granular control over what monitoring information the administrator can see, leave the Monitor option enabled and then enable or disable specific nodes on the tab as described in Provide Granular Access to the Monitor Tab.	Yes	No	Yes
Policies	Controls access to the **Policies** tab. If you disable this privilege, the administrator will not see the **Policies** tab and will not have access to any policy information. For more granular control over what policy information the administrator can see, for example to enable access to a specific type of policy or to enable read-only access to policy information, leave the **Policies** option enabled and then enable or disable specific nodes on the tab as	Yes	No	Yes

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Access Level	Description	Enable	Read Only	Disable
	described in Provide Granular Access to the Policy Tab.			
Objects	Controls access to the **Objects** tab. If you disable this privilege, the administrator will not see the **Objects** tab and will not have access to any objects, security profiles, log forwarding profiles, decryption profiles, or schedules. For more granular control over what objects the administrator can see, leave the **Objects** option enabled and then enable or disable specific nodes on the tab as described in Provide Granular Access to the Objects Tab.	Yes	No	Yes
Network	Controls access to the **Network** tab. If you disable this privilege, the administrator will not see the **Network** tab and will not have access to any interface, zone, VLAN, virtual wire, virtual router, IPsec tunnel, DHCP, DNS Proxy, GlobalProtect, or QoS configuration information or to the network profiles. For more granular control over what objects the administrator can see, leave the **Network** option enabled and then enable or disable specific nodes on the tab as described in Provide Granular Access to the Network Tab.	Yes	No	Yes
Device	Controls access to the **Device** tab. If you disable this privilege, the administrator will not see the **Device** tab and will not have access to any firewall-wide configuration information, such as User-ID, high availability, server profile or certificate configuration information. For more granular control over what objects the administrator can see, leave the **Objects** option enabled and then enable or disable specific nodes on the tab as described in Provide Granular Access to the Device Tab.	Yes	No	Yes

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Access Level	Description	Enable	Read Only	Disable
	*You cannot enable access to the **Admin Roles** or **Administrators** nodes for a role-based administrator even if you enable full access to the **Device** tab.*			

Provide Granular Access to the Monitor Tab

In some cases you might want to enable the administrator to view some but not all areas of the **Monitor** tab. For example, you might want to restrict operations administrators to the Config and System logs only, because they do not contain sensitive user data. Although this section of the administrator role definition specifies what areas of the **Monitor** tab the administrator can see, you can also couple privileges in this section with privacy privileges, such as disabling the ability to see usernames in logs and reports. One thing to keep in mind, however, is that any system-generated reports will still show usernames and IP addresses even if you disable that functionality in the role. For this reason, if you do not want the administrator to see any of the private user information, disable access to the specific reports as detailed in the following table.

The following table lists the **Monitor** tab access levels and the administrator roles for which they are available.

Device Group and Template roles can see log data only for the device groups that are within the access domains assigned to those roles.

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
Monitor	Enables or disables access to the **Monitor** tab. If disabled, the administrator will not see this tab or any of the associated logs or reports.	Firewall: Yes Panorama: Yes Device Group/ Template: Yes	Yes	No	Yes
Logs	Enables or disables access to all log files. You can also leave this privilege enabled and then disable specific logs that you do not want the administrator to see. Keep in mind that if you want to protect the privacy of your users while still providing access to one or more of the logs, you can disable the **Privacy** > **Show Full IP Addresses** option	Firewall: Yes Panorama: Yes Device Group/ Template: Yes	Yes	No	Yes

©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
	and/or the **Show User Names In Logs And Reports** option.				
Traffic	Specifies whether the administrator can see the traffic logs.	Firewall: Yes Panorama: Yes Device Group/ Template: Yes	Yes	No	Yes
Threat	Specifies whether the administrator can see the threat logs.	Firewall: Yes Panorama: Yes Device Group/ Template: Yes	Yes	No	Yes
URL Filtering	Specifies whether the administrator can see the URL filtering logs.	Firewall: Yes Panorama: Yes Device Group/ Template: Yes	Yes	No	Yes
WildFire Submissions	Specifies whether the administrator can see the WildFire logs. These logs are only available if you have a WildFire subscription.	Firewall: Yes Panorama: Yes Device Group/ Template: Yes	Yes	No	Yes
Data Filtering	Specifies whether the administrator can see the data filtering logs.	Firewall: Yes Panorama: Yes Device Group/ Template: Yes	Yes	No	Yes
HIP Match	Specifies whether the administrator can see the HIP Match logs. HIP Match logs are available only if you have a GlobalProtect license (subscription).	Firewall: Yes Panorama: Yes Device Group/ Template: Yes	Yes	No	Yes
GlobalProtect	Specifies whether the administrator can see the GlobalProtect logs. These logs are available only if you have a GlobalProtect license (subscription).	Firewall: Yes Panorama: Yes Device Group/ Template: Yes	Yes	No	Yes

©2022 Palo Alto Networks, Inc.
https://t.me/learningnets

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
User-ID	Specifies whether the administrator can see the User-ID logs.	Firewall: Yes Panorama: Yes Device Group/ Template: Yes	Yes	No	Yes
GTP	Specifies whether the mobile network operator can see GTP logs.	Firewall: Yes Panorama: Yes Device Group/ Template: Yes	Yes	No	Yes
Tunnel Inspection	Specifies whether the administrator can see the Tunnel Inspection logs.	Firewall: Yes Panorama: Yes Device Group/ Template: Yes	Yes	No	Yes
SCTP	Specifies whether the mobile network operator can see Stream Control Transmission Protocol (SCTP) logs. *You must enable SCTP on Panorama (**Device** > **Setup** > **Management**) before you can control Administrator access to SCTP logs, custom reports, or predefined reports for Panorama and Device Group/ Template.*	Firewall: Yes Panorama: Yes Device Group/ Template: Yes	Yes	No	Yes
Configuration	Specifies whether the administrator can see the configuration logs.	Firewall: Yes Panorama: Yes Device Group/ Template: No	Yes	No	Yes

©2022 Palo Alto Networks, Inc.
https://t.me/learningnets

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
System	Specifies whether the administrator can see the system logs.	Firewall: Yes Panorama: Yes Device Group/ Template: No	Yes	No	Yes
Alarms	Specifies whether the administrator can see system-generated alarms.	Firewall: Yes Panorama: Yes Device Group/ Template: Yes	Yes	No	Yes
Authentication	Specifies whether the administrator can see the Authentication logs.	Firewall: Yes Panorama: Yes Device Group/ Template: No	Yes	No	Yes
Automated Correlation Engine	Enables or disables access to the correlation objects and correlated event logs generated on the firewall.	Firewall: Yes Panorama: Yes Device Group/ Template: Yes	Yes	No	Yes
Correlation Objects	Specifies whether the administrator can view and enable/disable the correlation objects.	Firewall: Yes Panorama: Yes Device Group/ Template: Yes	Yes	No	Yes
Correlated Events	Specifies whether the administrator can view and enable/disable the correlation events.	Firewall: Yes Panorama: Yes Device Group/ Template: Yes	Yes	No	Yes
Packet Capture	Specifies whether the administrator can see packet captures (pcaps) from the **Monitor** tab. Keep in mind that packet captures are raw flow data and as such may contain user IP addresses. Disabling the **Show Full IP Addresses** privileges will not obfuscate the IP address in the pcap and	Firewall: Yes Panorama: No Device Group/ Template: No	Yes	Yes	Yes

©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
	you should therefore disable the Packet Capture privilege if you are concerned about user privacy.				
App Scope	Specifies whether the administrator can see the App Scope visibility and analysis tools. Enabling App Scope enables access to all of the **App Scope** charts.	Firewall: Yes Panorama: Yes Device Group/ Template: Yes	Yes	No	Yes
Session Browser	Specifies whether the administrator can browse and filter current running sessions on the firewall. Keep in mind that the session browser shows raw flow data and as such may contain user IP addresses. Disabling the **Show Full IP Addresses** privileges will not obfuscate the IP address in the session browser and you should therefore disable the **Session Browser** privilege if you are concerned about user privacy.	Firewall: Yes Panorama: No Device Group/ Template: No	Yes	No	Yes
Block IP List	Specifies whether the administrator can view the block list (Enable or Read Only) and delete entries from the list (Enable). If you disable the setting, the administrator won't be able to view or delete entries from the block list.	Firewall: Yes Panorama: under Context Switch UI: Yes Template: Yes	Yes	Yes	Yes
Botnet	Specifies whether the administrator can generate and view botnet analysis reports or view botnet reports in read-only mode. Disabling the **Show Full IP Addresses** privileges will not obfuscate the IP address in scheduled botnet reports and you should therefore disable	Firewall: Yes Panorama: No Device Group/ Template: No	Yes	Yes	Yes

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
	the **Botnet** privilege if you are concerned about user privacy.				
PDF Reports	Enables or disables access to all PDF reports. You can also leave this privilege enabled and then disable specific PDF reports that you do not want the administrator to see. Keep in mind that if you want to protect the privacy of your users while still providing access to one or more of the reports, you can disable the **Privacy** > **Show Full IP Addresses** option and/or the **Show User Names In Logs And Reports** option.	Firewall: Yes Panorama: Yes Device Group/ Template: Yes	Yes	No	Yes
Manage PDF Summary	Specifies whether the administrator can view, add or delete PDF summary report definitions. With read-only access, the administrator can see PDF summary report definitions, but not add or delete them. If you disable this option, the administrator can neither view the report definitions nor add/delete them.	Firewall: Yes Panorama: Yes Device Group/ Template: Yes	Yes	Yes	Yes
PDF Summary Reports	Specifies whether the administrator can see the generated PDF Summary reports in **Monitor** > **Reports**. If you disable this option, the **PDF Summary Reports** category will not display in the **Reports** node.	Firewall: Yes Panorama: Yes Device Group/ Template: Yes	Yes	No	Yes
User Activity Report	Specifies whether the administrator can view, add or delete User Activity report definitions and download the reports. With read-only access, the administrator can see User Activity report definitions, but not add, delete, or download them. If you disable this option,	Firewall: Yes Panorama: Yes Device Group/ Template: Yes	Yes	Yes	Yes

©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
	the administrator cannot see this category of PDF report.				
SaaS Application Usage Report	Specifies whether the administrator can view, add or delete a SaaS application usage report. With read-only access, the administrator can see the SaaS application usage report definitions, but cannot add or delete them. If you disable this option, the administrator can neither view the report definitions nor add or delete them.	Firewall: Yes Panorama: Yes Device Group/ Template: Yes	Yes	Yes	Yes
Report Groups	Specifies whether the administrator can view, add or delete report group definitions. With read-only access, the administrator can see report group definitions, but not add or delete them. If you disable this option, the administrator cannot see this category of PDF report.	Firewall: Yes Panorama: Yes Device Group/ Template: Yes	Yes	Yes	Yes
Email Scheduler	Specifies whether the administrator can schedule report groups for email. Because the generated reports that get emailed may contain sensitive user data that is not removed by disabling the **Privacy** > **Show Full IP Addresses** option and/ or the **Show User Names In Logs And Reports** options and because they may also show log data to which the administrator does not have access, you should disable the **Email Scheduler** option if you have user privacy requirements.	Firewall: Yes Panorama: Yes Device Group/ Template: Yes	Yes	Yes	Yes
Manage Custom Reports	Enables or disables access to all custom report functionality. You can also leave this privilege enabled and then	Firewall: Yes Panorama: Yes	Yes	No	Yes

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
	disable specific custom report categories that you do not want the administrator to be able to access. Keep in mind that if you want to protect the privacy of your users while still providing access to one or more of the reports, you can disable the **Privacy** > **Show Full IP Addresses** option and/or the **Show User Names In Logs And Reports** option. *Reports that are scheduled to run rather than run on demand will show IP address and user information. In this case, be sure to restrict access to the corresponding report areas. In addition, the custom report feature does not restrict the ability to generate reports that contain log data contained in logs that are excluded from the administrator role.*	Device Group/ Template: Yes			
Application Statistics	Specifies whether the administrator can create a custom report that includes data from the application statistics database.	Firewall: Yes Panorama: Yes Device Group/ Template: Yes	Yes	No	Yes
Data Filtering Log	Specifies whether the administrator can create a custom report that includes data from the Data Filtering logs.	Firewall: Yes Panorama: Yes Device Group/ Template: Yes	Yes	No	Yes

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
Threat Log	Specifies whether the administrator can create a custom report that includes data from the Threat logs.	Firewall: Yes Panorama: Yes Device Group/ Template: Yes	Yes	No	Yes
Threat Summary	Specifies whether the administrator can create a custom report that includes data from the Threat Summary database.	Firewall: Yes Panorama: Yes Device Group/ Template: Yes	Yes	No	Yes
Traffic Log	Specifies whether the administrator can create a custom report that includes data from the Traffic logs.	Firewall: Yes Panorama: Yes Device Group/ Template: Yes	Yes	No	Yes
Traffic Summary	Specifies whether the administrator can create a custom report that includes data from the Traffic Summary database.	Firewall: Yes Panorama: Yes Device Group/ Template: Yes	Yes	No	Yes
URL Log	Specifies whether the administrator can create a custom report that includes data from the URL Filtering logs.	Firewall: Yes Panorama: Yes Device Group/ Template: Yes	Yes	No	Yes
URL Summary	Specifies whether the administrator can create a custom report that includes data from the URL Summary database.	Firewall: Yes Panorama: Yes Device Group/ Template: Yes	Yes	No	Yes
HIP Match	Specifies whether the administrator can create a custom report that includes data from the HIP Match logs.	Firewall: Yes Panorama: Yes Device Group/ Template: Yes	Yes	No	Yes
GlobalProtect	Specifies whether the administrator can create a	Firewall: Yes Panorama: Yes	Yes	No	Yes

©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
	custom report that includes data from the GlobalProtect logs.	Device Group/ Template: Yes			
WildFire Log	Specifies whether the administrator can create a custom report that includes data from the WildFire logs.	Firewall: Yes Panorama: Yes Device Group/ Template: Yes	Yes	No	Yes
GTP Log	Specifies whether the mobile network operator can create a custom report that includes data from GTP logs.	Firewall: Yes Panorama: Yes Device Group/ Template: Yes	Yes	No	Yes
GTP Summary	Specifies whether the mobile network operator can create a custom report that includes data from GTP logs.	Firewall: Yes Panorama: Yes Device Group/ Template: Yes	Yes	No	Yes
Tunnel Log	Specifies whether the administrator can create a custom report that includes data from tunnel inspection logs.	Firewall: Yes Panorama: Yes Device Group/ Template: Yes	Yes	No	Yes
Tunnel Summary	Specifies whether the administrator can create a custom report that includes data from the Tunnel Summary database.	Firewall: Yes Panorama: Yes Device Group/ Template: Yes	Yes	No	Yes
SCTP Log	Specifies whether the mobile network operator can create a custom report that includes data from SCTP logs.	Firewall: Yes Panorama: Yes Device Group/ Template: Yes	Yes	No	Yes
SCTP Summary	Specifies whether the mobile network operator can create a custom report that includes data from the SCTP Summary database.	Firewall: Yes Panorama: Yes Device Group/ Template: Yes	Yes	No	Yes

©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
User-ID	Specifies whether the administrator can create a custom report that includes data from the User-ID logs.	Firewall: Yes Panorama: Yes Device Group/ Template: Yes	Yes	No	Yes
Authentication	Specifies whether the administrator can create a custom report that includes data from the Authentication logs.	Firewall: Yes Panorama: Yes Device Group/ Template: Yes	Yes	No	Yes
View Scheduled Custom Reports	Specifies whether the administrator can view a custom report that has been scheduled to generate.	Firewall: Yes Panorama: Yes Device Group/ Template: Yes	Yes	No	Yes
View Predefined Application Reports	Specifies whether the administrator can view Application Reports. Privacy privileges do not impact reports available on the **Monitor** > **Reports** node and you should therefore disable access to the reports if you have user privacy requirements.	Firewall: Yes Panorama: Yes Device Group/ Template: Yes	Yes	No	Yes
View Predefined Threat Reports	Specifies whether the administrator can view Threat Reports. Privacy privileges do not impact reports available on the **Monitor** > **Reports** node and you should therefore disable access to the reports if you have user privacy requirements.	Firewall: Yes Panorama: Yes Device Group/ Template: Yes	Yes	No	Yes
View Predefined URL Filtering Reports	Specifies whether the administrator can view URL Filtering Reports. Privacy privileges do not impact reports available on the **Monitor** > **Reports** node and you should therefore disable access to the reports if you have user privacy requirements.	Firewall: Yes Panorama: Yes Device Group/ Template: Yes	Yes	No	Yes

©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
View Predefined Traffic Reports	Specifies whether the administrator can view Traffic Reports. Privacy privileges do not impact reports available on the **Monitor** > **Reports** node and you should therefore disable access to the reports if you have user privacy requirements.	Firewall: Yes Panorama: Yes Device Group/ Template: Yes	Yes	No	Yes
View Predefined GTP Reports	Specifies whether the mobile network operator can view GTP Reports. Privacy privileges do not impact reports available on the **Monitor** > **Reports** node and you should therefore disable access to the reports if you have user privacy requirements.	Firewall: Yes Panorama: Yes Device Group/ Template: Yes	Yes	No	Yes
View Predefined SCTP Reports	Specifies whether the mobile network operator can view SCTP Reports. Privacy privileges do not impact reports available on the **Monitor** > **Reports** node and you should therefore disable access to the reports if you have user privacy requirements.	Firewall: Yes Panorama: Yes Device Group/ Template: Yes	Yes	No	Yes

Provide Granular Access to the Policy Tab

If you enable the Policy option in the Admin Role profile, you can then enable, disable, or provide read-only access to specific nodes within the tab as necessary for the role you are defining. By enabling access to a specific policy type, you enable the ability to view, add, or delete policy rules. By enabling read-only access to a specific policy, you enable the administrator to view the corresponding policy rule base, but not add or delete rules. Disabling access to a specific type of policy prevents the administrator from seeing the policy rule base.

Because policy that is based on specific users (by username or IP address) must be explicitly defined, privacy settings that disable the ability to see full IP addresses or usernames do not apply to the Policy tab. Therefore, you should only allow access to the Policy tab to administrators that are excluded from user privacy restrictions.

Access Level	Description	Enable	Read Only	Disable
Security	Enable this privilege to allow the administrator to view, add, and/or delete	Yes	Yes	Yes

https://t.me/learningnets

Access Level	Description	Enable	Read Only	Disable
	security rules. Set the privilege to read-only if you want the administrator to be able to see the rules, but not modify them. To prevent the administrator from seeing the security rulebase, disable this privilege.			
NAT	Enable this privilege to allow the administrator to view, add, and/or delete NAT rules. Set the privilege to read-only if you want the administrator to be able to see the rules, but not modify them. To prevent the administrator from seeing the NAT rulebase, disable this privilege.	Yes	Yes	Yes
QoS	Enable this privilege to allow the administrator to view, add, and/or delete QoS rules. Set the privilege to read-only if you want the administrator to be able to see the rules, but not modify them. To prevent the administrator from seeing the QoS rulebase, disable this privilege.	Yes	Yes	Yes
Policy Based Forwarding	Enable this privilege to allow the administrator to view, add, and/or delete Policy-Based Forwarding (PBF) rules. Set the privilege to read-only if you want the administrator to be able to see the rules, but not modify them. To prevent the administrator from seeing the PBF rulebase, disable this privilege.	Yes	Yes	Yes
Decryption	Enable this privilege to allow the administrator to view, add, and/or delete decryption rules. Set the privilege to read-only if you want the administrator to be able to see the rules, but not modify them. To prevent the administrator from seeing the decryption rulebase, disable this privilege.	Yes	Yes	Yes
Network Packet Broker	Enable this privilege to allow the administrator to view, add, and/or delete Network Packet Broker policy rules. Set the privilege to read-only if you want the administrator to be able to see the rules, but not modify them. To prevent the	Yes	Yes	Yes

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Access Level	Description	Enable	Read Only	Disable
	administrator from seeing the Network Packet Broker rulebase in the interface, disable this privilege.			
Tunnel Inspection	Enable this privilege to allow the administrator to view, add, and/or delete Tunnel Inspection rules. Set the privilege to read-only if you want the administrator to be able to see the rules, but not modify them. To prevent the administrator from seeing the Tunnel Inspection rulebase, disable this privilege.	Yes	Yes	Yes
Application Override	Enable this privilege to allow the administrator to view, add, and/or delete application override policy rules. Set the privilege to read-only if you want the administrator to be able to see the rules, but not modify them. To prevent the administrator from seeing the application override rulebase, disable this privilege.	Yes	Yes	Yes
Authentication	Enable this privilege to allow the administrator to view, add, and/or delete Authentication policy rules. Set the privilege to read-only if you want the administrator to be able to see the rules, but not modify them. To prevent the administrator from seeing the Authentication rulebase, disable this privilege.	Yes	Yes	Yes
DoS Protection	Enable this privilege to allow the administrator to view, add, and/or delete DoS protection rules. Set the privilege to read-only if you want the administrator to be able to see the rules, but not modify them. To prevent the administrator from seeing the DoS protection rulebase, disable this privilege.	Yes	Yes	Yes
SD-WAN	Enable this privilege to allow the administrator to view, add, and/or delete SD-WAN policy rules. Set the privilege to read-only if you want the administrator to be able to see the rules, but not modify them. To prevent the administrator from	Yes	Yes	Yes

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Access Level	Description	Enable	Read Only	Disable
	seeing the SD-WAN policy rulebase, disable this privilege.			

Provide Granular Access to the Objects Tab

An *object* is a container that groups specific policy filter values—such as IP addresses, URLs, applications, or services—for simplified rule definition. For example, an address object might contain specific IP address definitions for the web and application servers in your DMZ zone.

When deciding whether to allow access to the objects tab as a whole, determine whether the administrator will have policy definition responsibilities. If not, the administrator probably does not need access to the tab. If, however, the administrator will need to create policy, you can enable access to the tab and then provide granular access privileges at the node level.

By enabling access to a specific node, you give the administrator the privilege to view, add, and delete the corresponding object type. Giving read-only access allows the administrator to view the already defined objects, but not create or delete any. Disabling a node prevents the administrator from seeing the node in the web interface.

Access Level	Description	Enable	Read Only	Disable
Addresses	Specifies whether the administrator can view, add, or delete address objects for use in security policy.	Yes	Yes	Yes
Address Groups	Specifies whether the administrator can view, add, or delete address group objects for use in security policy.	Yes	Yes	Yes
Regions	Specifies whether the administrator can view, add, or delete regions objects for use in security, decryption, or DoS policy.	Yes	Yes	Yes
Applications	Specifies whether the administrator can view, add, or delete application objects for use in policy.	Yes	Yes	Yes
Application Groups	Specifies whether the administrator can view, add, or delete application group objects for use in policy.	Yes	Yes	Yes
Application Filters	Specifies whether the administrator can view, add, or delete application filters for simplification of repeated searches.	Yes	Yes	Yes

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Access Level	Description	Enable	Read Only	Disable
Services	Specifies whether the administrator can view, add, or delete service objects for use in creating policy rules that limit the port numbers an application can use.	Yes	Yes	Yes
Service Groups	Specifies whether the administrator can view, add, or delete service group objects for use in security policy.	Yes	Yes	Yes
Tags	Specifies whether the administrator can view, add, or delete tags that have been defined on the firewall.	Yes	Yes	Yes
GlobalProtect	Specifies whether the administrator can view, add, or delete HIP objects and profiles. You can restrict access to both types of objects at the GlobalProtect level, or provide more granular control by enabling the GlobalProtect privilege and restricting HIP Object or HIP Profile access.	Yes	No	Yes
HIP Objects	Specifies whether the administrator can view, add, or delete HIP objects, which are used to define HIP profiles. HIP Objects also generate HIP Match logs.	Yes	Yes	Yes
Clientless Apps	Specifies whether the administrator can view, add, modify, or delete GlobalProtect VPN Clientless applications.	Yes	Yes	Yes
Clientless App Groups	Specifies whether the administrator can view, add, modify, or delete GlobalProtect VPN Clientless application groups.	Yes	Yes	Yes
HIP Profiles	Specifies whether the administrator can view, add, or delete HIP Profiles for use in security policy and/or for generating HIP Match logs.	Yes	Yes	Yes
External Dynamic Lists	Specifies whether the administrator can view, add, or delete external dynamic lists for use in security policy.	Yes	Yes	Yes
Custom Objects	Specifies whether the administrator can see the custom spyware and vulnerability	Yes	No	Yes

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Access Level	Description	Enable	Read Only	Disable
	signatures. You can restrict access to either enable or disable access to all custom signatures at this level, or provide more granular control by enabling the Custom Objects privilege and then restricting access to each type of signature.			
Data Patterns	Specifies whether the administrator can view, add, or delete custom data pattern signatures for use in creating custom Vulnerability Protection profiles.	Yes	Yes	Yes
Spyware	Specifies whether the administrator can view, add, or delete custom spyware signatures for use in creating custom Vulnerability Protection profiles.	Yes	Yes	Yes
Vulnerability	Specifies whether the administrator can view, add, or delete custom vulnerability signatures for use in creating custom Vulnerability Protection profiles.	Yes	Yes	Yes
URL Category	Specifies whether the administrator can view, add, or delete custom URL categories for use in policy.	Yes	Yes	Yes
Security Profiles	Specifies whether the administrator can see security profiles. You can restrict access to either enable or disable access to all security profiles at this level, or provide more granular control by enabling the Security Profiles privilege and then restricting access to each type of profile.	Yes	No	Yes
Antivirus	Specifies whether the administrator can view, add, or delete antivirus profiles.	Yes	Yes	Yes
Anti-Spyware	Specifies whether the administrator can view, add, or delete Anti-Spyware profiles.	Yes	Yes	Yes
Vulnerability Protection	Specifies whether the administrator can view, add, or delete Vulnerability Protection profiles.	Yes	Yes	Yes

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Access Level	Description	Enable	Read Only	Disable
URL Filtering	Specifies whether the administrator can view, add, or delete URL filtering profiles.	Yes	Yes	Yes
File Blocking	Specifies whether the administrator can view, add, or delete file blocking profiles.	Yes	Yes	Yes
WildFire Analysis	Specifies whether the administrator can view, add, or delete WildFire analysis profiles.	Yes	Yes	Yes
Data Filtering	Specifies whether the administrator can view, add, or delete data filtering profiles.	Yes	Yes	Yes
DoS Protection	Specifies whether the administrator can view, add, or delete DoS protection profiles.	Yes	Yes	Yes
GTP Protection	Specifies whether the mobile network operator can view, add, or delete GTP Protection profiles.	Yes	Yes	Yes
SCTP Protection	Specifies whether the mobile network operator can view, add, or delete Stream Control Transmission Protocol (SCTP) Protection profiles.	Yes	Yes	Yes
Security Profile Groups	Specifies whether the administrator can view, add, or delete security profile groups.	Yes	Yes	Yes
Log Forwarding	Specifies whether the administrator can view, add, or delete log forwarding profiles.	Yes	Yes	Yes
Authentication	Specifies whether the administrator can view, add, or delete authentication enforcement objects.	Yes	Yes	Yes
Decryption Profile	Specifies whether the administrator can view, add, or delete decryption profiles.	Yes	Yes	Yes
SD-WAN Link Management	Specifies whether the administrator can add or delete Path Quality, SaaS Quality, Traffic Distribution, and Error Correction profiles.	Yes	No	Yes

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Access Level	Description	Enable	Read Only	Disable
Path Quality Profile	Specifies whether the administrator can view, add, or delete SD-WAN Path Quality profiles.	Yes	Yes	Yes
SaaS Quality Profile	Specifies whether the administrator can view, add, or delete SD-WAN SaaS Quality profiles.	Yes	Yes	Yes
Traffic Distribution Profile	Specifies whether the administrator can view, add, or delete SD-WAN Traffic Distribution profiles.	Yes	Yes	Yes
Error Correction Profile	Specifies whether the administrator can view, add, or delete SD-WAN Error Correction profiles.	Yes	Yes	Yes
Packet Broker Profile	Specifies whether the administrator can view, add, or delete Packet Broker profiles.	Yes	Yes	Yes
Schedules	Specifies whether the administrator can view, add, or delete schedules for limiting a security policy to a specific date and/or time range.	Yes	Yes	Yes

Provide Granular Access to the Network Tab

When deciding whether to allow access to the **Network** tab as a whole, determine whether the administrator will have network administration responsibilities, including GlobalProtect administration. If not, the administrator probably does not need access to the tab.

You can also define access to the **Network** tab at the node level. By enabling access to a specific node, you give the administrator the privilege to view, add, and delete the corresponding network configurations. Giving read-only access allows the administrator to view the already-defined configuration, but not create or delete any. Disabling a node prevents the administrator from seeing the node in the web interface.

Access Level	Description	Enable	Read Only	Disable
Interfaces	Specifies whether the administrator can view, add, or delete interface configurations.	Yes	Yes	Yes
Zones	Specifies whether the administrator can view, add, or delete zones.	Yes	Yes	Yes

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Access Level	Description	Enable	Read Only	Disable
VLANs	Specifies whether the administrator can view, add, or delete VLANs.	Yes	Yes	Yes
Virtual Wires	Specifies whether the administrator can view, add, or delete virtual wires.	Yes	Yes	Yes
Virtual Routers	Specifies whether the administrator can view, add, modify or delete virtual routers.	Yes	Yes	Yes
IPSec Tunnels	Specifies whether the administrator can view, add, modify, or delete IPSec Tunnel configurations.	Yes	Yes	Yes
GRE Tunnels	Specifies whether the administrator can view, add, modify, or delete GRE Tunnel configurations.	Yes	Yes	Yes
DHCP	Specifies whether the administrator can view, add, modify, or delete DHCP server and DHCP relay configurations.	Yes	Yes	Yes
DNS Proxy	Specifies whether the administrator can view, add, modify, or delete DNS proxy configurations.	Yes	Yes	Yes
GlobalProtect	Specifies whether the administrator can view, add, modify GlobalProtect portal and gateway configurations. You can disable access to the GlobalProtect functions entirely, or you can enable the GlobalProtect privilege and then restrict the role to either the portal or gateway configuration areas.	Yes	No	Yes
Portals	Specifies whether the administrator can view, add, modify, or delete GlobalProtect portal configurations.	Yes	Yes	Yes
Gateways	Specifies whether the administrator can view, add, modify, or delete GlobalProtect gateway configurations.	Yes	Yes	Yes
MDM	Specifies whether the administrator can view, add, modify, or delete GlobalProtect MDM server configurations.	Yes	Yes	Yes

https://t.me/learningnets

Access Level	Description	Enable	Read Only	Disable
Device Block List	Specifies whether the administrator can view, add, modify, or delete device block lists.	Yes	Yes	Yes
Clientless Apps	Specifies whether the administrator can view, add, modify, or delete GlobalProtect Clientless VPN applications.	Yes	Yes	Yes
Clientless App Groups	Specifies whether the administrator can view, add, modify, or delete GlobalProtect Clientless VPN application groups.	Yes	Yes	Yes
QoS	Specifies whether the administrator can view, add, modify, or delete QoS configurations.	Yes	Yes	Yes
LLDP	Specifies whether the administrator can view add, modify, or delete LLDP configurations.	Yes	Yes	Yes
Network Profiles	Sets the default state to enable or disable for all of the Network settings described below.	Yes	No	Yes
GlobalProtect IPSec Crypto	Controls access to the **Network Profiles** > **GlobalProtect IPSec Crypto** node. If you disable this privilege, the administrator will not see that node, or configure algorithms for authentication and encryption in VPN tunnels between a GlobalProtect gateway and clients. If you set the privilege to read-only, the administrator can view existing GlobalProtect IPSec Crypto profiles but cannot add or edit them.	Yes	Yes	Yes
IKE Gateways	Controls access to the **Network Profiles** > **IKE Gateways** node. If you disable this privilege, the administrator will not see the **IKE Gateways** node or define gateways that include the configuration information necessary to perform IKE protocol negotiation with peer gateway.	Yes	Yes	Yes

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Access Level	Description	Enable	Read Only	Disable
	If the privilege state is set to read-only, you can view the currently configured IKE Gateways but cannot add or edit gateways.			
IPSec Crypto	Controls access to the **Network Profiles** > **IPSec Crypto** node. If you disable this privilege, the administrator will not see the **Network Profiles** > **IPSec Crypto** node or specify protocols and algorithms for identification, authentication, and encryption in VPN tunnels based on IPSec SA negotiation. If the privilege state is set to read-only, you can view the currently configured IPSec Crypto configuration but cannot add or edit a configuration.	Yes	Yes	Yes
IKE Crypto	Controls how devices exchange information to ensure secure communication. Specify the protocols and algorithms for identification, authentication, and encryption in VPN tunnels based on IPsec SA negotiation (IKEv1 Phase-1).	Yes	Yes	Yes
Monitor	Controls access to the **Network Profiles** > **Monitor** node. If you disable this privilege, the administrator will not see the **Network Profiles** > **Monitor** node or be able to create or edit a monitor profile that is used to monitor IPSec tunnels and monitor a next-hop device for policy-based forwarding (PBF) rules. If the privilege state is set to read-only, you can view the currently configured monitor profile configuration but cannot add or edit a configuration.	Yes	Yes	Yes
Interface Mgmt	Controls access to the **Network Profiles** > **Interface Mgmt** node. If you disable this privilege, the administrator will not see the **Network Profiles** > **Interface Mgmt** node or be able to specify the protocols that are used to manage the firewall.	Yes	Yes	Yes

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Access Level	Description	Enable	Read Only	Disable
	If the privilege state is set to read-only, you can view the currently configured Interface management profile configuration but cannot add or edit a configuration.			
Zone Protection	Controls access to the **Network Profiles** > **Zone Protection** node. If you disable this privilege, the administrator will not see the **Network Profiles** > **Zone Protection** node or be able to configure a profile that determines how the firewall responds to attacks from specified security zones. If the privilege state is set to read-only, you can view the currently configured Zone Protection profile configuration but cannot add or edit a configuration.	Yes	Yes	Yes
QoS Profile	Controls access to the **Network Profiles** > **QoS** node. If you disable this privilege, the administrator will not see the **Network Profiles** > **QoS** node or be able to configure a QoS profile that determines how QoS traffic classes are treated. If the privilege state is set to read-only, you can view the currently configured QoS profile configuration but cannot add or edit a configuration.	Yes	Yes	Yes
LLDP Profile	Controls access to the **Network Profiles** > **LLDP** node. If you disable this privilege, the administrator will not see the **Network Profiles** > **LLDP** node or be able to configure an LLDP profile that controls whether the interfaces on the firewall can participate in the Link Layer Discovery Protocol. If the privilege state is set to read-only, you can view the currently configured LLDP profile configuration but cannot add or edit a configuration.	Yes	Yes	Yes
BFD Profile	Controls access to the **Network Profiles** > **BFD Profile** node. If you disable this privilege, the administrator will not see	Yes	Yes	Yes

https://t.me/learningnets

Access Level	Description	Enable	Read Only	Disable
	the **Network Profiles** > **BFD Profile** node or be able to configure a BFD profile. A Bidirectional Forwarding Detection (BFD) profile allows you to configure BFD settings to apply to one or more static routes or routing protocols. Thus, BFD detects a failed link or BFD peer and allows an extremely fast failover. If the privilege state is set to read-only, you can view the currently configured BFD profile but cannot add or edit a BFD profile.			
SD-WAN Interface Profile	Controls access to the **SD-WAN Interface Profile** node. If you disable this privilege, the administrator will not see the **SD-WAN Interface Profile** node or be able to configure an SD-WAN Interface Profile. An SD-WAN Interface Profile defines the characteristics of ISP connections and specifies the link speed and how frequently the firewall monitors the link. If the privilege state is set to read-only, you can view the currently configured SD-WAN Interface Profile but cannot add or edit one.	Yes	Yes	Yes

Provide Granular Access to the Device Tab

To define granular access privileges for the **Device** tab, when creating or editing an admin role profile (**Device** > **Admin Roles**), scroll down to the **Device** node on the **WebUI** tab.

Access Level	Description	Enable	Read Only	Disable
Setup	Controls access to the **Setup** node. If you disable this privilege, the administrator will not see the **Setup** node or have access to firewall-wide setup configuration information, such as Management, Operations, Service, Content-ID, WildFire or Session setup information.	Yes	Yes	Yes

https://t.me/learningnets

Access Level	Description	Enable	Read Only	Disable
	If the privilege state is set to read-only, you can view the current configuration but cannot make any changes.			
Management	Controls access to the **Management** node. If you disable this privilege, the administrator will not be able to configure settings such as the hostname, domain, timezone, authentication, logging and reporting, Panorama connections, banner, message, and password complexity settings, and more. If the privilege state is set to read-only, you can view the current configuration but cannot make any changes.	Yes	Yes	Yes
Operations	Controls access to the **Operations** and **Telemetry and Threat Intelligence** nodes. If you disable this privilege, the administrator cannot: • Load firewall configurations. • Save or revert the firewall configuration. *This privilege applies only to the **Device** > **Operations** options. The Save and Commit privileges control whether the administrator can save or revert configurations through the **Config** > **Save** and **Config** > **Revert** options.* • Create custom logos. • Configure SNMP monitoring of firewall settings. • Configure the Statistics Service feature. • Configure **Telemetry and Threat Intelligence** settings. Only administrators with the predefined Superuser role can export or import	Yes	Yes	Yes

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Access Level	Description	Enable	Read Only	Disable
	firewall configurations and shut down the firewall. Only administrators with the predefined Superuser or Device Administrator role can reboot the firewall or restart the dataplane. Administrators with a role that allows access only to specific virtual systems cannot load, save, or revert firewall configurations through the **Device** > **Operations** options.			
Services	Controls access to the **Services** node. If you disable this privilege, the administrator will not be able to configure services for DNS servers, an update server, proxy server, or NTP servers, or set up service routes. If the privilege state is set to read-only, you can view the current configuration but cannot make any changes.	Yes	Yes	Yes
Content-ID	Controls access to the **Content-ID** node. If you disable this privilege, the administrator will not be able to configure URL filtering or Content-ID. If the privilege state is set to read-only, you can view the current configuration but cannot make any changes.	Yes	Yes	Yes
WildFire	Controls access to the **WildFire** node. If you disable this privilege, the administrator will not be able to configure WildFire settings. If the privilege state is set to read-only, you can view the current configuration but cannot make any changes.	Yes	Yes	Yes
Session	Controls access to the **Session** node. If you disable this privilege, the administrator will not be able to configure session settings or timeouts for TCP, UDP or ICMP, or configure decryption or VPN session settings.	Yes	Yes	Yes

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Access Level	Description	Enable	Read Only	Disable
	If the privilege state is set to read-only, you can view the current configuration but cannot make any changes.			
HSM	Controls access to the **HSM** node. If you disable this privilege, the administrator will not be able to configure a Hardware Security Module. If the privilege state is set to read-only, you can view the current configuration but cannot make any changes.	Yes	Yes	Yes
High Availability	Controls access to the **High Availability** node. If you disable this privilege, the administrator will not see the **High Availability** node or have access to firewall-wide high availability configuration information such as General setup information or Link and Path Monitoring. If you set this privilege to read-only, the administrator can view High Availability configuration information for the firewall but is not allowed to perform any configuration procedures.	Yes	Yes	Yes
Config Audit	Controls access to the **Config Audit** node. If you disable this privilege, the administrator will not see the **Config Audit** node or have access to any firewall-wide configuration information.	Yes	No	Yes
Administrators	Controls access to the **Administrators** node. This function can only be allowed for read-only access. If you disable this privilege, the administrator will not see the **Administrators** node or have access to information about their own administrator account. If you set this privilege to read-only, the administrator can view the configuration information for their own administrator account. They will not see any information	No	Yes	Yes

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Access Level	Description	Enable	Read Only	Disable
	about other administrator accounts configured on the firewall.			
Admin Roles	Controls access to the **Admin Roles** node. This function can only be allowed for read-only access. If you disable this privilege, the administrator will not see the **Admin Roles** node or have access to any firewall-wide information concerning Admin Role profiles configuration. If you set this privilege to read-only, you can view the configuration information for all administrator roles configured on the firewall.	No	Yes	Yes
Authentication Profile	Controls access to the **Authentication Profile** node. If you disable this privilege, the administrator will not see the **Authentication Profile** node or be able to create or edit authentication profiles that specify RADIUS, TACACS +, LDAP, Kerberos, SAML, multi-factor authentication (MFA), or local database authentication settings. PAN-OS uses authentication profiles to authenticate firewall administrators and Authentication Portal or GlobalProtect end users. If you set this privilege to read-only, the administrator can view the **Authentication Profile** information but cannot create or edit authentication profiles.	Yes	Yes	Yes
Authentication Sequence	Controls access to the **Authentication Sequence** node. If you disable this privilege, the administrator will not see the **Authentication Sequence** node or be able to create or edit an authentication sequence. If you set this privilege to read-only, the administrator can view the **Authentication Profile** information but cannot create or edit an authentication sequence.	Yes	Yes	Yes

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Access Level	Description	Enable	Read Only	Disable
Virtual Systems	Controls access to the **Virtual Systems** node. If you disable this privilege, the administrator will not see or be able to configure virtual systems. If the privilege state is set to read-only, you can view the currently configured virtual systems but cannot add or edit a configuration.	Yes	Yes	Yes
Shared Gateways	Controls access to the **Shared Gateways** node. Shared gateways allow virtual systems to share a common interface for external communications. If you disable this privilege, the administrator will not see or be able to configure shared gateways. If the privilege state is set to read-only, you can view the currently configured shared gateways but cannot add or edit a configuration.	Yes	Yes	Yes
User Identification	Controls access to the **User Identification** node. If you disable this privilege, the administrator will not see the **User Identification** node or have access to firewall-wide User Identification configuration information, such as User Mapping, Connection Security, User-ID Agents, Terminal Server Agents, Group Mappings Settings, or Authentication Portal Settings. If you set this privilege to read-only, the administrator can view configuration information for the firewall but is not allowed to perform any configuration procedures.	Yes	Yes	Yes
VM Information Source	Controls access to the **VM Information Source** node that allows you to configure the firewall/Windows User-ID agent to collect VM inventory automatically. If you disable this privilege, the administrator will not see the **VM Information Source** node.	Yes	Yes	Yes

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Access Level	Description	Enable	Read Only	Disable
	If you set this privilege to read-only, the administrator can view the VM information sources configured but cannot add, edit, or delete any sources. *This privilege is not available to Device Group and Template administrators.*			
Certificate Management	Sets the default state to enable or disable for all of the Certificate settings described below.	Yes	No	Yes
Certificates	Controls access to the **Certificates** node. If you disable this privilege, the administrator will not see the **Certificates** node or be able to configure or access information regarding Device Certificates or Default Trusted Certificate Authorities. If you set this privilege to read-only, the administrator can view Certificate configuration information for the firewall but is not allowed to perform any configuration procedures.	Yes	Yes	Yes
Certificate Profile	Controls access to the **Certificate Profile** node. If you disable this privilege, the administrator will not see the **Certificate Profile** node or be able to create certificate profiles. If you set this privilege to read-only, the administrator can view Certificate Profiles that are currently configured for the firewall but is not allowed to create or edit a certificate profile.	Yes	Yes	Yes
OCSP Responder	Controls access to the **OCSP Responder** node. If you disable this privilege, the administrator will not see the **OCSP Responder** node or be able to define a server that will be used to verify the revocation status of certificates issues by the firewall. If you set this privilege to read-only, the administrator can view the **OCSP**	Yes	Yes	Yes

https://t.me/learningnets

Access Level	Description	Enable	Read Only	Disable
	Responder configuration for the firewall but is not allowed to create or edit an OCSP responder configuration.			
SSL/TLS Service Profile	Controls access to the **SSL/TLS Service Profile** node. If you disable this privilege, the administrator will not see the node or configure a profile that specifies a certificate and a protocol version or range of versions for firewall services that use SSL/TLS. If you set this privilege to read-only, the administrator can view existing SSL/TLS Service profiles but cannot create or edit them.	Yes	Yes	Yes
SCEP	Controls access to the **SCEP** node. If you disable this privilege, the administrator will not see the node or be able to define a profile that specifies simple certificate enrollment protocol (SCEP) settings for issuing unique device certificates. If you set this privilege to read-only, the administrator can view existing SCEP profiles but cannot create or edit them.	Yes	Yes	Yes
SSL Decryption Exclusion	Controls access to the **SSL Decryption Exclusion** node. If you disable this privilege, the administrator will not see the node or be able to add custom exclusions. If you set this privilege to read-only, the administrator can view existing SSL decryption exceptions but cannot create or edit them.	Yes	Yes	Yes
SSH Service Profile	Controls access to the **SSH Service Profile** node. If you disable this privilege, the administrator will be unable to see the node or configure a profile to specify parameters for SSH connections to your Palo Alto Networks management and high availability (HA) appliances.	Yes	Yes	Yes

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Access Level	Description	Enable	Read Only	Disable
	If you set this privilege to read-only, the administrator can view existing SSH service profiles but cannot edit or create them.			
Response Pages	Controls access to the **Response Pages** node. If you disable this privilege, the administrator will not see the **Response Page** node or be able to define a custom HTML message that is downloaded and displayed instead of a requested web page or file. If you set this privilege to read-only, the administrator can view the **Response Page** configuration for the firewall but is not allowed to create or edit a response page configuration.	Yes	Yes	Yes
Log Settings	Sets the default state to enable or disable for all of the Log settings described below.	Yes	No	Yes
System	Controls access to the **Log Settings** > **System** node. If you disable this privilege, the administrator cannot see the **Log Settings** > **System** node or specify which System logs the firewall forwards to Panorama or external services (such as a syslog server). If you set this privilege to read-only, the administrator can view the **Log Settings** > **System** settings for the firewall but cannot add, edit, or delete the settings.	Yes	Yes	Yes
Configuration	Controls access to the **Log Settings** > **Configuration** node. If you disable this privilege, the administrator cannot see the **Log Settings** > **Configuration** node or specify which Configuration logs the firewall forwards to Panorama or external services (such as a syslog server). If you set this privilege to read-only, the administrator can view the **Log Settings** > **Configuration** settings for the firewall but cannot add, edit, or delete the settings.	Yes	Yes	Yes

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Access Level	Description	Enable	Read Only	Disable
User-ID	Controls access to the **Log Settings** > **User-ID** node. If you disable this privilege, the administrator cannot see the **Log Settings** > **User-ID** node or specify which User-ID logs the firewall forwards to Panorama or external services (such as a syslog server). If you set this privilege to read-only, the administrator can view the **Log Settings** > **User-ID** settings for the firewall but cannot add, edit, or delete the settings.	Yes	Yes	Yes
HIP Match	Controls access to the **Log Settings** > **HIP Match** node. If you disable this privilege, the administrator cannot see the **Log Settings** > **HIP Match** node or specify which Host Information Profile (HIP) match logs the firewall forwards to Panorama or external services (such as a syslog server). HIP match logs provide information on Security policy rules that apply to GlobalProtect endpoints. If you set this privilege to read-only, the administrator can view the **Log Settings** > **HIP** settings for the firewall but cannot add, edit, or delete the settings.	Yes	Yes	Yes
GlobalProtect	Controls access to the **Log Settings** > **GlobalProtect** node. If you disable this privilege, the administrator cannot see the **Log Settings** > **GlobalProtect** node or specify which GlobalProtect logs the firewall forwards to Panorama or external services (such as a syslog server). If you set this privilege to read-only, the administrator can view the **Log Settings** > **GlobalProtect** settings for the firewall but cannot add, edit, or delete the settings.	Yes	Yes	Yes
Correlation	Controls access to the **Log Settings** > **Correlation** node. If you disable this privilege, the administrator cannot see the **Log Settings** > **Correlation** node or add, delete, or modify correlation log	Yes	Yes	Yes

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Access Level	Description	Enable	Read Only	Disable
	forwarding settings or tag source or destination IP addresses. If you set this privilege to read-only, the administrator can view the **Log Settings** > **Correlation** settings for the firewall but cannot add, edit, or delete the settings.			
Alarm Settings	Controls access to the **Log Settings** > **Alarm Settings** node. If you disable this privilege, the administrator cannot see the **Log Settings** > **Alarm Settings** node or configure notifications that the firewall generates when a Security policy rule (or group of rules) is hit repeatedly within a configurable time period. If you set this privilege to read-only, the administrator can view the **Log Settings** > **Alarm Settings** for the firewall but cannot edit the settings.	Yes	Yes	Yes
Manage Logs	Controls access to the **Log Settings** > **Manage Logs** node. If you disable this privilege, the administrator cannot see the **Log Settings** > **Manage Logs** node or clear the indicated logs. If you set this privilege to read-only, the administrator can view the **Log Settings** > **Manage Logs** information but cannot clear any of the logs.	Yes	Yes	Yes
Server Profiles	Sets the default state to enable or disable for all of the Server Profiles settings described below.	Yes	No	Yes
SNMP Trap	Controls access to the **Server Profiles** > **SNMP Trap** node. If you disable this privilege, the administrator will not see the **Server Profiles** > **SNMP Trap** node or be able to specify one or more SNMP trap destinations to be used for system log entries. If you set this privilege to read-only, the administrator can view the **Server Profiles**	Yes	Yes	Yes

https://t.me/learningnets

Access Level	Description	Enable	Read Only	Disable
	> **SNMP Trap Logs** information but cannot specify SNMP trap destinations.			
Syslog	Controls access to the **Server Profiles** > **Syslog** node. If you disable this privilege, the administrator will not see the **Server Profiles** > **Syslog** node or be able to specify one or more syslog servers. If you set this privilege to read-only, the administrator can view the **Server Profiles** > **Syslog** information but cannot specify syslog servers.	Yes	Yes	Yes
Email	Controls access to the **Server Profiles** > **Email** node. If you disable this privilege, the administrator will not see the **Server Profiles** > **Email** node or be able to configure an email profile that can be used to enable email notification for system and configuration log entries. If you set this privilege to read-only, the administrator can view the **Server Profiles** > **Email** information but cannot configure an email server profile.	Yes	Yes	Yes
HTTP	Controls access to the **Server Profiles** > **HTTP** node. If you disable this privilege, the administrator will not see the **Server Profiles** > **HTTP** node or be able to configure an HTTP server profile that can be used to enable log forwarding to HTTP destinations any log entries. If you set this privilege to read-only, the administrator can view the **Server Profiles** > **HTTP** information but cannot configure an HTTP server profile.	Yes	Yes	Yes
Netflow	Controls access to the **Server Profiles** > **Netflow** node. If you disable this privilege, the administrator will not see the **Server Profiles** > **Netflow** node or be able to define a NetFlow server profile, which specifies the frequency of the export along with the NetFlow servers that will receive the exported data.	Yes	Yes	Yes

https://t.me/learningnets

Access Level	Description	Enable	Read Only	Disable
	If you set this privilege to read-only, the administrator can view the **Server Profiles** > **Netflow** information but cannot define a Netflow profile.			
RADIUS	Controls access to the **Server Profiles** > **RADIUS** node. If you disable this privilege, the administrator will not see the **Server Profiles** > **RADIUS** node or be able to configure settings for the RADIUS servers that are identified in authentication profiles. If you set this privilege to read-only, the administrator can view the **Server Profiles** > **RADIUS** information but cannot configure settings for the RADIUS servers.	Yes	Yes	Yes
TACACS+	Controls access to the **Server Profiles** > **TACACS+** node. If you disable this privilege, the administrator will not see the node or configure settings for the TACACS + servers that authentication profiles reference. If you set this privilege to read-only, the administrator can view existing TACACS + server profiles but cannot add or edit them.	Yes	Yes	Yes
LDAP	Controls access to the **Server Profiles** > **LDAP** node. If you disable this privilege, the administrator will not see the **Server Profiles** > **LDAP** node or be able to configure settings for the LDAP servers to use for authentication by way of authentication profiles. If you set this privilege to read-only, the administrator can view the **Server Profiles** > **LDAP** information but cannot configure settings for the LDAP servers.	Yes	Yes	Yes
Kerberos	Controls access to the **Server Profiles** > **Kerberos** node. If you disable this privilege, the administrator will not see the **Server Profiles** > **Kerberos** node or	Yes	Yes	Yes

https://t.me/learningnets

Access Level	Description	Enable	Read Only	Disable
	configure a Kerberos server that allows users to authenticate natively to a domain controller. If you set this privilege to read-only, the administrator can view the **Server Profiles** > **Kerberos** information but cannot configure settings for Kerberos servers.			
SAML Identity Provider	Controls access to the **Server Profiles** > **SAML Identity Provider** node. If you disable this privilege, the administrator cannot see the node or configure SAML identity provider (IdP) server profiles. If you set this privilege to read-only, the administrator can view the **Server Profiles** > **SAML Identity Provider** information but cannot configure SAML IdP server profiles.	Yes	Yes	Yes
Multi Factor Authentication	Controls access to the **Server Profiles** > **Multi Factor Authentication** node. If you disable this privilege, the administrator cannot see the node or configure multi-factor authentication (MFA) server profiles. If you set this privilege to read-only, the administrator can view the **Server Profiles** > **SAML Identity Provider** information but cannot configure MFA server profiles.			
Local User Database	Sets the default state to enable or disable for all of the Local User Database settings described below.	Yes	No	Yes
Users	Controls access to the **Local User Database** > **Users** node. If you disable this privilege, the administrator will not see the **Local User Database** > **Users** node or set up a local database on the firewall to store authentication information for remote access users, firewall administrators, and Authentication Portal users. If you set this privilege to read-only, the administrator can view the **Local User Database** > **Users** information but cannot	Yes	Yes	Yes

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Access Level	Description	Enable	Read Only	Disable
	set up a local database on the firewall to store authentication information.			
User Groups	Controls access to the **Local User Database** > **Users** node. If you disable this privilege, the administrator will not see the **Local User Database** > **Users** node or be able to add user group information to the local database. If you set this privilege to read-only, the administrator can view the **Local User Database** > **Users** information but cannot add user group information to the local database.	Yes	Yes	Yes
Access Domain	Controls access to the **Access Domain** node. If you disable this privilege, the administrator will not see the **Access Domain** node or be able to create or edit an access domain. If you set this privilege to read-only, the administrator can view the **Access Domain** information but cannot create or edit an access domain.	Yes	Yes	Yes
Scheduled Log Export	Controls access to the **Scheduled Log Export** node. If you disable this privilege, the administrator will not see the **Scheduled Log Export** node or be able schedule exports of logs and save them to a File Transfer Protocol (FTP) server in CSV format or use Secure Copy (SCP) to securely transfer data between the firewall and a remote host. If you set this privilege to read-only, the administrator can view the **Scheduled Log Export Profile** information but cannot schedule the export of logs.	Yes	No	Yes
Software	Controls access to the **Software** node. If you disable this privilege, the administrator will not see the **Software** node or view the latest versions of the PAN-OS software available from Palo Alto Networks, read the release notes for each	Yes	Yes	Yes

https://t.me/learningnets

Access Level	Description	Enable	Read Only	Disable
	version, and select a release to download and install. If you set this privilege to read-only, the administrator can view the **Software** information but cannot download or install software.			
GlobalProtect Client	Controls access to the **GlobalProtect Client** node. If you disable this privilege, the administrator will not see the **GlobalProtect Client** node or view available GlobalProtect releases, download the code or activate the GlobalProtect app. If you set this privilege to read-only, the administrator can view the available **GlobalProtect Client** releases but cannot download or install the app software.	Yes	Yes	Yes
Dynamic Updates	Controls access to the **Dynamic Updates** node. If you disable this privilege, the administrator will not see the **Dynamic Updates** node or be able to view the latest updates, read the release notes for each update, or select an update to upload and install. If you set this privilege to read-only, the administrator can view the available **Dynamic Updates** releases, read the release notes but cannot upload or install the software.	Yes	Yes	Yes
Licenses	Controls access to the **Licenses** node. If you disable this privilege, the administrator will not see the **Licenses** node or be able to view the licenses installed or activate licenses. If you set this privilege to read-only, the administrator can view the installed **Licenses**, but cannot perform license management functions.	Yes	Yes	Yes
Support	Controls access to the **Support** node. If you disable this privilege, the administrator cannot see the **Support**	Yes	Yes	Yes

https://t.me/learningnets

Access Level	Description	Enable	Read Only	Disable
	node, activate support, or access production and security alerts from Palo Alto Networks. If you set this privilege to read-only, the administrator can see the **Support** node and access production and security alerts but cannot activate support. Only administrators with the predefined Superuser role can use the **Support** node to generate tech support files or generate and download stats dump and core files.			
Master Key and Diagnostics	Controls access to the **Master Key and Diagnostics** node. If you disable this privilege, the administrator will not see the **Master Key and Diagnostics** node or be able to specify a master key to encrypt private keys on the firewall. If you set this privilege to read-only, the administrator can view the **Master Key and Diagnostics** node and view information about master keys that have been specified but cannot add or edit a new master key configuration.	Yes	Yes	Yes
Policy Recommendation	Controls access to **IoT** and **SaaS** policy rule recommendations. If you disable these privileges, the administrator can't see the **Policy Recommendation** > **IoT** node, the **Policy Recommendation** > **SaaS** node, or both, depending on which privileges you disable. If you set these privileges to read-only, the administrator can view the nodes but cannot import policy rules or edit information.	Yes	Yes	Yes

Define User Privacy Settings in the Admin Role Profile

To define what private end user data an administrator has access to, when creating or editing an admin role profile (**Device** > **Admin Roles**), scroll down to the **Privacy** option on the **WebUI** tab.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Access Level	Description	Enable	Read Only	Disable
Privacy	Sets the default state to enable or disable for all of the privacy settings described below.	Yes	N/A	Yes
Show Full IP addresses	When disabled, full IP addresses obtained by traffic running through the Palo Alto firewall are not shown in logs or reports. In place of the IP addresses that are normally displayed, the relevant subnet is displayed. *Scheduled reports that are displayed in the interface through **Monitor** > **Reports** and reports that are sent via scheduled emails will still display full IP addresses. Because of this exception, we recommend that the following settings within the **Monitor** tab be set to disable: Custom Reports, Application Reports, Threat Reports, URL Filtering Reports, Traffic Reports and Email Scheduler.*	Yes	N/A	Yes
Show User Names in Logs and Reports	When disabled, usernames obtained by traffic running through the Palo Alto Networks firewall are not shown in logs or reports. Columns where the usernames would normally be displayed are empty.	Yes	N/A	Yes

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Access Level	Description	Enable	Read Only	Disable
	*Scheduled reports that are displayed in the interface through **Monitor** > **Reports** or reports that are sent via the email scheduler will still display usernames. Because of this exception, we recommend that the following settings within the Monitor tab be set to disable: Custom Reports, Application Reports, Threat Reports, URL Filtering Reports, Traffic Reports and Email Scheduler.*			
View PCAP Files	When disabled, packet capture files that are normally available within the Traffic, Threat and Data Filtering logs are not displayed.	Yes	N/A	Yes

Restrict Administrator Access to Commit and Validate Functions

To restrict access to commit (and revert), save, and validate functions when creating or editing an Admin Role profile (**Device** > **Admin Roles**), scroll down to the **Commit**, **Save**, and **Validate** options on the **WebUI** tab.

Access Level	Description	Enable	Read Only	Disable
Commit	Sets the default state to enabled or disabled for all of the commit and revert privileges described below.	Yes	N/A	Yes
Device	When disabled, an administrator cannot commit or revert changes that any administrator made to the firewall configuration, including his or her own changes.	Yes	N/A	Yes
Commit For Other Admins	When disabled, an administrator cannot commit or revert changes that other administrators made to the firewall configuration.	Yes	N/A	Yes

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Access Level	Description	Enable	Read Only	Disable
Save	Sets the default state to enabled or disabled for all of the save operation privileges described below.	Yes	N/A	Yes
Partial save	When disabled, an administrator cannot save changes that any administrator made to the firewall configuration, including his or her own changes.	Yes	N/A	Yes
Save For Other Admins	When disabled, an administrator cannot save changes that other administrators made to the firewall configuration.	Yes	N/A	Yes
Validate	When disabled, an administrator cannot validate a configuration.	Yes	N/A	Yes

Provide Granular Access to Global Settings

To define what global settings and administrator has access to, when creating or editing an admin role profile (**Device** > **Admin Roles**), scroll down to the **Global** option on the **WebUI** tab.

Access Level	Description	Enable	Read Only	Disable
Global	Sets the default state to enable or disable for all of the global settings described below. In effect, this setting is only for System Alarms at this time.	Yes	N/A	Yes
System Alarms	When disabled, an administrator cannot view or acknowledge alarms that are generated.	Yes	N/A	Yes

Provide Granular Access to the Panorama Tab

The following table lists the **Panorama** tab access levels and the custom Panorama administrator roles for which they are available. Firewall administrators cannot access any of these privileges.

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
Setup	Specifies whether the administrator can view or edit Panorama setup information,	Panorama: Yes	Yes	Yes	Yes

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
	including **Management**, **Operations and Telemetry**, **Services**, Content-ID, **WildFire**, Session, or **HSM**. If you set the privilege to: • read-only, the administrator can see the information but cannot edit it. • disable this privilege, the administrator cannot see or edit the information.	Device Group/ Template: No			
High Availability	Specifies whether the administrator can view and manage high availability (HA) settings for the Panorama management server. If you set this privilege to read-only, the administrator can view HA configuration information for the Panorama management server but can't manage the configuration. If you disable this privilege, the administrator can't see or manage HA configuration settings for the Panorama management server.	Panorama: Yes Device Group/ Template: No	Yes	Yes	Yes
Config Audit	Specifies whether the administrator can run Panorama configuration audits. If you disable this privilege, the administrator can't run Panorama configuration audits.	Panorama: Yes Device Group/ Template: No	Yes	No	Yes
Administrators	Specifies whether the administrator can view Panorama administrator account details. You can't enable full access to this function: just read-only access. (Only Panorama administrators with a dynamic	Panorama: Yes Device Group/ Template: No	No	Yes	Yes

https://t.me/learningnets

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
	role can add, edit, or delete Panorama administrators.) With read-only access, the administrator can see information about his or her own account but no other Panorama administrator accounts. If you disable this privilege, the administrator can't see information about any Panorama administrator account, including his or her own.				
Admin Roles	Specifies whether the administrator can view Panorama administrator roles. You can't enable full access to this function: just read-only access. (Only Panorama administrators with a dynamic role can add, edit, or delete custom Panorama roles.) With read-only access, the administrator can see Panorama administrator role configurations but can't manage them. If you disable this privilege, the administrator can't see or manage Panorama administrator roles.	Panorama: Yes Device Group/ Template: No	No	Yes	Yes
Access Domain	Specifies whether the administrator can view, add, edit, delete, or clone access domain configurations for Panorama administrators. (This privilege controls access	Panorama: Yes Device Group/ Template: No	Yes	Yes	Yes

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
	only to the configuration of access domains, not access to the device groups, templates, and firewall contexts that are assigned to access domains.) If you set this privilege to read-only, the administrator can view Panorama access domain configurations but can't manage them. If you disable this privilege, the administrator can't see or manage Panorama access domain configurations.	*You assign access domains to Device Group and Template administrators so they can access the configuration and monitoring data within the device groups, templates, and firewall contexts that are assigned to those access domains.*			
Authentication Profile	Specifies whether the administrator can view, add, edit, delete, or clone authentication profiles for Panorama administrators. If you set this privilege to read-only, the administrator can view Panorama authentication profiles but can't manage them. If you disable this privilege, the administrator can't see or manage Panorama authentication profiles.	Panorama: Yes Device Group/ Template: No	Yes	Yes	Yes
Authentication Sequence	Specifies whether the administrator can view,	Panorama: Yes	Yes	Yes	Yes

https://t.me/learningnets

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
	add, edit, delete, or clone authentication sequences for Panorama administrators. If you set this privilege to read-only, the administrator can view Panorama authentication sequences but can't manage them. If you disable this privilege, the administrator can't see or manage Panorama authentication sequences.	Device Group/ Template: No			
User Identification	Specifies whether the administrator can configure User-ID connection security and view, add, edit, or delete data redistribution points (such as User-ID agents). If you set this privilege to read-only, the administrator can view settings for User-ID connection security and redistribution points but can't manage the settings. If you disable this privilege, the administrator can't see or manage settings for User-ID connection security or redistribution points.	Panorama: Yes Device Group/ Template: No	Yes	Yes	Yes
Managed Devices	Specifies whether the administrator can view, add, edit, or delete firewalls as managed devices, and install software or content updates on them. If you set this privilege to read-only, the administrator can see managed firewalls but can't add, delete, tag, or install updates on them. If you disable this privilege, the administrator can't view,	Panorama: Yes Device Group/ Template: Yes	Yes (No for Device Group and Template roles)	Yes	Yes

©2022 Palo Alto Networks, Inc.
https://t.me/learningnets

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
	add, edit, tag, delete, or install updates on managed firewalls. *An administrator with Device Deployment privileges can still select **Panorama** > **Device Deployment** to install updates on managed firewalls.*				
Templates	Specifies whether the administrator can view, edit, add, or delete templates and template stacks. If you set the privilege to read-only, the administrator can see template and stack configurations but can't manage them. If you disable this privilege, the administrator can't see or manage template and stack configurations.	Panorama: Yes Device Group/ Template: Yes *Device Group and Template administrators can see only the templates and stacks that are within the access domains assigned to those administrators.*	Yes (No for Device Group and Template admins)	Yes	Yes
Device Groups	Specifies whether the administrator can view, edit, add, or delete device groups. If you set this privilege to read-only, the administrator can see	Panorama: Yes Device Group/ Template: Yes	Yes	Yes	Yes

https://t.me/learningnets

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
	device group configurations but can't manage them. If you disable this privilege, the administrator can't see or manage device group configurations.	*Device Group and Template administrators can access only the device groups that are within the access domains assigned to those administrators.*			
Managed Collectors	Specifies whether the administrator can view, edit, add, or delete managed collectors. If you set this privilege to read-only, the administrator can see managed collector configurations but can't manage them. If you disable this privilege, the administrator can't view, edit, add, or delete managed collector configurations. *An administrator with Device Deployment privileges can still use the **Panorama** > **Device Deployment** options to install updates on managed collectors.*	Panorama: Yes Device Group/ Template: No	Yes	Yes	Yes

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
Collector Groups	Specifies whether the administrator can view, edit, add, or delete Collector Groups. If you set this privilege to read-only, the administrator can see Collector Groups but can't manage them. If you disable this privilege, the administrator can't see or manage Collector Groups.	Panorama: Yes Device Group/ Template: No	Yes	Yes	Yes
VMware Service Manager	Specifies whether the administrator can view and edit VMware Service Manager settings. If you set this privilege to read-only, the administrator can see the settings but can't perform any related configuration or operational procedures. If you disable this privilege, the administrator can't see the settings or perform any related configuration or operational procedures.	Panorama: Yes Device Group/ Template: No	Yes	Yes	Yes
Certificate Management	Sets the default state, enabled or disabled, for all of the Panorama certificate management privileges.	Panorama: Yes Device Group/ Template: No	Yes	No	Yes
Certificates	Specifies whether the administrator can view, edit, generate, delete, revoke, renew, or export certificates. This privilege also specifies whether the administrator can import or export HA keys. If you set this privilege to read-only, the administrator can see Panorama certificates but can't manage the certificates or HA keys.	Panorama: Yes Device Group/ Template: No	Yes	Yes	Yes

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
	If you disable this privilege, the administrator can't see or manage Panorama certificates or HA keys.				
Certificate Profile	Specifies whether the administrator can view, add, edit, delete or clone Panorama certificate profiles. If you set this privilege to read-only, the administrator can see Panorama certificate profiles but can't manage them. If you disable this privilege, the administrator can't see or manage Panorama certificate profiles.	Panorama: Yes Device Group/ Template: No	Yes	Yes	Yes
SSL/TLS Service Profile	Specifies whether the administrator can view, add, edit, delete or clone SSL/TLS Service profiles. If you set this privilege to read-only, the administrator can see SSL/TLS Service profiles but can't manage them. If you disable this privilege, the administrator can't see or manage SSL/TLS Service profiles.	Panorama: Yes Device Group/ Template: No	Yes	Yes	Yes
Log Settings	Sets the default state, enabled or disabled, for all the log setting privileges.	Panorama: Yes Device Group/ Template: No	Yes	No	Yes
System	Specifies whether the administrator can see and configure the settings that control the forwarding of System logs to external services (syslog, email, SNMP trap, or HTTP servers). If you set this privilege to read-only, the administrator can	Panorama: Yes Device Group/ Template: No	Yes	Yes	Yes

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
	see the System log forwarding settings but can't manage them. If you disable this privilege, the administrator can't see or manage the settings. *This privilege pertains only to System logs that Panorama and Log Collectors generate. The Collector Groups privilege (**Panorama** > **Collector Groups**) controls forwarding for System logs that Log Collectors receive from firewalls. The **Device** > **Log Settings** > System privilege controls log forwarding from firewalls directly to external services (without aggregation on Log Collectors).*				
Config	Specifies whether the administrator can see and configure the settings that control the forwarding of Config logs to external services (syslog, email, SNMP trap, or HTTP servers). If you set this privilege to read-only, the administrator can see the Config log forwarding settings but can't manage them.	Panorama: Yes Device Group/ Template: No	Yes	Yes	Yes

https://t.me/learningnets

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
	If you disable this privilege, the administrator can't see or manage the settings. *This privilege pertains only to Config logs that Panorama and Log Collectors generate. The Collector Groups privilege (**Panorama** > **Collector Groups**) controls forwarding for Config logs that Log Collectors receive from firewalls. The **Device** > **Log Settings** > Configuration privilege controls log forwarding from firewalls directly to external services (without aggregation on Log Collectors).*				
User-ID	Specifies whether the administrator can see and configure the settings that control the forwarding of User-ID logs to external services (syslog, email, SNMP trap, or HTTP servers). If you set this privilege to read-only, the administrator can see the Config log forwarding settings but can't manage them.	Panorama: Yes Device Group/ Template: No	Yes	Yes	Yes

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
	If you disable this privilege, the administrator can't see or manage the settings. *This privilege pertains only to User-ID logs that Panorama generates. The Collector Groups privilege (**Panorama > Collector Groups**) controls forwarding for User-ID logs that Log Collectors receive from firewalls. The **Device > Log Settings** > User-ID privilege controls log forwarding from firewalls directly to external services (without aggregation on Log Collectors).*				
HIP Match	Specifies whether the administrator can see and configure the settings that control the forwarding of HIP Match logs from a Panorama virtual appliance in Legacy mode to external services (syslog, email, SNMP trap, or HTTP servers). If you set this privilege to read-only, the administrator can see the forwarding settings of HIP Match logs but can't manage them.	Panorama: Yes Device Group/ Template: No	Yes	Yes	Yes

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
	If you disable this privilege, the administrator can't see or manage the settings. *The Collector Groups privilege (**Panorama** > **Collector Groups**) controls forwarding for HIP Match logs that Log Collectors receive from firewalls. The **Device** > **Log Settings** > HIP Match privilege controls log forwarding from firewalls directly to external services (without aggregation on Log Collectors).*				
GlobalProtect	Specifies whether the administrator can see and configure the settings that control the forwarding of GlobalProtect logs from a Panorama virtual appliance in Legacy mode to external services (syslog, email, SNMP trap, or HTTP servers). If you set this privilege to read-only, the administrator can see the forwarding settings of GlobalProtect logs but can't manage them. If you disable this privilege, the administrator can't see or manage the settings.	Panorama: Yes Device Group/ Template: No	Yes	Yes	Yes

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
	*The Collector Groups privilege (**Panorama > Collector Groups**) controls forwarding for GlobalProtect logs that Log Collectors receive from firewalls. The **Device > Log Settings > GlobalProtect** privilege controls log forwarding from firewalls directly to external services (without aggregation on Log Collectors).*				
Correlation	Specifies whether the administrator can see and configure the settings that control the forwarding of Correlation logs from a Panorama virtual appliance in Legacy mode to external services (syslog, email, SNMP trap, or HTTP servers). If you set this privilege to read-only, the administrator can see the Correlation log forwarding settings but can't manage them. If you disable this privilege, the administrator can't see or manage the settings.	Panorama: Yes Device Group/ Template: No	Yes	Yes	Yes

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
	*The Collector Groups privilege (**Panorama > Collector Groups**) controls forwarding of Correlation logs from a Panorama M-Series appliance or Panorama virtual appliance in Panorama mode.*				
Traffic	Specifies whether the administrator can see and configure the settings that control the forwarding of Traffic logs from a Panorama virtual appliance in Legacy mode to external services (syslog, email, SNMP trap, or HTTP servers). If you set this privilege to read-only, the administrator can see the forwarding settings of Traffic logs but can't manage them. If you disable this privilege, the administrator can't see or manage the settings.	Panorama: Yes Device Group/ Template: No	Yes	Yes	Yes

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
	*The Collector Groups privilege (**Panorama > Collector Groups**) controls forwarding for Traffic logs that Log Collectors receive from firewalls. The Log Forwarding privilege (**Objects > Log Forwarding**) controls forwarding from firewalls directly to external services (without aggregation on Log Collectors).*				
Threat	Specifies whether the administrator can see and configure the settings that control the forwarding of Threat logs from a Panorama virtual appliance in Legacy mode to external services (syslog, email, SNMP trap, or HTTP servers). If you set this privilege to read-only, the administrator can see the forwarding settings of Threat logs but can't manage them. If you disable this privilege, the administrator can't see or manage the settings.	Panorama: Yes Device Group/ Template: No	Yes	Yes	Yes

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
	*The Collector Groups privilege (**Panorama > Collector Groups**) controls forwarding for Threat logs that Log Collectors receive from firewalls. The Log Forwarding privilege (**Objects > Log Forwarding**) controls forwarding from firewalls directly to external services (without aggregation on Log Collectors).*				
WildFire	Specifies whether the administrator can see and configure the settings that control the forwarding of WildFire logs from a Panorama virtual appliance in Legacy mode to external services (syslog, email, SNMP trap, or HTTP servers). If you set this privilege to read-only, the administrator can see the forwarding settings of WildFire logs but can't manage them. If you disable this privilege, the administrator can't see or manage the settings.	Panorama: Yes Device Group/ Template: No	Yes	Yes	Yes

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
	*The Collector Groups privilege (**Panorama** > **Collector Groups**) controls the forwarding for WildFire logs that Log Collectors receive from firewalls. The Log Forwarding privilege (**Objects** > **Log Forwarding**) controls forwarding from firewalls directly to external services (without aggregation on Log Collectors).*				
Server Profiles	Sets the default state, enabled or disabled, for all the server profile privileges.	Panorama: Yes Device Group/ Template: No	Yes	No	Yes

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
	*These privileges pertain only to the server profiles that are used for forwarding logs from Panorama or Log Collectors and the server profiles that are used for authenticating Panorama administrators. The **Device** > Server Profiles privileges control access to the server profiles that are used for forwarding logs directly from firewalls to external services and for authenticating firewall administrators.*				
SNMP Trap	Specifies whether the administrator can see and configure SNMP trap server profiles. If you set this privilege to read-only, the administrator can see SNMP trap server profiles but can't manage them. If you disable this privilege, the administrator can't see or manage SNMP trap server profiles.	Panorama: Yes Device Group/ Template: No	Yes	Yes	Yes
Syslog	Specifies whether the administrator can see and configure Syslog server profiles. If you set this privilege to read-only, the administrator can see	Panorama: Yes Device Group/ Template: No	Yes	Yes	Yes

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
	Syslog server profiles but can't manage them. If you disable this privilege, the administrator can't see or manage Syslog server profiles.				
Email	Specifies whether the administrator can see and configure email server profiles. If you set this privilege to read-only, the administrator can see email server profiles but can't manage them. If you disable this privilege, the administrator can't see or manage email server profiles.	Panorama: Yes Device Group/ Template: No	Yes	Yes	Yes
RADIUS	Specifies whether the administrator can see and configure the RADIUS server profiles that are used to authenticate Panorama administrators. If you set this privilege to read-only, the administrator can see the RADIUS server profiles but can't manage them. If you disable this privilege, the administrator can't see or manage the RADIUS server profiles.	Panorama: Yes Device Group/ Template: No	Yes	Yes	Yes
TACACS+	Specifies whether the administrator can see and configure the TACACS+ server profiles that are used to authenticate Panorama administrators. If you disable this privilege, the administrator can't see the node or configure settings for the TACACS+ servers that authentication profiles reference.	Panorama: Yes Device Group/ Template: No	Yes	Yes	Yes

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
	If you set this privilege to read-only, the administrator can view existing TACACS+ server profiles but can't add or edit them.				
LDAP	Specifies whether the administrator can see and configure the LDAP server profiles that are used to authenticate Panorama administrators. If you set this privilege to read-only, the administrator can see the LDAP server profiles but can't manage them. If you disable this privilege, the administrator can't see or manage the LDAP server profiles.	Panorama: Yes Device Group/ Template: No	Yes	Yes	Yes
Kerberos	Specifies whether the administrator can see and configure the Kerberos server profiles that are used to authenticate Panorama administrators. If you set this privilege to read-only, the administrator can see the Kerberos server profiles but can't manage them. If you disable this privilege, the administrator can't see or manage the Kerberos server profiles.	Panorama: Yes Device Group/ Template: No	Yes	Yes	Yes
SAML Identity Provider	Specifies whether the administrator can see and configure the SAML Identity Provider (IdP) server profiles that are used to authenticate Panorama administrators. If you set this privilege to read-only, the administrator can see	Panorama: Yes Device Group/ Template: No	Yes	Yes	Yes

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
	the SAML IdP server profiles but can't manage them. If you disable this privilege, the administrator can't see or manage the SAML IdP server profiles.				
Scheduled Config Export	Specifies whether the administrator can view, add, edit, delete, or clone scheduled Panorama configuration exports. If you set this privilege to read-only, the administrator can view the scheduled exports but can't manage them. If you disable this privilege, the administrator can't see or manage the scheduled exports.	Panorama: Yes Device Group/ Template: No	Yes	No	Yes
Software	Specifies whether the administrator can: view information about software updates installed on the Panorama management server; download, upload, or install the updates; and view the associated release notes. If you set this privilege to read-only, the administrator can view information about Panorama software updates and view the associated release notes but can't perform any related operations. If you disable this privilege, the administrator can't see Panorama software updates, see the associated release notes, or perform any related operations.	Panorama: Yes Device Group/ Template: No	Yes	Yes	Yes

https://t.me/learningnets

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
	 *The **Panorama** > Device Deployment > Software privilege controls access to PAN-OS software deployed on firewalls and Panorama software deployed on Dedicated Log Collectors.*				
Dynamic Updates	Specifies whether the administrator can: view information about content updates installed on the Panorama management server (for example, WildFire updates); download, upload, install, or revert the updates; and view the associated release notes. If you set this privilege to read-only, the administrator can view information about Panorama content updates and view the associated release notes but can't perform any related operations. If you disable this privilege, the administrator can't see Panorama content updates, see the associated release notes, or perform any related operations.	Panorama: Yes Device Group/ Template: No	Yes	Yes	Yes

https://t.me/learningnets

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
	*The **Panorama** > Device Deployment > Dynamic Updates privilege controls access to content updates deployed on firewalls and Dedicated Log Collectors.*				
Support	Specifies whether the administrator can: view Panorama support license information, product alerts, and security alerts; activate a support license, and manage cases. Only a superuser admin can generate Tech Support files. If you set this privilege to read-only, the administrator can view Panorama support information, product alerts, and security alerts, but can't activate a support license, generate Tech Support files, or manage cases. If you disable this privilege, the administrator can't: see Panorama support information, product alerts, or security alerts; activate a support license, generate Tech Support files, or manage cases.	Panorama: Yes Device Group/ Template: No	Yes	Yes	Yes
Device Deployment	Sets the default state, enabled or disabled, for all the privileges associated with deploying licenses and software or content updates to firewalls and Log Collectors.	Panorama: Yes Device Group/ Template: Yes	Yes	No	Yes

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
	*The **Panorama** > Software and **Panorama** > Dynamic Updates privileges control the software and content updates installed on a Panorama management server.*				
Software	Specifies whether the administrator can: view information about the software updates installed on firewalls and Log Collectors; download, upload, or install the updates; and view the associated release notes. If you set this privilege to read-only, the administrator can see information about the software updates and view the associated release notes but can't deploy the updates to firewalls or dedicated Log Collectors. If you disable this privilege, the administrator can't see information about the software updates, see the associated release notes, or deploy the updates to firewalls or Dedicated Log Collectors.	Panorama: Yes Device Group/ Template: Yes	Yes	Yes	Yes
GlobalProtect Client	Specifies whether the administrator can: view information about GlobalProtect app software updates on firewalls; download, upload, or activate the updates; and view the associated release notes.	Panorama: Yes Device Group/ Template: Yes	Yes	Yes	Yes

https://t.me/learningnets

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
	If you set this privilege to read-only, the administrator can see information about GlobalProtect app software updates and view the associated release notes but can't activate the updates on firewalls. If you disable this privilege, the administrator can't see information about GlobalProtect app software updates, see the associated release notes, or activate the updates on firewalls.				
Dynamic Updates	Specifies whether the administrator can: view information about the content updates (for example, Applications updates) installed on firewalls and Dedicated Log Collectors; download, upload, or install the updates; and view the associated release notes. If you set this privilege to read-only, the administrator can see information about the content updates and view the associated release notes but can't deploy the updates to firewalls or Dedicated Log Collectors. If you disable this privilege, the administrator can't see information about the content updates, see the associated release notes, or deploy the updates to firewalls or Dedicated Log Collectors.	Panorama: Yes Device Group/ Template: Yes	Yes	Yes	Yes
Licenses	Specifies whether the administrator can view, refresh, and activate firewall licenses.	Panorama: Yes Device Group/ Template: Yes	Yes	Yes	Yes

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
	If you set this privilege to read-only, the administrator can view firewall licenses but can't refresh or activate those licenses. If you disable this privilege, the administrator can't view, refresh, or activate firewall licenses.				
Master Key and Diagnostics	Specifies whether the administrator can view and configure a master key by which to encrypt private keys on Panorama. If you set this privilege to read-only, the administrator can view the Panorama master key configuration but can't change it. If you disable this privilege, the administrator can't see or edit the Panorama master key configuration.	Panorama: Yes Device Group/ Template: No	Yes	Yes	Yes

Provide Granular Access to Operations Settings

To define which operations settings an administrator has access to, when creating or editing an admin role profile for a firewall (**Device** > **Admin Roles**), scroll down to the **Operations** option on the **Web UI** tab.

Access Level	Description	Enable	Read Only	Disable
Reboot	Restart the firewall. The firewall logs out all users, reloads the PAN-OS software and active configuration, closes and logs existing sessions, and creates a system log entry that shows the name of the administrator that initiated the reboot.	Yes	N/A	Yes
Generate Tech Support File	Generate a tech support system file that the Palo Alto Networks support team can	Yes	N/A	Yes

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Access Level	Description	Enable	Read Only	Disable
	use to troubleshoot issues that you may be experiencing with the firewall.			
Generate Stats Dump File	Generate and download a set of XML reports that summarizes network traffic over the last seven days for the firewall.	Yes	N/A	Yes
Download Core Files	If the firewall experiences a system process failure, a core file is automatically generated that contains details about the process and why it failed. You can download this core file to upload to your Palo Alto Networks support case to obtain further assistance in resolving the issue.	Yes	N/A	Yes

Panorama Web Interface Access Privileges

The custom Panorama administrator roles allow you to define access to the options on Panorama and the ability to only allow access to Device Groups and Templates (**Policies**, **Objects**, **Network**, **Device** tabs).

The administrator roles you can create are **Panorama** and **Device Group and Template**. You can't assign CLI access privileges to a **Device Group and Template** Admin Role profile. If you assign superuser privileges for the CLI to a **Panorama** Admin Role profile, administrators with that role can access all features regardless of the web interface privileges you assign.

Access Level	Description	Enable	Read Only	Disable
Dashboard	Controls access to the **Dashboard** tab. If you disable this privilege, the administrator will not see the tab and will not have access to any of the Dashboard widgets.	Yes	No	Yes
ACC	Controls access to the Application Command Center (ACC). If you disable this privilege, the **ACC** tab will not display in the web interface. Keep in mind that if you want to protect the privacy of your users while still providing access to the ACC, you can disable the **Privacy** > **Show Full IP Addresses** option and/or the **Show User Names In Logs And Reports** option.	Yes	No	Yes

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Access Level	Description	Enable	Read Only	Disable
Monitor	Controls access to the **Monitor** tab. If you disable this privilege, the administrator will not see the **Monitor** tab and will not have access to any of the logs, packet captures, session information, reports or to App Scope. For more granular control over what monitoring information the administrator can see, leave the Monitor option enabled and then enable or disable specific nodes on the tab as described in Provide Granular Access to the Monitor Tab.	Yes	No	Yes
Policies	Controls access to the **Policies** tab. If you disable this privilege, the administrator will not see the **Policies** tab and will not have access to any policy information. For more granular control over what policy information the administrator can see, for example to enable access to a specific type of policy or to enable read-only access to policy information, leave the **Policies** option enabled and then enable or disable specific nodes on the tab as described in Provide Granular Access to the Policy Tab.	Yes	No	Yes
Objects	Controls access to the **Objects** tab. If you disable this privilege, the administrator will not see the **Objects** tab and will not have access to any objects, security profiles, log forwarding profiles, decryption profiles, or schedules. For more granular control over what objects the administrator can see, leave the **Objects** option enabled and then enable or disable specific nodes on the tab as described in Provide Granular Access to the Objects Tab.	Yes	No	Yes
Network	Controls access to the **Network** tab. If you disable this privilege, the administrator will not see the **Network** tab and will not have access to any interface, zone, VLAN, virtual wire, virtual router, IPsec tunnel, DHCP, DNS Proxy, GlobalProtect, or QoS configuration information	Yes	No	Yes

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Access Level	Description	Enable	Read Only	Disable
	or to the network profiles. For more granular control over what objects the administrator can see, leave the **Network** option enabled and then enable or disable specific nodes on the tab as described in Provide Granular Access to the Network Tab.			
Device	Controls access to the **Device** tab. If you disable this privilege, the administrator will not see the **Device** tab and will not have access to any firewall-wide configuration information, such as User-ID, High Availability, server profile or certificate configuration information. For more granular control over what objects the administrator can see, leave the **Device** option enabled and then enable or disable specific nodes on the tab as described in Provide Granular Access to the Device Tab. *You can't enable access to the **Admin Roles** or **Administrators** nodes for a role-based administrator even if you enable full access to the **Device** tab.*	Yes	No	Yes
Panorama	Controls access to the **Panorama** tab. If you disable this privilege, the administrator will not see the **Panorama** tab and will not have access to any Panorama-wide configuration information, such as Managed Devices, Managed Collectors, or Collector Groups. For more granular control over what objects the administrator can see, leave the **Panorama** option enabled and then enable or disable specific nodes on the tab as described in Provide Granular Access to the Panorama Tab.	Yes	No	Yes

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Access Level	Description	Enable	Read Only	Disable
Privacy	Controls access to the privacy settings described in Define User Privacy Settings in the Admin Role Profile.	Yes	No	Yes
Validate	When disabled, an administrator cannot validate a configuration.	Yes	No	Yes
Save	Sets the default state (enabled or disabled) for all the save privileges described below (Partial Save and Save For Other Admins).	Yes	No	Yes
• Partial Save	When disabled, an administrator cannot save changes that any administrator made to the Panorama configuration.	Yes	No	Yes
• Save For Other Admins	When disabled, an administrator cannot save changes that other administrators made to the Panorama configuration.	Yes	No	Yes
Commit	Sets the default state (enabled or disabled) for all the commit, push, and revert privileges described below (Panorama, Device Groups, Templates, Force Template Values, Collector Groups, WildFire Appliance Clusters).	Yes	No	Yes
• Panorama	When disabled, an administrator cannot commit or revert configuration changes that any administrators made, including his or her own changes.	Yes	No	Yes
• Commit for Other Admins	When disabled, an administrator cannot commit or revert configuration changes that other administrators made.	Yes	No	Yes
Device Groups	When disabled, an administrator cannot push changes to device groups.	Yes	No	Yes
Templates	When disabled, an administrator cannot push changes to templates.	Yes	No	Yes
Force Template Values	This privilege controls access to the **Force Template Values** option in the Push Scope Selection dialog.	Yes	No	Yes

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Access Level	Description	Enable	Read Only	Disable
	When disabled, an administrator cannot replace overridden settings in local firewall configurations with settings that Panorama pushes from a template. *If you push a configuration with **Force Template Values** enabled, all overridden values on the firewall are replaced with values from the template. Before you use this option, check for overridden values on the firewalls to ensure your commit does not result in any unexpected network outages or issues caused by replacing those overridden values.*			
Collector Groups	When disabled, an administrator cannot push changes to Collector Groups.	Yes	No	Yes
WildFire Appliance Clusters	When disabled, an administrator cannot push changes to WildFire appliance clusters.	Yes	No	Yes
Tasks	When disabled, an administrator cannot access the Task Manager.	Yes	No	Yes
Global	Controls access to the global settings (system alarms) described in Provide Granular Access to Global Settings.	Yes	No	Yes

https://t.me/learningnets

Reference: Port Number Usage

The following tables list the ports that firewalls and Panorama use to communicate with each other, or with other services on the network.

- Ports Used for Management Functions
- Ports Used for HA
- Ports Used for Panorama
- Ports Used for GlobalProtect
- Ports Used for User-ID
- Ports Used for IPSec
- Ports Used for Routing
- Ports Used for DHCP
- Ports Used for Infrastructure

Ports Used for Management Functions

The firewall and Panorama use the following ports for management functions.

Destination Port	Protocol	Description
22	TCP	Used for communication from a client system to the firewall CLI interface.
80	TCP	The port the firewall listens on for Online Certificate Status Protocol (OCSP) updates when acting as an OCSP responder.
123	UDP	Port the firewall uses for NTP updates.
443	TCP	Used for communication from a client system to the firewall web interface. This is also the port the firewall and User-ID agent listens on for updates when you Enable VM Monitoring to Track Changes on the Virtual Network. For monitoring an AWS environment, this is the only port that is used. For monitoring a VMware vCenter/ESXi environment, the listening port defaults to 443, but it is configurable.
4443	TCP	Used as an alternative SSL port for HTTPS.
162	UDP	Port the firewall, Panorama, or a Log Collector uses to Forward Traps to an SNMP Manager.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Destination Port	Protocol	Description
		This port doesn't need to be open on the Palo Alto Networks firewall. You must configure the Simple Network Management Protocol (SNMP) manager to listen on this port. For details, refer to the documentation of your SNMP management software.
161	UDP	Port the firewall listens on for polling requests (GET messages) from the SNMP manager.
514 514 6514	TCP UDP SSL	Port that the firewall, Panorama, or a Log Collector uses to send logs to a syslog server if you Configure Syslog Monitoring, and the ports that the PAN-OS integrated User-ID agent or Windows-based User-ID agent listens on for authentication syslog messages.
2055	UDP	Default port the firewall uses to send NetFlow records to a NetFlow collector if you Configure NetFlow Exports, but this is configurable.
5008	TCP	Port the GlobalProtect Mobile Security Manager listens on for HIP requests from the GlobalProtect gateways. If you are using a third-party MDM system, you can configure the gateway to use a different port as required by the MDM vendor.
6080 6081 6082	TCP TLS 1.2 TCP	Ports used for User-ID™ Authentication Portal: • 6080 for NT LAN Manager (NTLM) authentication • 6081 for Authentication Portal without an SSL/TLS Server Profile • 6082 for Authentication Portal with an SSL/TLS Server Profile
10443	SSL	Port that the firewall and Panorama use to provide contextual information about a threat or to seamlessly shift your threat investigation to the Threat Vault and AutoFocus.

Ports Used for HA

Firewalls configured as High Availability (HA) peers must be able to communicate with each other to maintain state information (HA1 control link) and synchronize data (HA2 data link). In Active/Active HA deployments the peer firewalls must also forward packets to the HA peer that owns the session. The HA3 link is a Layer 2 (MAC-in-MAC) link and it does not support Layer 3 addressing or encryption.

https://t.me/learningnets

Destination Port	Protocol	Description
28769 28260	TCP TCP	Used for the HA1 control link for clear text communication between the HA peer firewalls. The HA1 link is a Layer 3 link and requires an IP address.
28	TCP	Used for the HA1 control link for encrypted communication (SSH over TCP) between the HA peer firewalls.
28770	TCP	Listening port for HA1 backup links.
28771	TCP	Used for heartbeat backups. Palo Alto Networks recommends enabling heartbeat backup on the MGT interface if you use an in-band port for the HA1 or the HA1 backup links.
99 29281	IP UDP	Used for the HA2 link to synchronize sessions, forwarding tables, IPSec security associations and ARP tables between firewalls in an HA pair. Data flow on the HA2 link is always unidirectional (except for the HA2 keep-alive); it flows from the active firewall (Active/Passive) or active-primary (Active/Active) to the passive firewall (Active/Passive) or active-secondary (Active/Active). The HA2 link is a Layer 2 link, and it uses ether type 0x7261 by default. The HA data link can also be configured to use either IP (protocol number 99) or UDP (port 29281) as the transport, and thereby allow the HA data link to span subnets.

Ports Used for Panorama

Panorama uses the following ports.

Destination Port	Protocol	Description
22	TCP	Used for communication from a client system to the Panorama CLI interface.
443	TCP	Used for communication from a client system to the Panorama web interface.
444	TCP	Used for communication between Panorama and Cortex Data Lake.
3978	TCP	Used for communication between Panorama and managed firewalls or managed collectors, as well as for

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Destination Port	Protocol	Description
		communication among managed collectors in a Collector Group: • For communication between Panorama and firewalls. This connection is initiated from the managed firewall to Panorama and facilitates a bi-directional data exchange on which the firewalls forward logs to Panorama and Panorama pushes configuration changes to the firewalls. Context switching commands are sent over the same connection. • Log Collectors use this destination port to forward logs to Panorama. • For communication with the default Log Collector on an M-Series appliance in Panorama mode and with Dedicated Log Collectors.
28443	TCP	Used for managed devices (firewalls and Log Collectors) to retrieve software and content updates from Panorama. *Only devices that run PAN-OS 8.x and later releases retrieve updates from Panorama over this port. For devices running earlier releases, Panorama pushes the update packages over port 3978.*
28769 (5.1 and later) 28260 (5.0 and later) 49160 (5.0 and earlier)	TCP TCP TCP	Used for the HA connectivity and synchronization between Panorama HA peers using clear text communication. Communication can be initiated by either peer.
28	TCP	Used for the HA connectivity and synchronization between Panorama HA peers using encrypted communication (SSH over TCP). Communication can be initiated by either peer. Used for communication between Log Collectors in a Collector Group for log distribution.
28270 (6.0 and later) 49190 (5.1 and earlier)	TCP	Used for communication among Log Collectors in a Collector Group for log distribution.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Destination Port	Protocol	Description
2049	TCP	Used by the Panorama virtual appliance to write logs to the NFS datastore.
10443	SSL	Port that Panorama uses to provide contextual information about a threat or to seamlessly shift your threat investigation to the Threat Vault and AutoFocus.
23000 to 23999	TCP, UDP, or SSL	Used for Syslog communication between Panorama and the Traps ESM components.

Ports Used for GlobalProtect

GlobalProtect uses the following ports.

Destination Port	Protocol	Description
443	TCP	Used for communication between GlobalProtect apps and portals, or GlobalProtect apps and gateways and for SSL tunnel connections. GlobalProtect gateways also use this port to collect host information from GlobalProtect apps and perform host information profile (HIP) checks.
4501	UDP	Used for IPSec tunnel connections between GlobalProtect apps and gateways.

For tips on how to use a loopback interface to provide access to GlobalProtect on different ports and addresses, refer to Can GlobalProtect Portal Page be Configured tobe Accessed on any Port?

Ports Used for User-ID

User-ID is a feature that enables mapping of user IP addresses to usernames and group memberships, enabling user- or group-based policy and visibility into user activity on your network (for example, to be able to quickly track down a user who may be the victim of a threat). To perform this mapping, the firewall, the User-ID agent (either installed on a Windows-based system or the PAN-OS integrated agent running on the firewall), and/or the Terminal Server agent must be able to connect to directory services on your network to perform Group Mapping and User Mapping. Additionally, if the agents are running on systems external to the firewall, they must be able to connect to the firewall to communicate the IP address to username mappings to the firewall. The following table lists the communication requirements for User-ID along with the port numbers required to establish connections.

https://t.me/learningnets

Destination Port	Protocol	Description
389	TCP	Port the firewall uses to connect to an LDAP server (plaintext or Start Transport Layer Security (Start TLS) to Map Users to Groups.
3268	TCP	Port the firewall uses to connect to an Active Directory global catalog server (plaintext or Start TLS) to Map Users to Groups.
636	TCP	Port the firewall uses for LDAP over SSL connections with an LDAP server to Map Users to Groups.
3269	TCP	Port the firewall uses for LDAP over SSL connections with an Active Directory global catalog server to Map Users to Groups.
514 6514	TCP UDP SSL	Port the User-ID agent listens on for authentication syslog messages if you Configure User-ID to Monitor Syslog Senders for User Mapping. The port depends on the type of agent and protocol: • PAN-OS integrated User-ID agent—Port 6514 for SSL and port 514 for UDP. • Windows-based User-ID agent—Port 514 for both TCP and UDP.
5007	TCP	Port the firewall listens on for user mapping information from the User-ID or Terminal Server agent. The agent sends the IP address and username mapping along with a timestamp whenever it learns of a new or updated mapping. In addition, it connects to the firewall at regular intervals to refresh known mappings.
5006	TCP	Port the User-ID agent listens on for XML API requests. The source for this communication is typically the system running a script that invokes the API.
88	UDP/TCP	Port the User-ID agent uses to authenticate to a Kerberos server. The firewall tries UDP first and falls back to TCP.
1812	UDP	Port the User-ID agent uses to authenticate to a RADIUS server.
49	TCP	Port the User-ID agent uses to authenticate to a TACACS+ server.
135	TCP	Port the User-ID agent uses to establish TCP-based WMI connections with the Microsoft Remote Procedure Call (RPC) Endpoint Mapper. The Endpoint Mapper then assigns the agent a randomly assigned port in the 49152-65535 port range. The agent uses this connection to make RPC queries for Exchange Server or AD server security logs, session tables. This is also the port used to access Terminal Servers.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Destination Port	Protocol	Description
		The User-ID agent also uses this port to connect to client systems to perform Windows Management Instrumentation (WMI) probing.
139	TCP	Port the User-ID agent uses to establish TCP-based NetBIOS connections to the AD server so that it can send RPC queries for security logs and session information. The User-ID agent also uses this port to connect to client systems for NetBIOS probing (supported on the Windows-based User-ID agent only).
445	TCP	Port the User-ID agent uses to connect to the Active Directory (AD) using TCP-based SMB connections to the AD server for access to user logon information (print spooler and Net Logon).
5985	HTTP	Port the User-ID agent uses to monitor security logs and session information with the WinRM protocol over HTTP.
5986	HTTPS	Port the User-ID agent uses to monitor security logs and session information with the WinRM protocol over HTTPS.
5009	TCP	Port the firewall uses to connect to the Terminal Server Agent.

Ports Used for IPSec

The firewall and Panorama use the following ports for IPSec functions.

Destination Port	Protocol	Description
500	UDP	Port used by IKE on the management plane to connect with remote IKE peers.
4500	UDP	Port used by IKE on the management plane to connect with remote IKE peers.
4510	UDP	Port used by the dataplane to send requests to IKE.
4511	UDP	Port used by the dataplane to send requests to `keymgr`.

Ports Used for Routing

The firewall and Panorama use the following ports for routing functions.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Destination Port	Protocol	Description
179	TCP	Port used by BGP to connect to peers.
3784 3785 4784	UDP	Ports used by BGP to connect to peers.
520	UDP	Port used for RIPv2.
89	IP	Port used for OSPF.
103	IP	Port used for Protocol Independent Multicast (PIM).

Ports Used for DHCP

The firewall and Panorama use the following ports for DHCP functions.

Destination Port	Protocol	Description
67 68 546 547	UDP	Ports used as DHCP server listening ports.

Ports Used for Infrastructure

The firewall and Panorama use the following ports for infrastructure functions.

Destination Port	Protocol	Description
111	TCP/UDP	Port used as a port mapper.
23	TCP/UDP	Port used for the Telnet application protocol.
69	TCP/UDP	Port used for TFTP.
2049	TCP/UDP	Port used for the Network File System (NFS).

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Destination Port	Protocol	Description
28260	TCP	Port used by internal sysd IPC communication for internal processes.
28261	TCP	Port used by internal masterd applications to manage internal processes.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Reset the Firewall to Factory Default Settings

Resetting the firewall to factory defaults will result in the loss of all configuration settings and logs.

STEP 1 | Set up a console connection to the firewall.

1. Connect a serial cable from your computer to the Console port and connect to the firewall using terminal emulation software (9600-8-N-1).

   If your computer does not have a 9-pin serial port, use a USB-to-serial port connector.

2. Enter your login credentials.
3. Enter the following CLI command:

   `debug system maintenance-mode`

   The firewall will reboot in the maintenance mode.

STEP 2 | Reset the system to factory default settings.

1. When the firewall reboots, press **`Enter`** to continue to the maintenance mode menu.
2. Select **`Factory Reset`** and press **`Enter`**.
3. Select **`Factory Reset`** and press **`Enter`** again.

   The firewall will reboot without any configuration settings. The default username and password to log in to the firewall is admin/admin.

   To perform initial configuration on the firewall and to set up network connectivity, see Integrate the Firewall into Your Management Network.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Bootstrap the Firewall

Bootstrapping speeds up the process of configuring and licensing the firewall to make it operational on the network with or without Internet access. Bootstrapping allows you to choose whether to configure the firewall with a basic configuration file (init-cfg.txt) so that it can connect to Panorama and obtain the complete configuration or to fully configure the firewall with the basic configuration and the optional bootstrap.xml file.

- USB Flash Drive Support
- Sample init-cfg.txt Files
- Prepare a USB Flash Drive for Bootstrapping a Firewall
- Bootstrap a Firewall Using a USB Flash Drive

USB Flash Drive Support

The USB flash drive that bootstraps a hardware-based Palo Alto Networks firewall must support one of the following:

- File Allocation Table 32 (FAT32)
- Third Extended File System (ext3)

The firewall can bootstrap from the following flash drives with USB2.0 or USB3.0 connectivity:

Supported USB Flash Drives
Kingston • Kingston SE9 8GB (2.0) • Kingston SE9 16GB (3.0) • Kingston SE9 32GB (3.0)
SanDisk • SanDisk Cruzer Fit CZ33 8GB (2.0) • SanDisk Cruzer Fit CZ33 16GB (2.0) • SanDisk Cruzer CZ36 16GB (2.0) • SanDisk Cruzer CZ36 32GB (2.0) • SanDisk Extreme CZ80 32GB (3.0)
Silicon Power • Silicon Power Jewel 32GB (3.0) • Silicon Power Blaze 16GB (3.0)
PNY • PNY Attache 16GB (2.0)

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Supported USB Flash Drives
• PNY Turbo 32GB (3.0)

Sample init-cfg.txt Files

An init-cfg.txt file is required for the bootstrap process; this file is a basic configuration file that you create using a text editor. To create this file, see 5 The following sample init-cfg.txt files show the parameters that are supported in the file; the parameters that you must provide are in bold.

Sample init-cfg.txt (Static IP Address)

```
type=static
ip-address=10.5.107.19
default-gateway=10.5.107.1
netmask=255.255.255.0
ipv6-address=2001:400:f00::1/64
ipv6-default-
gateway=2001:400:f00::2
hostname=Ca-FW-DC1
panorama-server=10.5.107.20
panorama-server-2=10.5.107.21
tplname=FINANCE_TG4
dgname=finance_dg
dns-primary=10.5.6.6
dns-secondary=10.5.6.7
op-command-modes=multi-
vsys,jumbo-frame
dhcp-send-hostname=no
dhcp-send-client-id=no
dhcp-accept-server-hostname=no
dhcp-accept-server-domain=no
```

Sample init-cfg.txt (DHCP Client)

```
type=dhcp-client
ip-address=
default-gateway=
netmask=
ipv6-address=
ipv6-default-gateway=
hostname=Ca-FW-DC1
panorama-server=10.5.107.20
panorama-server-2=10.5.107.21
tplname=FINANCE_TG4
dgname=finance_dg
dns-primary=10.5.6.6
dns-secondary=10.5.6.7
op-command-modes=multi-
vsys,jumbo-frame
dhcp-send-hostname=yes
dhcp-send-client-id=yes
dhcp-accept-server-
hostname=yes
dhcp-accept-server-domain=yes
```

The following table describes the fields in the init-cfg.txt file. The type is required; if the type is static, the IP address, default gateway and netmask are required, or the IPv6 address and IPv6 default gateway are required.

Field	Description
type	(Required) Type of management IP address: static or dhcp-client.
ip-address	(Required for IPv4 static management address) IPv4 address. The firewall ignores this field if the type is dhcp-client.
default-gateway	(Required for IPv4 static management address) IPv4 default gateway for the management interface. The firewall ignores this field if the type is dhcp-client.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Field	Description
netmask	(Required for IPv4 static management address) IPv4 netmask. The firewall ignores this field if the type is dhcp-client.
ipv6-address	(Required for IPv6 static management address) IPv6 address and / prefix length of the management interface. The firewall ignores this field if the type is dhcp-client.
ipv6-default-gateway	(Required for IPv6 static management address) IPv6 default gateway for the management interface. The firewall ignores this field if the type is dhcp-client.
hostname	(Optional) Host name for the firewall.
panorama-server	(Recommended) IPv4 or IPv6 address of the primary Panorama server.
panorama-server-2	(Optional) IPv4 or IPv6 address of the secondary Panorama server.
tplname	(Recommended) Panorama template name.
dgname	(Recommended) Panorama device group name.
dns-primary	(Optional) IPv4 or IPv6 address of the primary DNS server.
dns-secondary	(Optional) IPv4 or IPv6 address of the secondary DNS server.
vm-auth-key	(VM-Series firewalls only) Virtual machine authentication key.
op-command-modes	(Optional) Enter multi-vsys, jumbo-frame, or both separated by a comma only. Enables multiple virtual systems and jumbo frames while bootstrapping.
dhcp-send-hostname	(DHCP client type only) The DHCP server determines a value of yes or no. If yes, the firewall sends its hostname to the DHCP server.
dhcp-send-client-id	(DHCP client type only) The DHCP server determines a value of yes or no. If yes, the firewall sends its client ID to the DHCP server.
dhcp-accept-server-hostname	(DHCP client type only) The DHCP server determines a value of yes or no. If yes, the firewall accepts its hostname from the DHCP server.
dhcp-accept-server-domain	(DHCP client type only) The DHCP server determines a value of yes or no. If yes, the firewall accepts its DNS server from the DHCP server.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Prepare a USB Flash Drive for Bootstrapping a Firewall

You can use a USB flash drive to bootstrap a physical firewall. However, to do so you must be running a PAN-OS 7.1.0 or later image and Reset the Firewall to Factory Default Settings. For security reasons, you can bootstrap a firewall only when it is in factory default state or has all private data deleted.

STEP 1 | Obtain serial numbers (S/Ns) and auth codes for support subscriptions from your order fulfillment email.

STEP 2 | Register S/Ns of new firewalls on the Customer Support portal.

1. Go to support.paloaltonetworks.com, log in, and select **Assets** > **Devices** > **Register New Device** > **Register device using Serial Number or Authorization Code**.
2. Follow the steps to Register the Firewall.
3. Click **Submit**.

STEP 3 | Activate authorization codes on the Customer Support portal, which creates license keys.

1. Go to support.paloaltonetworks.com, log in, and select the **Assets** > **Devices** on the left-hand navigation pane.
2. For each device S/N you just registered, click the **Action** link (the pencil icon).
3. Under Activate Licenses, select **Activate Auth-Code**.
4. Enter the **Authorization code** and click **Agree** and **Submit**.

STEP 4 | Add the S/Ns in Panorama.

Complete Step 1 in Add a Firewall as a Managed Device in the Panorama Administrator's Guide.

STEP 5 | Create the init-cfg.txt file.

Create the init-cfg.txt file, a mandatory file that provides bootstrap parameters. The fields are described in Sample init-cfg.txt Files.

If the init-cfg.txt file is missing, the bootstrap process will fail and the firewall will boot up with the default configuration in the normal boot-up sequence.

There are no spaces between the key and value in each field; do not add spaces because they cause failures during parsing on the management server side.

You can have multiple init-cfg.txt files—one each for different remote sites—by prepending the S/N to the file name. For example:

0008C200105-init-cfg.txt

0008C200107-init-cfg.txt

If no prepended filename is present, the firewall uses the init-cfg.txt file and proceeds with bootstrapping.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 6 | (Optional) Create the bootstrap.xml file.

The optional bootstrap.xml file is a complete firewall configuration that you can export from an existing production firewall.

1. Select **Device** > **Setup** > **Operations** > **Export named configuration snapshot**.
2. Select the **Name** of the saved or the running configuration.
3. Click **OK**.
4. Rename the file as **`bootstrap.xml`**.

STEP 7 | Create and download the bootstrap bundle from the Customer Support portal.

For a physical firewall, the bootstrap bundle requires only the /license and /config directories.

Use one of the following methods to create and download the bootstrap bundle:

- Use **Method 1** to create a bootstrap bundle specific to a remote site (you have only one init-cfg.txt file).
- Use **Method 2** to create one bootstrap bundle for multiple sites.

Method 1

1. On your local system, go to support.paloaltonetworks.com and log in.
2. Select **Assets**.
3. Select the S/N of the firewall you want to bootstrap.
4. Select **Bootstrap Container**.
5. Click **Select**.
6. Upload and **Open** the init-cfg.txt file you created.
7. (Optional) Select the bootstrap.xml file you created and **Upload Files**.

   You must use a bootstrap.xml file from a firewall of the same model and PAN-OS version.
8. Select **Bootstrap Container Download** to download a tar.gz file named **`bootstrap_<S/N>_<date>.tar.gz`** to your local system. This bootstrap container includes the license keys associated with the S/N of the firewall.

Method 2

Create a tar.gz file on your local system with two top-level directories: /license and /config. Include all licenses and all init-cfg.txt files with S/Ns prepended to the filenames.

The license key files you download from the Customer Support portal have the S/N in the license file name. PAN-OS checks the S/N in the file name against the firewall S/N while executing the bootstrap process.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 8 | Import the tar.gz file you created (to a firewall running a PAN-OS 7.1.0 or later image) using Secure Copy (SCP) or TFTP.

Access the CLI and enter one of the following commands:

- **`tftp import bootstrap-bundle file <path and filename> from <host IP address>`**

  For example:

  `tftp import bootstrap-bundle file /home/userx/bootstrap/devices/pa5000.tar.gz from 10.1.2.3`

- **`scp import bootstrap-bundle from <<user>@<host>:<path to file>>`**

  For example:

  `scp import bootstrap-bundle from userx@10.1.2.3:/home/userx/bootstrap/devices/pa200_bootstrap_bundle.tar.gz`

STEP 9 | Prepare the USB flash drive.

1. Insert the USB flash drive into the firewall that you used in the prior step.
2. Enter the following CLI operational command, using your tar.gz filename in place of "**`pa5000.tar.gz`**". This command formats the USB flash drive, unzips the file, and validates the USB flash drive:

   `request system bootstrap-usb prepare from pa5000.tar.gz`
3. Press **`y`** to continue. The following message displays when the USB drive is ready:

   `USB prepare completed successfully.`
4. Remove the USB flash drive from the firewall.
5. You can prepare as many USB flash drives as needed.

STEP 10 | Deliver the USB flash drive to your remote site.

If you used Method 2 to create the bootstrap bundle, you can use the same USB flash drive content for bootstrapping firewalls at multiple remote sites. You can translate the content into multiple USB flash drives or a single USB flash drive used multiple times.

Bootstrap a Firewall Using a USB Flash Drive

After you receive a new Palo Alto Networks firewall and a USB flash drive loaded with bootstrap files, you can bootstrap the firewall.



Microsoft Windows and Apple Mac operating systems are unable to read the bootstrap USB flash drive because the drive is formatted using an ext4 file system. You must install third-party software or use a Linux system to read the USB drive.

STEP 1 | The firewall must be in a factory default state or must have all private data deleted.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 2 | To ensure connectivity with your corporate headquarters, cable the firewall by connecting the management interface (MGT) using an Ethernet cable to one of the following:

- An upstream modem
- A port on the switch or router
- An Ethernet jack in the wall

STEP 3 | Insert the USB flash drive into the USB port on the firewall and power on the firewall. The factory default firewall bootstraps itself from the USB flash drive.

The firewall Status light turns from yellow to green when the firewall is configured; autocommit is successful.

STEP 4 | Verify bootstrap completion. You can see basic status logs on the console during the bootstrap and you can verify that the process is complete.

1. If you included Panorama values (panorama-server, tplname, and dgname) in your init-cfg.txt file, check Panorama managed devices, device group, and template name.
2. Verify the general system settings and configuration by accessing the web interface and selecting **Dashboard** > **Widgets** > **System** or by using the CLI operational commands **`show system info`** and **`show config running`**.
3. Verify the license installation by selecting **Device** > **Licenses** or by using the CLI operational command **`request license info`**.
4. If you have Panorama configured, manage the content versions and software versions from Panorama. If you do not have Panorama configured, use the web interface to manage content versions and software versions.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 5 | (Panorama managed firewalls only) Create a device registration authentication key and add it to the firewall.

This is required to successfully add a bootstrapped firewall to Panorama management. The device registration authentication key has a finite lifetime and including the device registration authentication key in the init-cfg.txt file is not supported.

1. Log in to the Panorama web interface.
2. Select **Panorama** > **Device Registration Auth Key** and **Add** a new authentication key.
3. Configure the authentication key.
   - **Name**—Add a descriptive name for the authentication key.
   - **Lifetime**—Specify the key lifetime to limit how long you can use the authentication key to onboard new firewalls.
   - **Count**—Specify how many times you can use the authentication key to onboard new firewalls.
   - **Device Type**—Specify that this authentication key is used to authenticate only a **Firewall**.

     *You can select **Any** to use the device registration authentication key to onboard firewalls, Log Collectors, and WildFire appliances.*
   - (Optional) **Devices**—Enter one or more device serial numbers to specify for which firewalls the authentication key is valid.
4. Click **OK**.

   When prompted, **Copy Auth Key** and **Close**.
5. Log in to the firewall web interface.
6. Select **Device** > **Setup** > **Management** and edit the Panorama Settings.
7. Paste the device registration authentication key you copied in the previous step and click **OK**.
8. **Commit**.
9. Log in to the Panorama web interface and select **Panorama** > **Managed Devices** > **Summary** to verify the firewall is `Connected` to Panorama

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Device Telemetry

Device telemetry collects data about your next-generation firewall or Panorama, and shares it with Palo Alto Networks by uploading the data to Cortex Data Lake. This data is used to power telemetry apps, and for sharing threat intelligence.

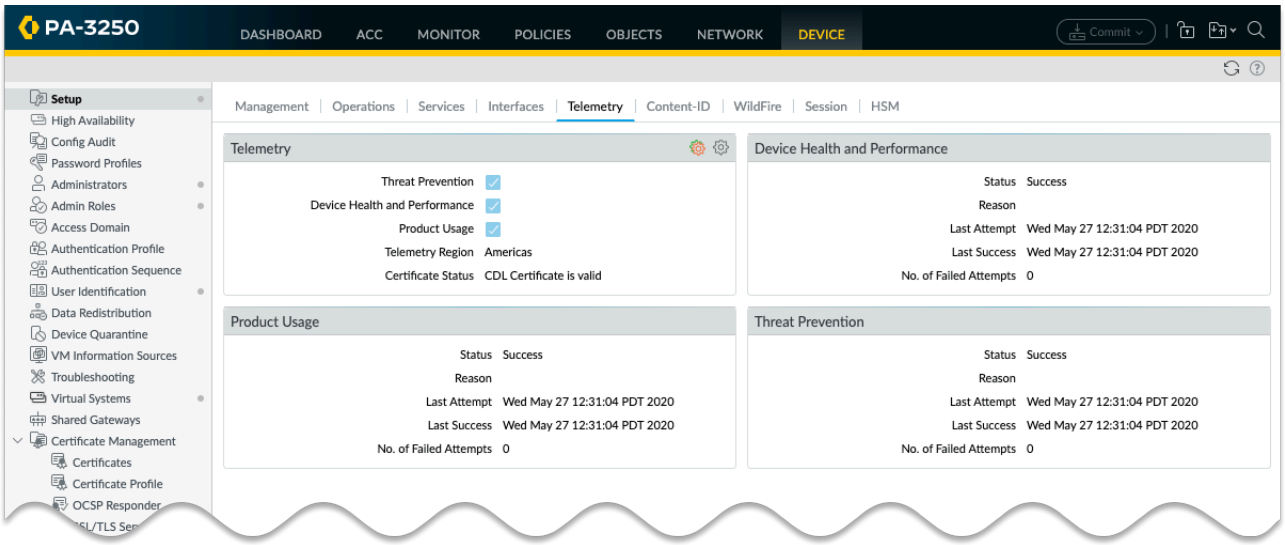
> Device Telemetry Overview
> Device Telemetry Collection and Transmission Intervals
> Manage Device Telemetry
> Monitor Device Telemetry
> Sample the Data that Device Telemetry Collects

https://t.me/learningnets

Device Telemetry Overview

Device telemetry collects data about your next-generation firewall or Panorama and shares it with Palo Alto Networks by uploading the data to Cortex Data Lake. This data is used to power telemetry apps, which are cloud-based applications that make it easy to monitor and manage your next-generation firewalls and Panoramas. These apps improve your visibility into device health, performance, capacity planning, and configuration. Through these apps, you can maximize the benefits you enjoy from the products and services that Palo Alto Networks delivers.

Telemetry data is also used for sharing threat intelligence, providing enhanced intrusion prevention, evaluation of threat signatures, as well as improved malware detection within PAN-DB URL filtering, DNS-based command-and-control (C2) signatures, and WildFire.



Telemetry data is collected and stored locally on your device for a limited period of time. This data is shared with Palo Alto Networks only if you configure a destination region for the data. If your organization has a Cortex Data Lake license, then you can only send the data to the same region as where your Cortex Data Lake instance resides. If your organization does not have a Cortex Data Lake license, then you must install a device certificate in order to share this data. In this case, you can choose any available region, although you must conform to all applicable local laws regarding privacy and data storage.

Telemetry data is collected and shared with Palo Alto Networks on predefined collection intervals. You can control whether data is collected and shared by enabling/disabling categories of data. You can also monitor the current status of data collection and transmission.

Finally, you can obtain a live sample of the data that your firewall is collecting for telemetry purposes. For a complete description of all the telemetry metrics that can be shared with Palo Alto Networks, including the privacy implication for each metric, see the PAN-OS Device Telemetry Metrics Reference Guide.

*The automatically created user **_cliadmin** may appear under **Logged in Admins** on the dashboard while telemetry is enabled. This user is created only for telemetry collection.*

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Device Telemetry Collection and Transmission Intervals

PAN-OS collects and sends telemetry data on fixed intervals. Collection is defined on a metric by metric basis, and can be one of:

- Every 20 minutes.
- Every hour.
- Daily.

Telemetry is collected into data bundles. Each bundle is an aggregation of all the data collected up to the point of data transmission. These bundles are stored on the device until a transmission event, which occur once every 1 hour. When a bundle has been successfully sent to Palo Alto Networks, it is deleted from the device.

If an error occurs while sending a bundle to Palo Alto Networks, the firewall waits 10 minutes and then tries again. The firewall will continue to try to send the bundle until it is either successful, or it needs the storage space to collect new telemetry data.

At every regular transmission interval, the firewall begins by sending the bundles scheduled for that event. After a successful transfer of those bundles, the firewall sends any failed bundles that it might have stored from previous transmission events.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Manage Device Telemetry

To manage device telemetry you can:

- Enable Device Telemetry
- Disable Device Telemetry
- Manage the Data that Device Telemetry Collects
- Manage Historical Device Telemetry

Enable Device Telemetry

By default, your device does not share data with Palo Alto Networks. If sharing is enabled, you can stop sharing all device telemetry by: **Device** > **Setup** > **Telemetry**, uncheck the **Enable Telemetry** box, and then commit your change.

To enable Device Telemetry so that data is shared with Palo Alto Networks:

STEP 1 | Enable Cortex Data Lake.

1. If your organization does not have a Cortex Data Lake license, install a device certificate if one is not already installed on your device.

   If your organization does have a Cortex Data Lake license, make sure it is activated.
2. Make sure that your network is properly configured so that the firewall can send data to Cortex Data Lake.

STEP 2 | Navigate to **Device** > **Setup** > **Telemetry**

STEP 3 | Edit the **Telemetry** widget.

STEP 4 | In **Telemetry Destination**, select your region. If your organization is using Cortex Data Lake, you must use the region that your Cortex Data Lake is configured to use.

STEP 5 | Click **OK**, and then commit your changes.

Disable Device Telemetry

If your next-generation firewall is configured to share data with Palo Alto Networks, you can disable this sharing by:

STEP 1 | Navigate to **Device** > **Setup** > **Telemetry**

STEP 2 | Edit the **Telemetry** widget.

STEP 3 | Uncheck the **Enable Telemetry** box.

STEP 4 | Click **OK**, and then commit your changes.

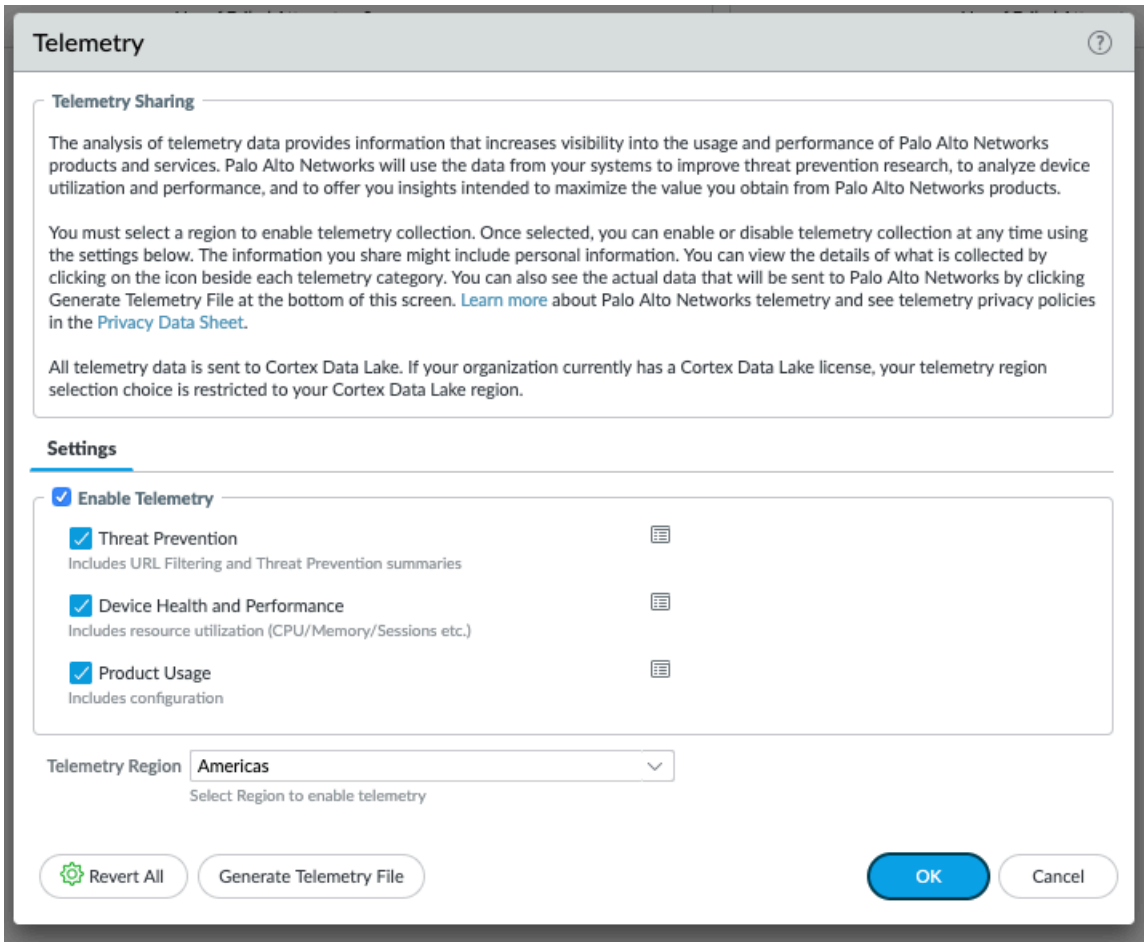
STEP 5 | Any telemetry data currently stored in Cortex Data Lake is automatically purged one year after your firewall uploaded it. Optionally, if you do not want the data to reside in Cortex Data Lake for this amount of time after you disable telemetry, open a support ticket and ask Palo Alto Networks to purge your telemetry data.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Manage the Data the Device Telemetry Collects

Select **Device** > **Setup** > **Telemetry** to see the currently collected telemetry categories. To change these categories, edit the Telemetry widget. Deselect any categories that you don't want the firewall to collect, **OK**, and then commit the change.



*To stop sharing all device telemetry, uncheck the **Enable Telemetry** box, and then commit your change.*

Manage Historical Device Telemetry

Device Telemetry changed significantly for the PAN-OS 10.1 release. Prior to 10.0, telemetry data was mostly of interest for threat intelligence purposes. As of 10.0, threat intelligence metrics are still a large portion the data collected by the device, but a great deal more data involving the health, performance, and configuration of the device is collected as well.

In other words, PAN-OS 10.1 device telemetry extends the data that was collected for previous releases. PAN-OS 10.1 also sends telemetry data to a different cloud location than did prior releases. But the historical telemetry support still exists for next-generation firewalls running PAN-OS 10.0. The only difference is that the 10.1 device telemetry user interface is not capable of managing this historical data collection.

https://t.me/learningnets

If you have an existing next-generation firewall, and you have any of the historical telemetry data categories enabled, then when you upgrade to PAN-OS 10.1 your firewall will continue to collect and share this information. If you want to turn this telemetry data sharing off, use the following CLI commands:

```
set deviceconfig system update-schedule statistics-service
 application-reports no
set deviceconfig system update-schedule statistics-service threat-
prevention-reports no
set deviceconfig system update-schedule statistics-service threat-
prevention-information no
set deviceconfig system update-schedule statistics-service threat-
prevention-pcap no
set deviceconfig system update-schedule statistics-service passive-
dns-monitoring no
set deviceconfig system update-schedule statistics-service url-
reports no
set deviceconfig system update-schedule statistics-service health-
performance-reports no
set deviceconfig system update-schedule statistics-service file-
identification-reports no
```

If you have a 10.1 firewall and this telemetry sharing is turned off, but you want to share this data with Palo Alto Networks, then you can turn it on using:

```
set deviceconfig system update-schedule statistics-service
 application-reports yes
set deviceconfig system update-schedule statistics-service threat-
prevention-reports yes
set deviceconfig system update-schedule statistics-service threat-
prevention-information yes
set deviceconfig system update-schedule statistics-service threat-
prevention-pcap yes
set deviceconfig system update-schedule statistics-service passive-
dns-monitoring yes
set deviceconfig system update-schedule statistics-service url-
reports yes
set deviceconfig system update-schedule statistics-service health-
performance-reports yes
set deviceconfig system update-schedule statistics-service file-
identification-reports yes
```

You can see whether your device is collecting and sharing this historical telemetry data using the following CLI command:

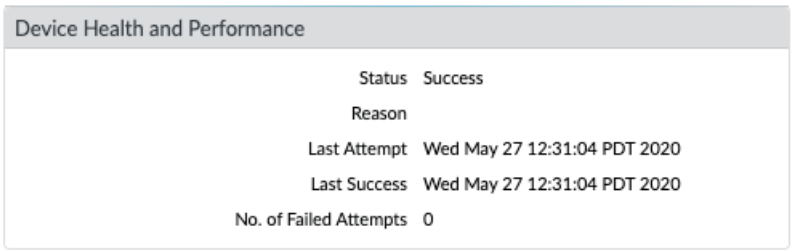
```
show deviceconfig system update-schedule statistics-service
```

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Monitor Device Telemetry

PAN-OS shows you the sharing status for each telemetry category. Widgets for each metrics category are available at **Device** > **Setup** > **Telemetry**.



In the event of a failure, your device will retry the send attempt at the next transmission time. If the problem persists, check to make sure that your devices are properly configured to send data to Cortex Data Lake:

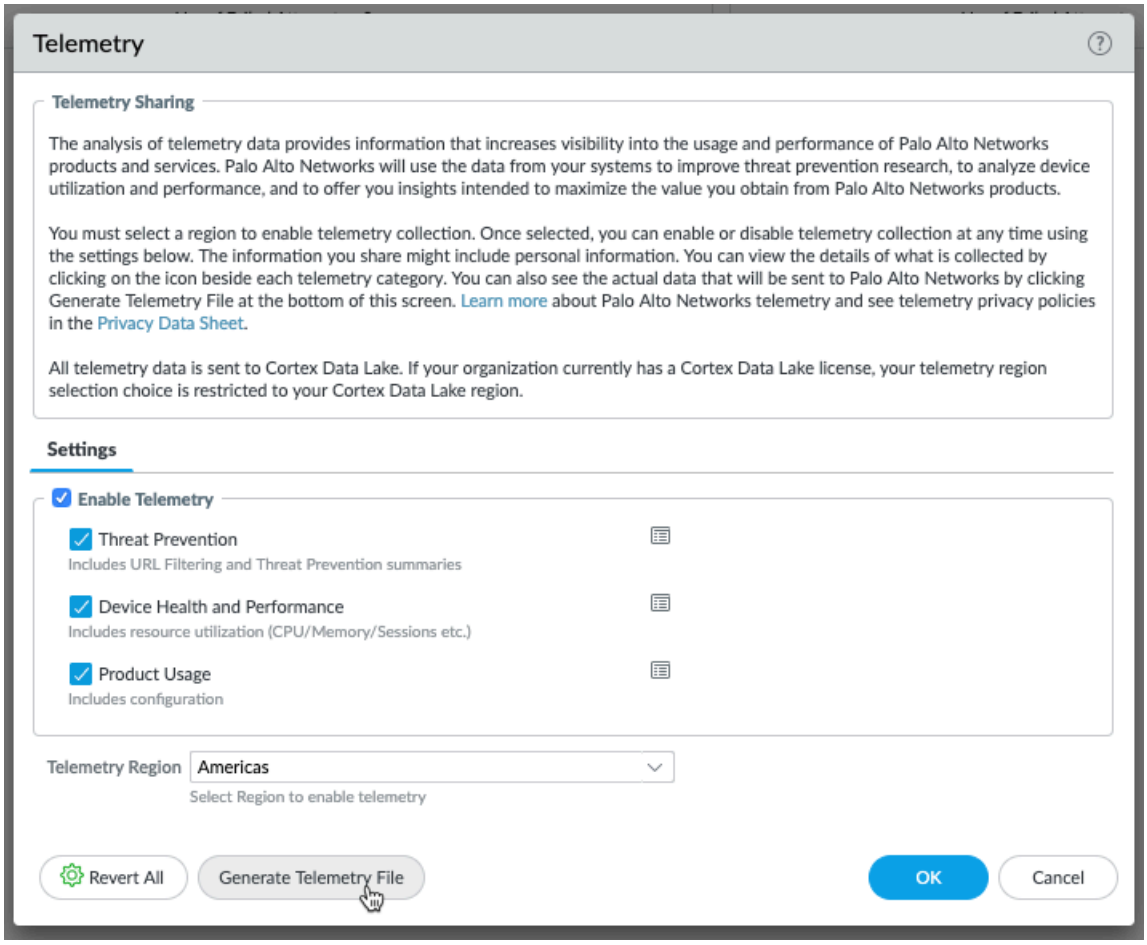
- If your organization has a Cortex Data Lake license, then make sure your Cortex Data Lake license has been activated, and that your firewall is configured to use Cortex Data Lake.
- If your organization does not have a Cortex Data Lake license, then make sure you have installed a device certificate, and that your network is configured to allow traffic to Cortex Data Lake.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Sample the Data that Device Telemetry Collects

You can download a live example of the data that device telemetry collects and shares with Palo Alto Networks. To do this, go to **Device** > **Setup** > **Telemetry**, and edit the **Telemetry** widget. Then click **Generate Telemetry File**.



The data collection will take a few minutes, depending on the speed of your firewall. When the process completes, click **Download Device Telemetry Data**. The telemetry bundle is a compressed tar ball, and it is placed in your default browser download directory.

For a description of every metric that device telemetry collects and shares with Palo Alto Networks, see the PAN-OS Device Telemetry Metrics Reference Guide.

https://t.me/learningnets

Authentication

Authentication is a method for protecting services and applications by verifying the identities of users so that only legitimate users have access. Several firewall and Panorama features require authentication. Administrators authenticate to access the web interface, CLI, or XML API of the firewall and Panorama. End users authenticate through Authentication Portal or GlobalProtect to access various services and applications. You can choose from several authentication services to protect your network and to accommodate your existing security infrastructure while ensuring a smooth user experience.

If you have a public key infrastructure, you can deploy certificates to enable authentication without users having to manually respond to login challenges (see Certificate Management). Alternatively, or in addition to certificates, you can implement interactive authentication, which requires users to authenticate using one or more methods. The following topics describe how to implement, test, and troubleshoot the different types of interactive authentication:

- Authentication Types
- Plan Your Authentication Deployment
- Configure Multi-Factor Authentication
- Configure SAML Authentication
- Configure Kerberos Single Sign-On
- Configure Kerberos Server Authentication
- Configure TACACS+ Authentication
- Configure RADIUS Authentication
- Configure LDAP Authentication
- Connection Timeouts for Authentication Servers
- Configure Local Database Authentication
- Configure an Authentication Profile and Sequence
- Test Authentication Server Connectivity
- Authentication Policy
- Troubleshoot Authentication Issues

https://t.me/learningnets

Authentication Types

- External Authentication Services
- Multi-Factor Authentication
- SAML
- Kerberos
- TACACS+
- RADIUS
- LDAP
- Local Authentication

External Authentication Services

The firewall and Panorama can use external servers to control administrative access to the web interface and end user access to services or applications through Authentication Portal and GlobalProtect. In this context, any authentication service that is not local to the firewall or Panorama is considered external, regardless of whether the service is internal (such as Kerberos) or external (such as a SAML identity provider) relative to your network. The server types that the firewall and Panorama can integrate with include Multi-Factor Authentication (MFA), SAML, Kerberos, TACACS+, RADIUS, and LDAP. Although you can also use the Local Authentication services that the firewall and Panorama support, usually external services are preferable because they provide:

- Central management of all user accounts in an external identity store. All the supported external services provide this option for end users and administrators.
- Central management of account authorization (role and access domain assignments). SAML, TACACS+, and RADIUS support this option for administrators.
- Single sign-on (SSO), which enables users to authenticate only once for access to multiple services and applications. SAML and Kerberos support SSO.
- Multiple authentication challenges of different types (factors) to protect your most sensitive services and applications. MFA services support this option.

Authentication through an external service requires a server profile that defines how the firewall connects to the service. You assign the server profile to authentication profiles, which define settings that you customize for each application and set of users. For example, you can configure one authentication profile for administrators who access the web interface and another profile for end users who access a GlobalProtect portal. For details, see Configure an Authentication Profile and Sequence.

Multi-Factor Authentication

You can Configure Multi-Factor Authentication (MFA) to ensure that each user authenticates using multiple methods (factors) when accessing highly sensitive services and applications. For example, you can force users to enter a login password and then enter a verification code that they receive by phone before allowing access to important financial documents. This approach helps to prevent attackers from accessing every service and application in your network just by stealing passwords.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Of course, not every service and application requires the same degree of protection, and MFA might not be necessary for less sensitive services and applications that users access frequently. To accommodate a variety of security needs, you can Configure Authentication Policy rules that trigger MFA or a single authentication factor (such as login credentials or certificates) based on specific services, applications, and end users.

When choosing how many and which types of authentication factors to enforce, it's important to understand how policy evaluation affects the user experience. When a user requests a service or application, the firewall first evaluates Authentication policy. If the request matches an Authentication policy rule with MFA enabled, the firewall displays a Authentication Portal web form so that users can authenticate for the first factor. If authentication succeeds, the firewall displays an MFA login page for each additional factor. Some MFA services prompt the user to choose one factor out of two to four, which is useful when some factors are unavailable. If authentication succeeds for all factors, the firewall evaluates Security policy for the requested service or application.



To reduce the frequency of authentication challenges that interrupt the user workflow, configure the first factor to use Kerberos or SAML single sign-on (SSO) authentication.

To implement MFA for GlobalProtect, refer to Configure GlobalProtect to facilitate multi-factor authentication notifications.

You cannot use MFA authentication profiles in authentication sequences.

For end-user authentication via Authentication Policy, the firewall directly integrates with several MFA platforms (Duo v2, Okta Adaptive, PingID, and RSA SecurID), as well as integrating through RADIUS or SAML for all other MFA platforms. For remote user authentication to GlobalProtect portals and gateways and for administrator authentication to the Panorama and PAN-OS web interface, the firewall integrates with MFA vendors using RADIUS and SAML only.

The firewall supports the following MFA factors:

Factor	Description
Push	An endpoint device (such as a phone or tablet) prompts the user to allow or deny authentication.
Short message service (SMS)	An SMS message on the endpoint device prompts the user to allow or deny authentication. In some cases, the endpoint device provides a code that the user must enter in the MFA login page.
Voice	An automated phone call prompts the user to authenticate by pressing a key on the phone or entering a code in the MFA login page.
One-time password (OTP)	An endpoint device provides an automatically generated alphanumeric string, which the user enters in the MFA login page to enable authentication for a single transaction or session.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

SAML

You can use Security Assertion Markup Language (SAML) 2.0 to authenticate administrators who access the firewall or Panorama web interface and end users who access web applications that are internal or external to your organization. In environments where each user accesses many applications and authenticating for each one would impede user productivity, you can configure SAML single sign-on (SSO) to enable one login to access multiple applications. Likewise, SAML single logout (SLO) enables a user to end sessions for multiple applications by logging out of just one session. SSO is available to administrators who access the web interface and to end users who access applications through GlobalProtect or Authentication Portal. SLO is available to administrators and GlobalProtect end users, but not to Authentication Portal end users. When you configure SAML authentication on the firewall or on Panorama, you can specify SAML attributes for administrator authorization. SAML attributes enable you to quickly change the roles, access domains, and user groups of administrators through your directory service, which is often easier than reconfiguring settings on the firewall or Panorama.

Administrators cannot use SAML to authenticate to the CLI on the firewall or Panorama.

You cannot use SAML authentication profiles in authentication sequences.

SAML authentication requires a *service provider* (the firewall or Panorama), which controls access to applications, and an *identity provider* (IdP) such as PingFederate, which authenticates users. When a user requests a service or application, the firewall or Panorama intercepts the request and redirects the user to the IdP for authentication. The IdP then authenticates the user and returns a *SAML assertion*, which indicates authentication succeeded or failed. SAML Authentication for Authentication Portal End Users illustrates SAML authentication for an end user who accesses applications through Authentication Portal.

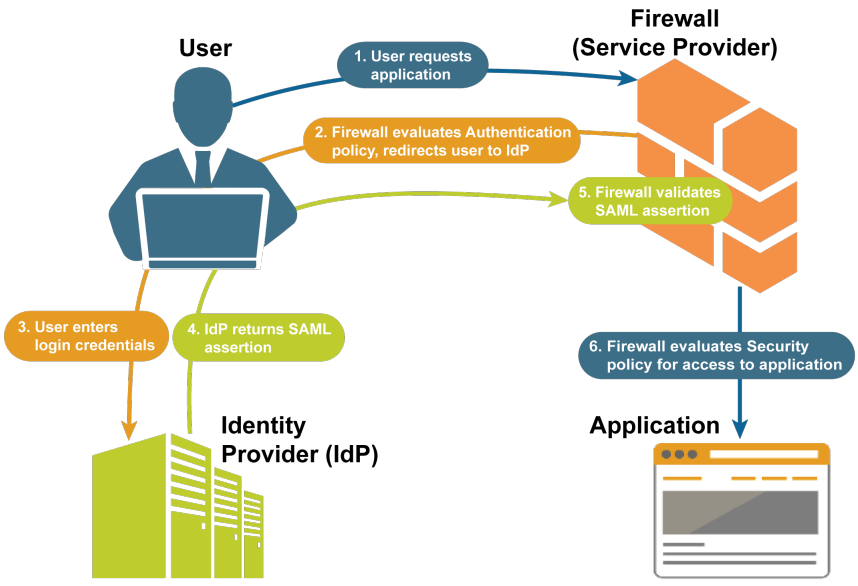


Figure 1: SAML Authentication for Authentication Portal End Users

Kerberos

Kerberos is an authentication protocol that enables a secure exchange of information between parties over an insecure network using unique keys (called tickets) to identify the parties. The

https://t.me/learningnets

firewall and Panorama support two types of Kerberos authentication for administrators and end users:

- **Kerberos server authentication**—A Kerberos server profile enables users to natively authenticate to an Active Directory domain controller or a Kerberos V5-compliant authentication server. This authentication method is interactive, requiring users to enter usernames and passwords. For the configuration steps, see Configure Kerberos Server Authentication.
- **Kerberos single sign-on (SSO)**—A network that supports Kerberos V5 SSO prompts a user to log in only for initial access to the network (such as logging in to Microsoft Windows). After this initial login, the user can access any browser-based service in the network (such as the firewall web interface) without having to log in again until the SSO session expires. (Your Kerberos administrator sets the duration of SSO sessions.) If you enable both Kerberos SSO and another external authentication service (such as a TACACS+ server), the firewall first tries SSO and, only if that fails, falls back to the external service for authentication. To support Kerberos SSO, your network requires:
  - A Kerberos infrastructure, including a key distribution center (KDC) with an authentication server (AS) and ticket-granting service (TGS).
  - A Kerberos account for the firewall or Panorama that will authenticate users. An account is required to create a Kerberos keytab, which is a file that contains the principal name and hashed password of the firewall or Panorama. The SSO process requires the keytab.

  For the configuration steps, see Configure Kerberos Single Sign-On.

  Kerberos SSO is available only for services and applications that are internal to your Kerberos environment. To enable SSO for external services and applications, use SAML.

TACACS+

Terminal Access Controller Access-Control System Plus (TACACS+) is a family of protocols that enable authentication and authorization through a centralized server. TACACS+ encrypts usernames and passwords, making it more secure than RADIUS, which encrypts only passwords. TACACS+ is also more reliable because it uses TCP, whereas RADIUS uses UDP. You can configure TACACS+ authentication for end users or administrators on the firewall and for administrators on Panorama. Optionally, you can use TACACS+ Vendor-Specific Attributes (VSAs) to manage administrator authorization. TACACS+ VSAs enable you to quickly change the roles, access domains, and user groups of administrators through your directory service instead of reconfiguring settings on the firewall and Panorama.

The firewall and Panorama support the following TACACS+ attributes and VSAs. Refer to your TACACS+ server documentation for the steps to define these VSAs on the TACACS+ server.

Name	Value
service	This attribute is required to identify the VSAs as specific to Palo Alto Networks. You must set the value to **`PaloAlto`**.

https://t.me/learningnets

Name	Value
protocol	This attribute is required to identify the VSAs as specific to Palo Alto Networks devices. You must set the value to **`firewall`**.
PaloAlto-Admin-Role	A default (dynamic) administrative role name or a custom administrative role name on the firewall.
PaloAlto-Admin-Access-Domain	The name of an access domain for firewall administrators (configured in the **Device** > **Access Domains** page). Define this VSA if the firewall has multiple virtual systems.
PaloAlto-Panorama-Admin-Role	A default (dynamic) administrative role name or a custom administrative role name on Panorama.
PaloAlto-Panorama-Admin-Access-Domain	The name of an access domain for Device Group and Template administrators (configured in the **Panorama** > **Access Domains** page).
PaloAlto-User-Group	The name of a user group in the Allow List of an authentication profile.

RADIUS

Remote Authentication Dial-In User Service (RADIUS) is a broadly supported networking protocol that provides centralized authentication and authorization. You can configure RADIUS authentication for end users or administrators on the firewall and for administrators on Panorama. Optionally, you can use RADIUS Vendor-Specific Attributes (VSAs) to manage administrator authorization. RADIUS VSAs enable you to quickly change the roles, access domains, and user groups of administrators through your directory service instead of reconfiguring settings on the firewall and Panorama. You can also configure the firewall to use a RADIUS server for:

- Collecting VSAs from GlobalProtect endpoints.
- Implementing Multi-Factor Authentication.

When sending authentication requests to a RADIUS server, the firewall and Panorama use the authentication profile name as the network access server (NAS) identifier, even if the profile is assigned to an authentication sequence for the service (such as administrative access to the web interface) that initiates the authentication process.

The firewall and Panorama support the following RADIUS VSAs. To define VSAs on a RADIUS server, you must specify the vendor code (25461 for Palo Alto Networks firewalls or Panorama) and the VSA name and number. Some VSAs also require a value. Refer to your RADIUS server documentation for the steps to define these VSAs.

Alternatively, you can download the Palo Alto Networks RADIUS dictionary, which defines the authentication attributes that the Palo Alto Networks firewall and a RADIUS server use to

https://t.me/learningnets

communicate with each other, and install it on your RADIUS server to map the attributes to the RADIUS binary data.

*When you predefine dynamic administrator roles for users on the server, use lower-case to specify the role (for example, enter **superuser**, not **SuperUser**).*

*When configuring the advanced vendor options on a Cisco Secure Access Control Server (ACS), you must set both the **Vendor Length Field Size** and **Vendor Type Field Size** to **1**. Otherwise, authentication will fail.*

Name	Number	Value
VSAs for administrator account management and authentication		
PaloAlto-Admin-Role	1	A default (dynamic) administrative role name or a custom administrative role name on the firewall.
PaloAlto-Admin-Access-Domain	2	The name of an access domain for firewall administrators (configured in the **Device** > **Access Domains** page). Define this VSA if the firewall has multiple virtual systems.
PaloAlto-Panorama-Admin-Role	3	A default (dynamic) administrative role name or a custom administrative role name on Panorama.
PaloAlto-Panorama-Admin-Access-Domain	4	The name of an access domain for Device Group and Template administrators (configured in the **Panorama** > **Access Domains** page).
PaloAlto-User-Group	5	The name of a user group that an authentication profile references.
VSAs forwarded from GlobalProtect endpoints to the RADIUS server		
PaloAlto-User-Domain	6	Don't specify a value when you define these VSAs.
PaloAlto-Client-Source-IP	7	
PaloAlto-Client-OS	8	
PaloAlto-Client-Hostname	9	
PaloAlto-GlobalProtect-Client-Version	10	

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

LDAP

Lightweight Directory Access Protocol (LDAP) is a standard protocol for accessing information directories. You can Configure LDAP Authentication for end users and for firewall and Panorama administrators.

Configuring the firewall to connect to an LDAP server also enables you to define policy rules based on users and user groups instead of just IP addresses. For the steps, see Map Users to Groups and Enable User- and Group-Based Policy.

Local Authentication

Although the firewall and Panorama provide local authentication for administrators and end users, External Authentication Services are preferable in most cases because they provide central account management. However, you might require special user accounts that you don't manage through the directory servers that your organization reserves for regular accounts. For example, you might define a superuser account that is local to the firewall so that you can access the firewall even if the directory server is down. In such cases, you can use the following local authentication methods:

- **(Firewall only) Local database authentication**—To Configure Local Database Authentication, you create a database that runs locally on the firewall and contains user accounts (usernames and passwords or hashed passwords) and user groups. This type of authentication is useful for creating user accounts that reuse the credentials of existing Unix accounts in cases where you know only the hashed passwords, not the plaintext passwords. Because local database authentication is associated with authentication profiles, you can accommodate deployments where different sets of users require different authentication settings, such as Kerberos single sign-on (SSO) or Multi-Factor Authentication (MFA). (For details, see Configure an Authentication Profile and Sequence). For administrator accounts that use an authentication profile, password complexity and expiration settings are not applied. This authentication method is available to administrators who access the firewall (but not Panorama) and end users who access services and applications through Authentication Portal or GlobalProtect.
- **Local authentication without a database**—You can configure firewall administrative accounts or Panorama administrative accounts without creating a database of users and user groups that runs locally on the firewall or Panorama. Because this method is not associated with authentication profiles, you cannot combine it with Kerberos SSO or MFA. However, this is the only authentication method that allows password profiles, which enable you to associate individual accounts with password expiration settings that differ from the global settings. (For details, see Define password complexity and expiration settings)

https://t.me/learningnets

Plan Your Authentication Deployment

The following are key questions to consider before you implement an authentication solution for administrators who access the firewall and end users who access services and applications through Authentication Portal.

For both end users and administrators, consider:

- How can you leverage your existing security infrastructure? Usually, integrating the firewall with an existing infrastructure is faster and cheaper than setting up a new, separate solution just for firewall services. The firewall can integrate with Multi-Factor Authentication, SAML, Kerberos, TACACS+, RADIUS, and LDAP servers. If your users access services and applications that are external to your network, you can use SAML to integrate the firewall with an identity provider (IdP) that controls access to both external and internal services and applications.
- How can you optimize the user experience? If you don't want users to authenticate manually and you have a public key infrastructure, you can implement certificate authentication. Another option is to implement Kerberos or SAML single sign-on (SSO) so that users can access multiple services and applications after logging in to just one. If your network requires additional security, you can combine certificate authentication with interactive (challenge-response) authentication.
- Do you require special user accounts that you don't manage through the directory servers that your organization reserves for regular accounts? For example, you might define a superuser account that is local to the firewall so that you can access the firewall even if the directory server is down. You can configure Local Authentication for these special-purpose accounts.

  External Authentication Services *are usually preferable to local authentication because they provide central account management, reliable authentication services, and usually logging and troubleshooting features.*
- Are the user names for your user accounts properly formatted? Leveraging SAML, Kerberos, TACACS+, RADIUS, and LDAP authentication requires all user names adhere to the regular expression Linux login name rule. User names must have the format **`[a-zA-Z0-9_.][a-zA-Z0-9_.-]{0,30}[a-zA-Z0-9_.$-]`**.

  This means that:

  - The first character of the user name must be an upper or lower case alphabetical letter, a number (0-9), or either _ (underscore) or **`.`** (period).
  - Other than the first and last characters, the user name may contain upper or lower case alphabetical characters, numbers (0-9), and _ (underscore), **`.`** (period), or **`-`** (dash). The maximum length is 30 characters excluding the first and last characters.
  - The last character of the user name may be an upper or lower case alphabetical letter, a number (0-9), or _ (underscore), **`.`** (period), **`$`**, or **`-`** (dash).

  Adhering to the regular expression Linux login name rule is required for PAN-OS administrators only. It is not required for GlobalProtect and Captive Portal users.

For end users only, consider:

- Which services and applications are more sensitive than others? For example, you might want stronger authentication for key financial documents than for search engines. To protect

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

your most sensitive services and applications, you can configure Multi-Factor Authentication (MFA) to ensure that each user authenticates using multiple methods (factors) when accessing those services and applications. To accommodate a variety of security needs, Configure Authentication Policy rules that trigger MFA or single factor authentication (such as login credentials or certificates) based on specific services, applications, and end users. Other ways to reduce your attack surface include network segmentation and user groups for allowed applications.

For administrators only, consider:

- Do you use an external server to centrally manage authorization for all administrative accounts? By defining Vendor-Specific Attributes (VSAs) on the external server, you can quickly change administrative role assignments through your directory service instead of reconfiguring settings on the firewall. VSAs also enable you to specify access domains for administrators of firewalls with multiple virtual systems. SAML, TACACS+, and RADIUS support external authorization.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Configure Multi-Factor Authentication

To use Multi-Factor Authentication (MFA) for protecting sensitive services and applications, you must configure Authentication Portal to display a web form for the first authentication factor and to record Authentication Timestamps. The firewall uses the timestamps to evaluate the timeouts for Authentication Policy rules. To enable additional authentication factors, you can integrate the firewall with MFA vendors through RADIUS or vendor APIs. After evaluating Authentication policy, the firewall evaluates Security policy, so you must configure rules for both policy types.

Palo Alto Networks provides support for MFA vendors through Applications content updates. This means that if you use Panorama to push device group configurations to firewalls, you must install the same Applications updates on the firewalls as on Panorama to avoid mismatches in vendor support.

MFA vendor API integrations are supported for end-user authentication through Authentication Policy only. For remote user authentication to GlobalProtect portals or gateways or for administrator authentication to the PAN-OS or Panorama web interface, you can only use MFA vendors supported through RADIUS or SAML; MFA services through vendor APIs are not supported in these use cases.

STEP 1 | Configure Authentication Portal in **Redirect** mode to display a web form for the first authentication factor, to record authentication timestamps, and to update user mappings.

STEP 2 | Configure one of the following server profiles to define how the firewall will connect to the service that authenticates users for the first authentication factor.

- Add a RADIUS server profile. This is required if the firewall integrates with an MFA vendor through RADIUS. In this case, the MFA vendor provides the first and all additional authentication factors, so you can skip the next step (configuring an MFA server profile). If the firewall integrates with an MFA vendor through an API, you can still use a RADIUS server profile for the first factor but MFA server profiles are required for the additional factors.
- Add a SAML IdP server profile.
- Add a Kerberos server profile.
- Add a TACACS+ server profile.
- Add an LDAP server profile.

In most cases, an external service is recommended for the first authentication factor. However, you can configure Configure Local Database Authentication as an alternative.

STEP 3 | Add an MFA server profile.

The profile defines how the firewall connects to the MFA server. Add a separate profile for each authentication factor after the first factor. The firewall integrates with these MFA servers

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

through vendor APIs. You can specify up to three additional factors. Each MFA vendor provides one factor, though some vendors let users choose one factor out of several.

1. Select **Device** > **Server Profiles** > **Multi Factor Authentication** and **Add** a profile.
2. Enter a **Name** to identify the MFA server.
3. Select the **Certificate Profile** that the firewall will use to validate the MFA server certificate when establishing a secure connection to the MFA server.
4. Select the **MFA Vendor** you deployed.
5. Configure the **Value** of each vendor attribute.

   The attributes define how the firewall connects to the MFA server. Each vendor **Type** requires different attributes and values; refer to your vendor documentation for details.
6. Click **OK** to save the profile.

STEP 4 | Configure an authentication profile.

The profile defines the order of the authentication factors that users must respond to.

1. Select **Device** > **Authentication Profile** and **Add** a profile.
2. Enter a **Name** to identify the authentication profile.
3. Select the **Type** for the first authentication factor and select the corresponding **Server Profile**.
4. Select **Factors**, **Enable Additional Authentication Factors**, and **Add** the MFA server profiles you configured.

   The firewall will invoke each MFA service in the listed order, from top to bottom.
5. Click **OK** to save the authentication profile.

STEP 5 | Configure an authentication enforcement object.

The object associates each authentication profile with an Authentication Portal method. The method determines whether the first authentication challenge (factor) is transparent or requires a user response.

Select the **Authentication Profile** you configured and enter a **Message** that tells users how to authenticate for the first factor. The message displays in the Authentication Portal web form.

*If you set the **Authentication Method** to **browser-challenge**, the Authentication Portal web form displays only if Kerberos SSO authentication fails. Otherwise, authentication for the first factor is automatic; users won't see the web form.*

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 6 | Configure an Authentication policy rule.

The rule must match the services and applications you want to protect and the users who must authenticate.

1. Select **Policies** > **Authentication** and **Add** a rule.
2. Enter a **Name** to identify the rule.
3. Select **Source** and **Add** specific zones and IP addresses or select **Any** zones or IP addresses.

   The rule applies only to traffic coming from the specified IP addresses or from interfaces in the specified zones.
4. Select **User** and select or **Add** the source users and user groups to which the rule applies (default is **any**).
5. Select **Destination** and **Add** specific zones and IP addresses or select **any** zones or IP addresses.

   The IP addresses can be resources (such as servers) for which you want to control access.
6. Select **Service/URL Category** and select or **Add** the services and service groups for which the rule controls access (default is **service-http**).
7. Select or **Add** the URL Categories for which the rule controls access (default is **any**). For example, you can create a custom URL category that specifies your most sensitive internal sites.
8. Select **Actions** and select the **Authentication Enforcement** object you created.
9. Specify the **Timeout** period in minutes (default 60) during which the firewall prompts the user to authenticate only once for repeated access to services and applications.

   ***Timeout** is a tradeoff between tighter security (less time between authentication prompts) and the user experience (more time between authentication prompts). More frequent authentication is often the right choice for access to critical systems and sensitive areas such as a data center. Less frequent authentication is often the right choice at the network perimeter and for businesses for which the user experience is key.*
10. Click **OK** to save the rule.

STEP 7 | Customize the MFA login page.

The firewall displays this page to tell users how to authenticate for MFA factors and to indicate the authentication status (in progress, succeeded, or failed).

1. Select **Device** > **Response Pages** and select **MFA Login Page**.
2. Select the **Predefined** response page and **Export** the page to your client system.
3. On your client system, use an HTML editor to customize the downloaded response page and save it with a unique filename.
4. Return to the MFA Login Page dialog on the firewall, **Import** your customized page, **Browse** to select the **Import File**, select the **Destination** (virtual system or **shared** location), click **OK**, and click **Close**.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 8 | Configure a Security policy rule that allows users to access the services and applications that require authentication.

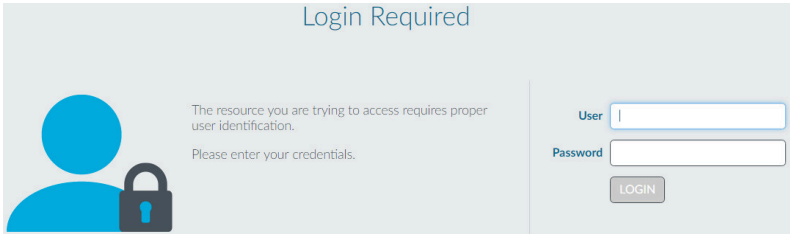
1. Create a Security Policy Rule.
2. **Commit** your changes.

   *The automated correlation engine on the firewall uses several correlation objects to detect events on your network that could indicate credential abuse relating to MFA. To review the events, select **Monitor** > **Automated Correlation Engine** > **Correlated Events**.*

STEP 9 | Verify that the firewall enforces MFA.

1. Log in to your network as one of the source users specified in the Authentication rule.
2. Request a service or application that matches one of the services or applications specified in the rule.

   The firewall displays the Authentication Portal web form for the first authentication factor. The page contains the message you entered in the authentication enforcement object. For example:



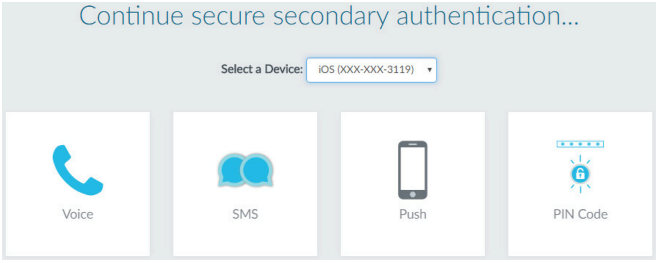
3. Enter your user credentials for the first authentication challenge.

   The firewall then displays an MFA login page for the next authentication factor. For example, the MFA service might prompt you to select the Voice, SMS, push, or PIN code

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

(OTP) authentication method. If you select push, your phone prompts you to approve the authentication.



4. Authenticate for the next factor.

   The firewall displays an authentication success or failure message. If authentication succeeded, the firewall displays an MFA login page for the next authentication factor, if any.

   Repeat this step for each MFA factor. After you authenticate for all the factors, the firewall evaluates Security policy to determine whether to allow access to the service or application.
5. End the session for the service or application you just accessed.
6. Start a new session for the same service or application. Be sure to perform this step within the **Timeout** period you configured in the Authentication rule.

   The firewall allows access without re-authenticating.
7. Wait until the **Timeout** period expires and request the same service or application.

   The firewall prompts you to re-authenticate.

Configure MFA Between RSA SecurID and the Firewall

Multi-factor authentication allows you to protect company assets by using multiple factors to verify a user's identity before allowing them to access network resources. To enable multi-factor authentication (MFA) between the firewall and the RSA SecurID Access Cloud Authentication Service, you must first configure the RSA SecurID Service so that you have the details that you need to configure the firewall to authenticate users using multiple factors. After you have performed the required configuration on the RSA SecurID Access Console, you can configure the firewall to integrate with RSA SecurID.

The Palo Alto Networks next-generation firewall integrates with the RSA SecurID Access Cloud Authentication Service. The MFA API integration with RSA SecurID is supported for cloud-based services only and does not support two-factor authentication for the on-premise Authentication Manager when the second factor uses the Vendor Specific API. The minimum content version required for this integration is 752 and PAN-OS 8.0.2.

- Get the RSA SecurID Access Cloud Authentication Service Details
- Configure the Firewall for MFA with RSA SecurID

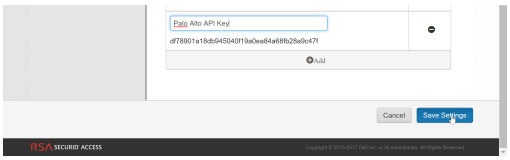
Get the RSA SecurID Access Cloud Authentication Service Details

In order to securely pass user authentication requests to and from the firewall and the RSA SecurID Access Cloud Authentication Service, you must first go to the RSA SecurID Access

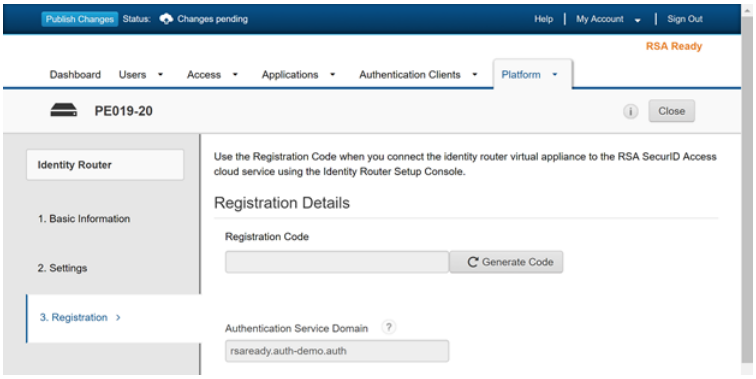
https://t.me/learningnets

Console and configure the RSA Access ID, the authentication service URL, and the client API key that the firewall needs to authenticate to and interact with the service. The firewall also needs the Access Policy ID that uses either the RSA Approve or RSA Tokencode authentication method to authenticate to the identity source.

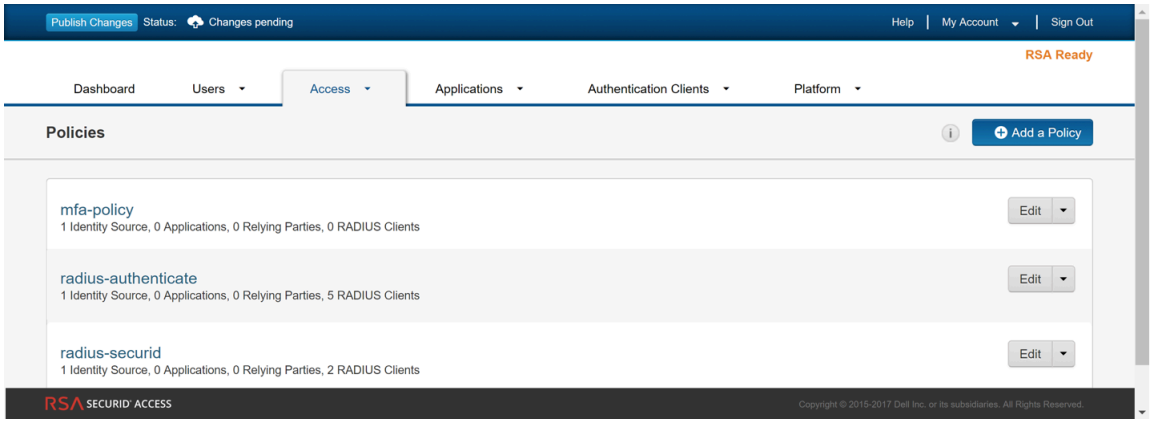
- **Generate the RSA SecurID API key**—Log on to RSA SecurID Access Console and select **My Account** > **Company Settings** > **Authentication API Keys**. **Add** a new key and then **Save Settings** and **Publish Changes**.



- **Get the RSA SecurID Access endpoint API (Authentication Service Domain) to which the firewall must connect**—Select **Platform** > **Identity Routers**, pick an Identity Router to **Edit** and jot down the **Authentication Service Domain**. In this example it is https://rsaready.auth-demo.auth.



- **Get the Access Policy ID**—Select **Access** > **Policies** and jot down the name of the access policy that will allow the firewall to act as an authentication client to the RSA SecurID service. The policy must be configured to use either the RSA Approve or the RSA Tokencode authentication methods only.



https://t.me/learningnets

Configure the Firewall for MFA with RSA SecurID

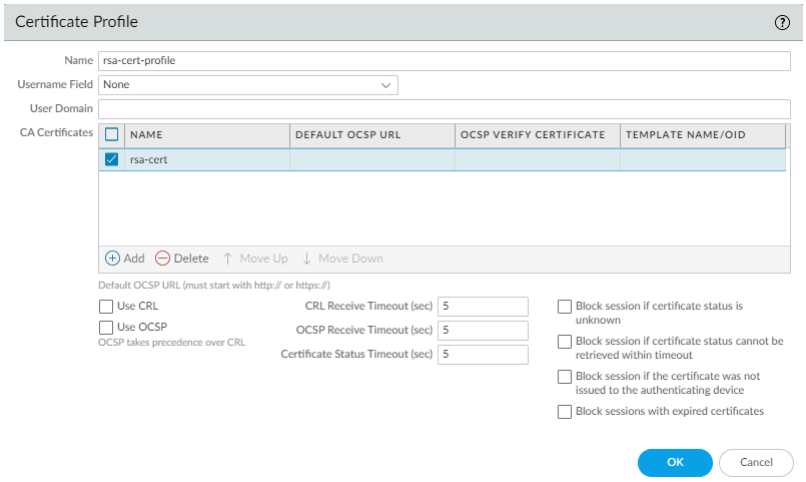
After you Get the RSA SecurID Access Cloud Authentication Service Details, you can configure the firewall to prompt users for an RSA SecurID token when MFA is invoked.

STEP 1 | Configure the firewall to trust the SSL certificate provided by the RSA SecurID Access endpoint API.

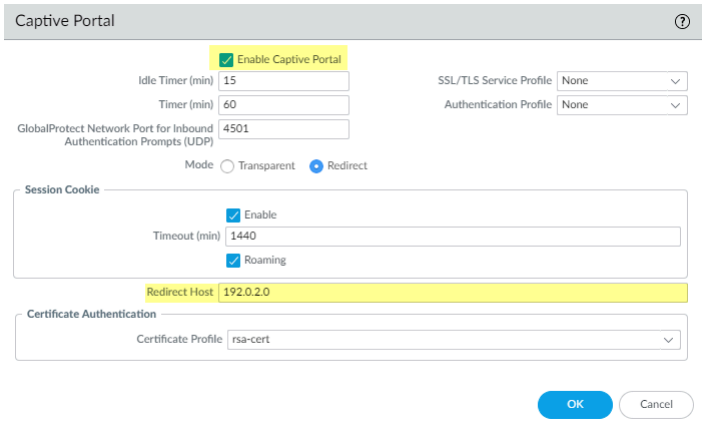
1. Export the SSL certificate from the RSA SecurID Access endpoint and import it into the firewall.

   To enable trust between the firewall and the RSA SecurID Access endpoint API, you must either import a self-signed certificate, or the CA certificate used to sign the certificate.

2. Configure a Certificate Profile (**Device** > **Certificate Management** > **Certificate Profile** and click **Add**).



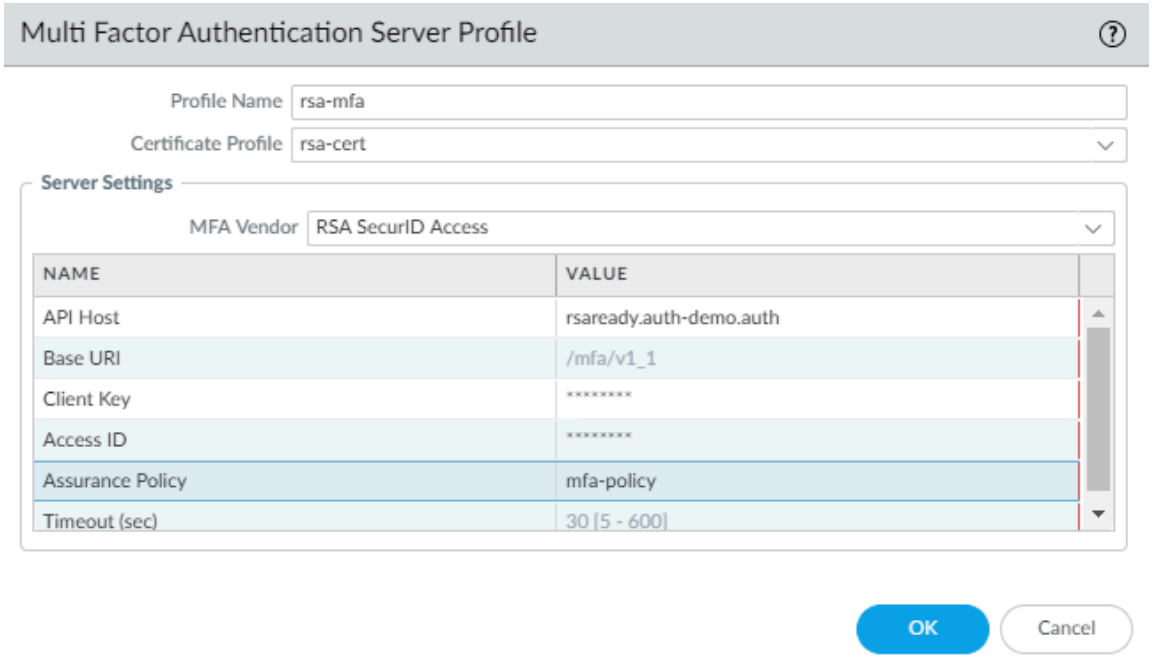
STEP 2 | Configure Authentication Portal (**Device** > **User Identification** > **Authentication Portal Settings**) in Redirect mode to display a web form for authenticating to RSA SecureID. Make sure to specify the Redirect Host as an IP address or a hostname (with no periods in its name) that resolves to the IP address of the Layer 3 interface on the firewall to which web requests are redirected.



https://t.me/learningnets

STEP 3 | Configure a multi-factor authentication server profile to specify how the firewall must connect with the RSA SecurID cloud service **(Device** > **Server Profiles** > **Multi Factor Authentication** and click **Add**).

1. Enter a **Name** to identify the MFA server profile.
2. Select the **Certificate Profile** that you created earlier, rsa-cert-profile in this example. The firewall will use this certificate when establishing a secure connection with RSA SecurID cloud service.
3. In the **MFA Vendor** drop-down, select **RSA SecurID Access**.
4. Configure the **Value** for each attribute that you noted in Get the RSA SecurID Access Cloud Authentication Service Details:
   - **API Host**—Enter the hostname or IP address of the RSA SecurID Access API endpoint to which the firewall must connect, rsaready.auth-demo.auth in this example.
   - **Base URI** —Do not modify the default value (/mfa/v1_1)
   - **Client Key**—Enter the RSA SecurID Client Key.
   - **Access ID**—Enter the RSA SecurID Access ID.
   - **Assurance Policy**—Enter the RSA SecurID Access Policy name, mfa-policy in this example.
   - **Timeout**—The default is 30 seconds.



5. Save the profile.

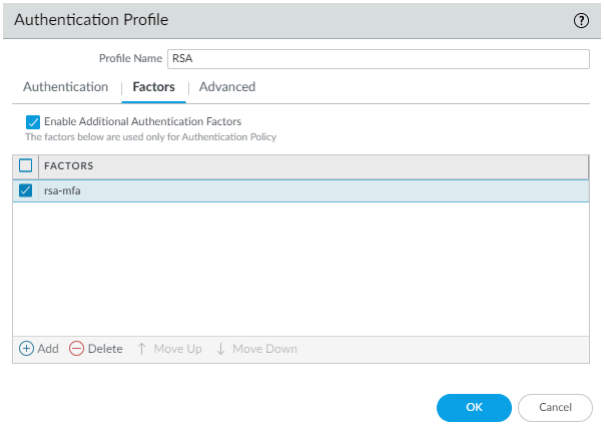
 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 4 | Configure an authentication profile (**Device** > **Authentication Profile** and click **Add**).

The profile defines the order of the authentication factors that users must respond to.

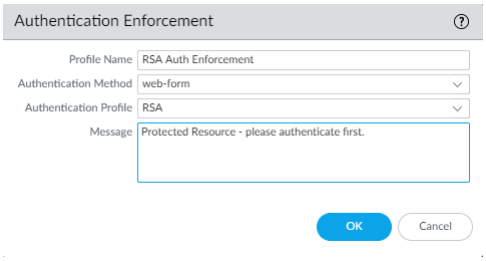
1. Select the **Type** for the first authentication factor and select the corresponding **Server Profile**.
2. Select **Factors**, **Enable Additional Authentication Factors**, and **Add** the rsa-mfa server profile you created earlier in this example.



3. Click **OK** to save the authentication profile.

STEP 5 | Configure an authentication enforcement object. (**Objects** > **Authentication** and click **Add**).

Make sure to select the authentication profile you just defined called RSA in this example.

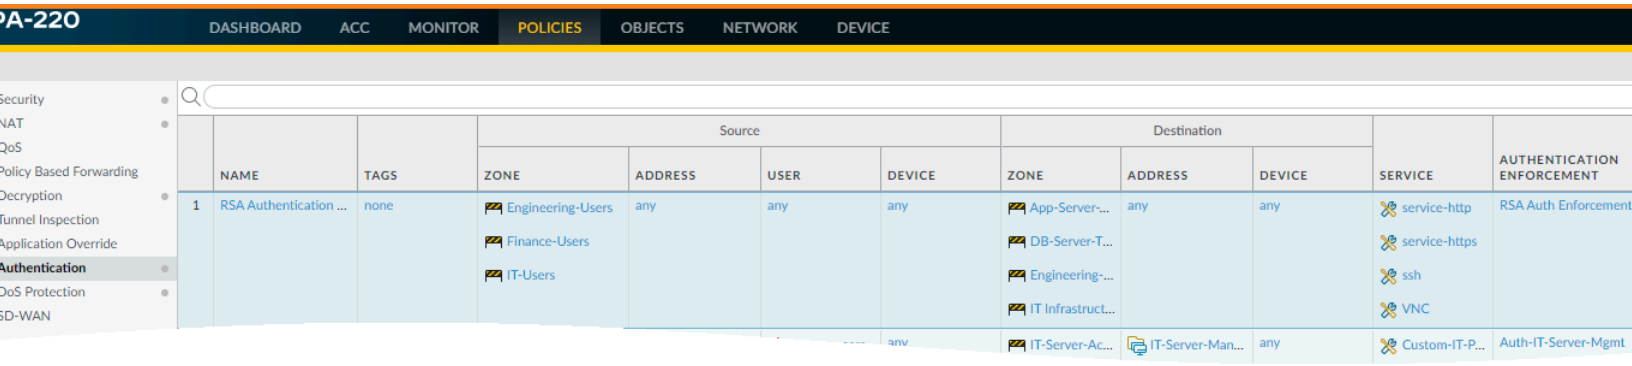


STEP 6 | Configure an Authentication policy rule. (**Policies** > **Authentication** and click **Add**)

Your authentication policy rule must match the services and applications you want to protect, specify the users who must authenticate, and include the authentication enforcement object that triggers the authentication profile. In this example, RSA SecurID Access authenticates all users who accessing HTTP, HTTPS, SSH, and VNC traffic with the authentication enforcement

https://t.me/learningnets

object called RSA Auth Enforcement (in **Actions**, select the **Authentication Enforcement** object).



STEP 7 | **Commit** your changes on the firewall.

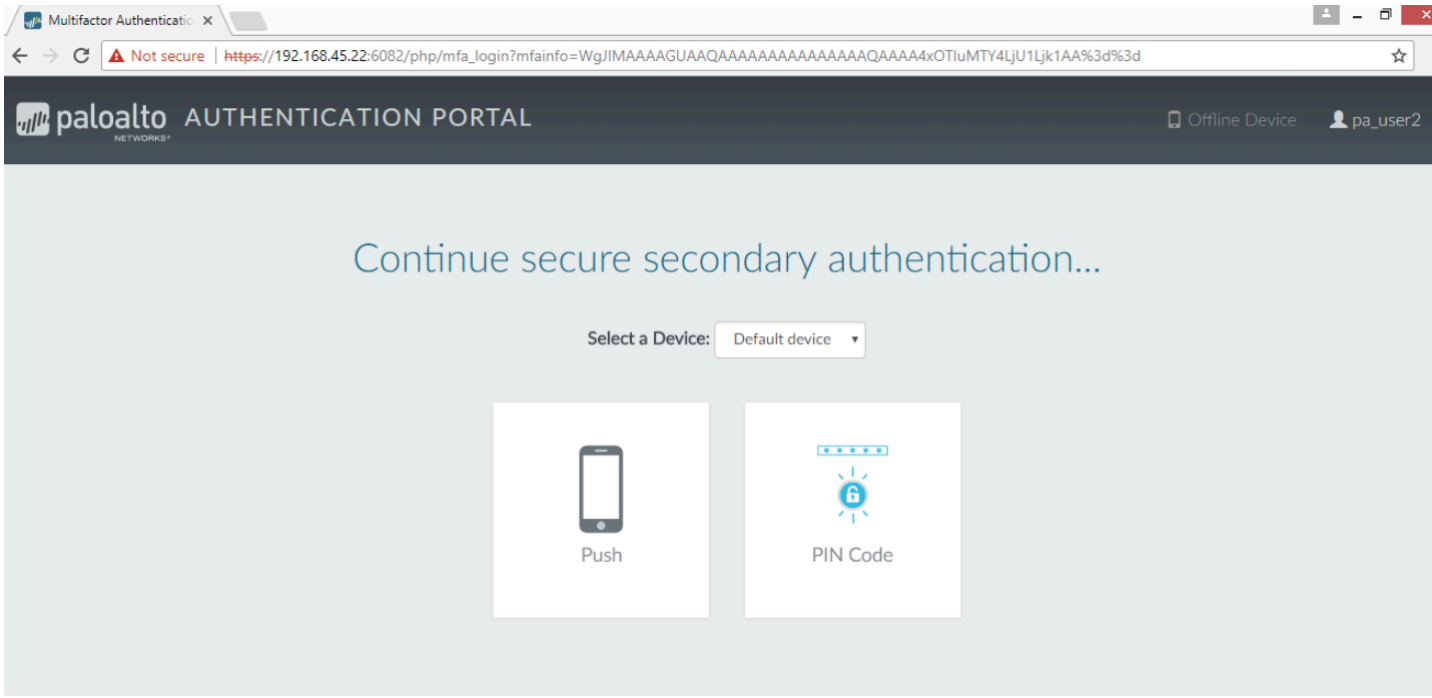
 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 8 | Verify that users on your network are being secured using RSA SecurID using the Push or PIN Code authentication method you enabled.

1. **Push authentication**

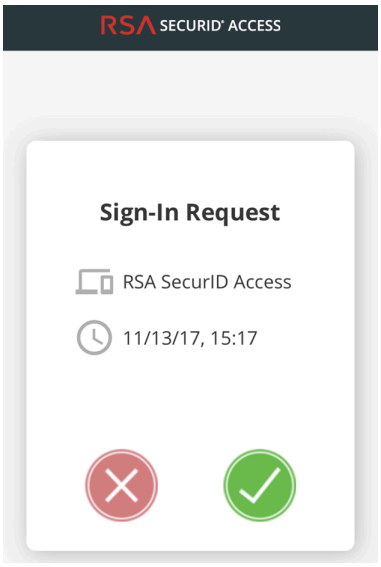
   1. Ask a user on your network to launch a web browser and access a website. The Authentication Portal page with the IP address or hostname for the Redirect Host you defined earlier should display.
   2. Verify that the user enters the credentials for the first authentication factor and then continues to the secondary authentication factor, and selects **Push**.



   3. Check for a **Sign-In request** on the RSA SecurID Access application on the user's mobile device.
   4. Ask the user to **Accept** the Sign-In Request on the mobile device, and wait for a few seconds for the firewall to receive the notification of successful authentication. The user should be able to access the requested website.

 ©2022 Palo Alto Networks, Inc.

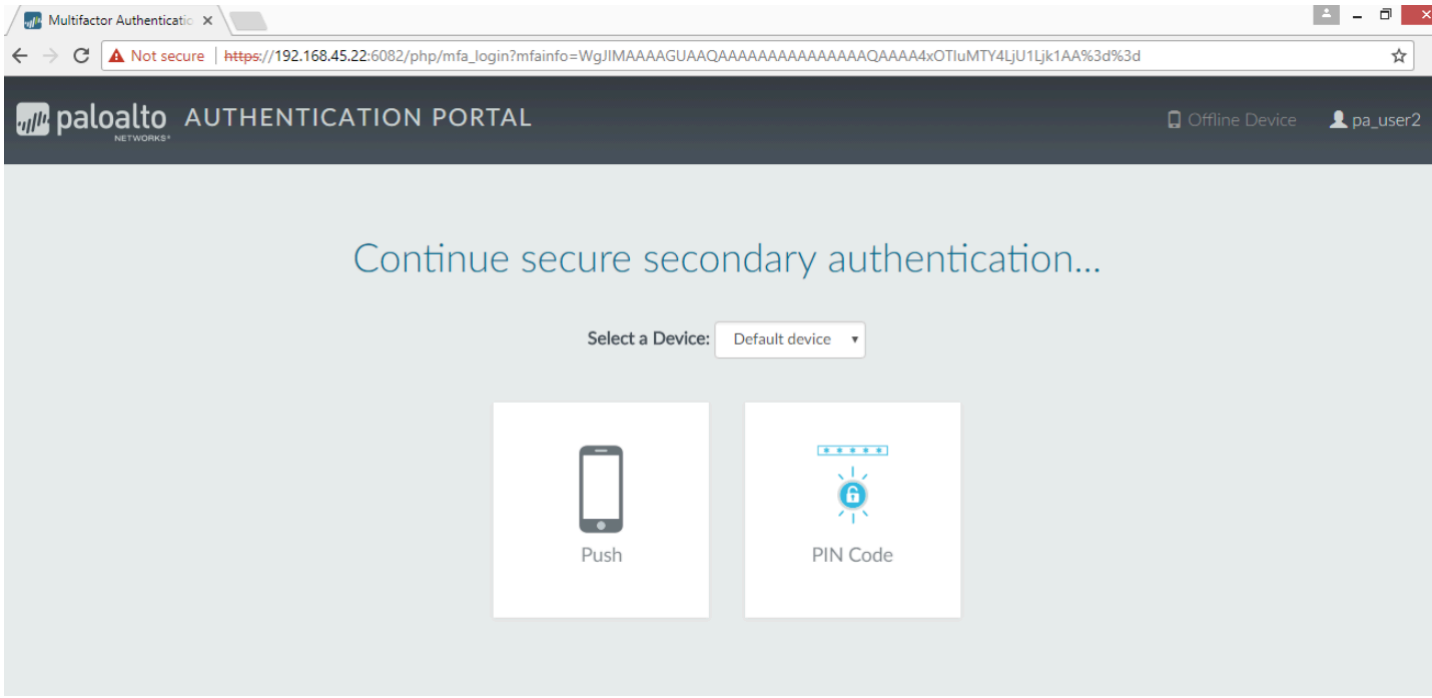
https://t.me/learningnets



*To test an authentication failure, **Decline** the sign-in request on the mobile device.*

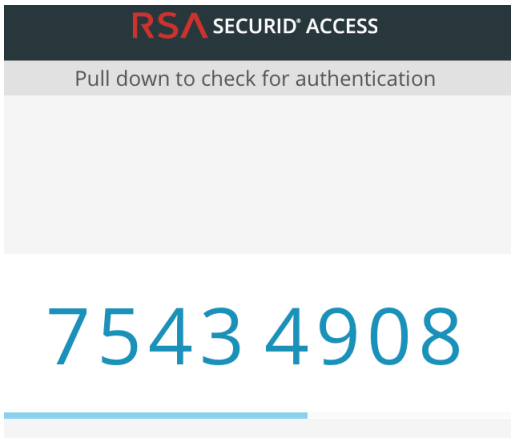
2. **PIN Code authentication**

   1. Ask a user on your network to launch a web browser and access a website. The Authentication Portal page with the IP address or hostname for the redirect host you defined earlier should display.
   2. Verify that the user enters the credentials for the first authentication factor and then continues to the secondary authentication factor, and selects **PIN Code**.



   3. Check that a **PIN Code** displays on the RSA SecurID Access application on the user's mobile device.

https://t.me/learningnets



4. Ask the user to copy the PIN code in the **Enter the PIN...** prompt of the web browser and click **Submit**. Wait for a few seconds for the firewall to receive the notification of successful authentication. The user should be able to access the requested website.

Configure MFA Between Okta and the Firewall

Multi-factor authentication allows you to protect company assets by using multiple factors to verify the identity of users before allowing them to access network resources.

To enable multi-factor authentication (MFA) between the firewall and the Okta identity management service:

- Configure Okta
- Configure the firewall to integrate with Okta
- Verify MFA with Okta

Configure Okta

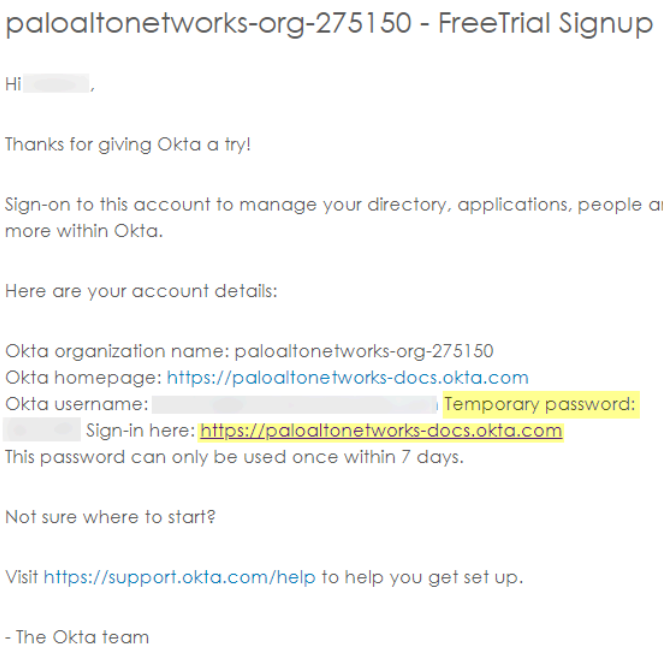
Log in to the Okta Admin Portal to create your user accounts, define your Okta MFA policy, and obtain the token information required to configure MFA with Okta on the firewall.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 1 | Create your Okta Admin user account.

1. Submit your email address and name, then click **Get Started**.
2. Click the link in the confirmation email and use the included temporary password to log in to the Okta Admin Portal.



paloaltonetworks-org-275150 - FreeTrial Signup

Hi [redacted],

Thanks for giving Okta a try!

Sign-on to this account to manage your directory, applications, people and more within Okta.

Here are your account details:

Okta organization name: paloaltonetworks-org-275150
Okta homepage: https://paloaltonetworks-docs.okta.com
Okta username: [redacted] Temporary password: [redacted]
Sign-in here: https://paloaltonetworks-docs.okta.com
This password can only be used once within 7 days.

Not sure where to start?

Visit https://support.okta.com/help to help you get set up.

- The Okta team

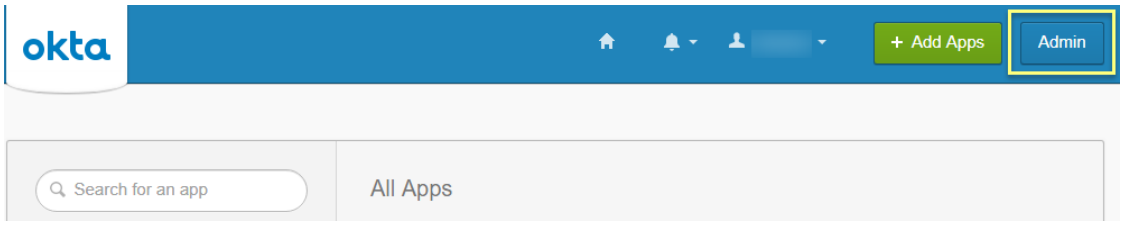
3. Create a new password that includes at least 8 characters, one lowercase letter, one uppercase letter, a number, and does not include any part of your username.
4. Select a password reminder question and enter the answer.
5. Select a security image, then **Create My Account**.

 ©2022 Palo Alto Networks, Inc.

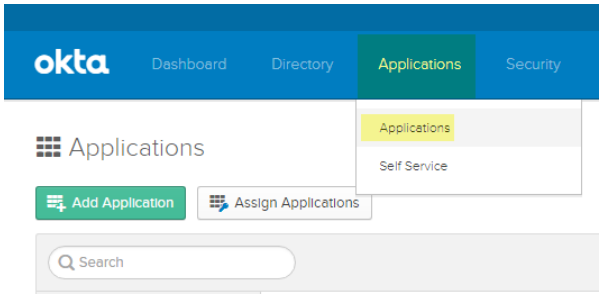
https://t.me/learningnets

STEP 2 | Configure your Okta service.

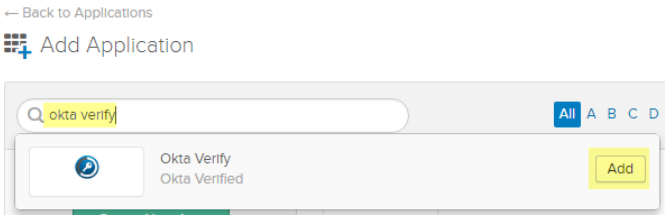
*If you log in and are not redirected to the Okta Admin Portal, select **Admin** at the upper right.*



1. From the Okta Dashboard, log in with your Okta Admin credentials, then select **Applications** > **Applications**.



2. Select **Add Application**.
3. Search for **`Okta Verify`**.
4. Select **Add**, then **Done**.

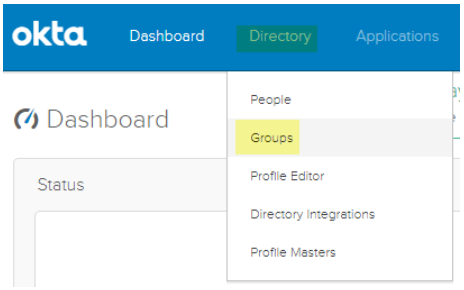


 ©2022 Palo Alto Networks, Inc.

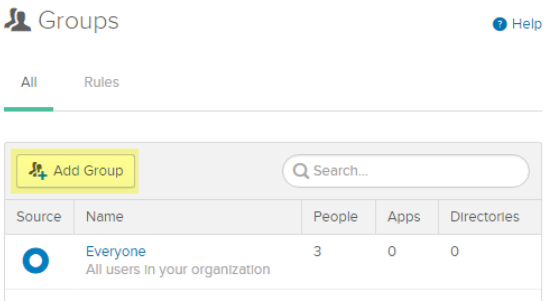
https://t.me/learningnets

STEP 3 | Create one or more user groups to categorize your users (for example, by device, by policy, or by department) and assign the Okta Verify application.

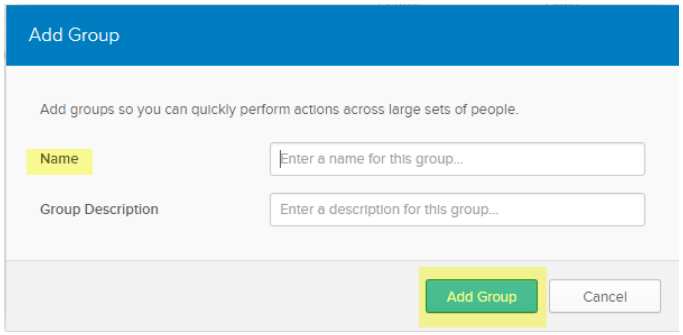
1. Select **Directory** > **Groups**.



2. Click **Add Group**.

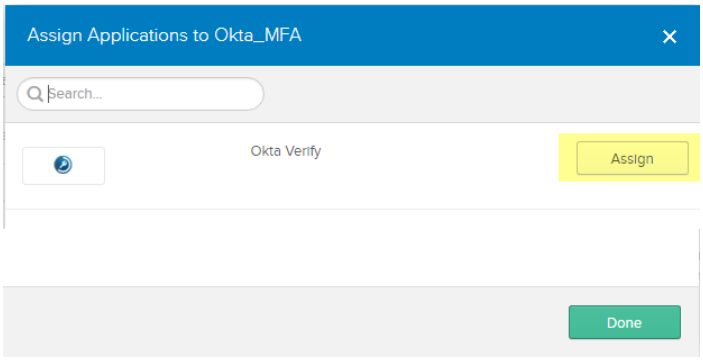


3. Enter a group **Name** and optionally a **Group Description**, then **Add Group**.



*The default group **Everyone** includes all users configured for your organization during the first step in Configure Okta.*

4. Select the group you created, then select **Manage Apps**.
5. **Assign** the Okta Verify application you added in Step 2.



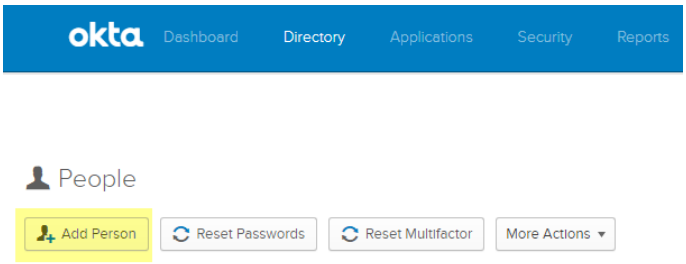
 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

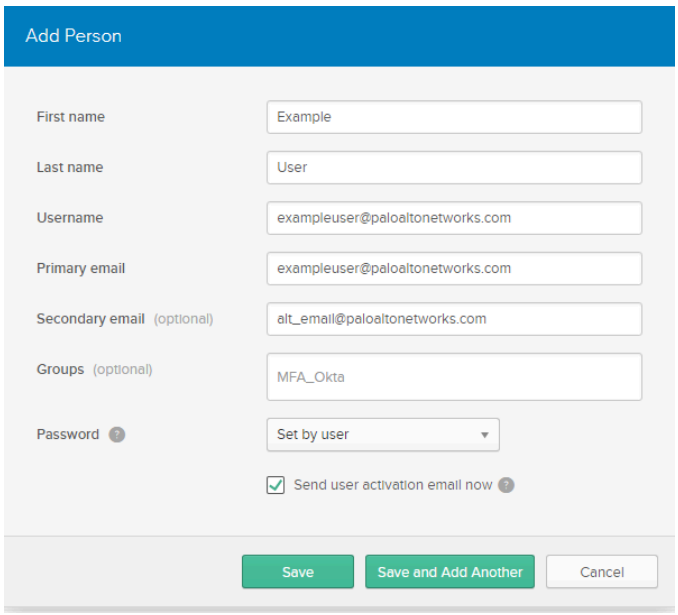
6. After the application has been **Assigned**, click **Done**.
7. Repeat this process for all groups that will use the Okta Verify application for MFA.

STEP 4 | Add users and assign them to a group.

1. From the Okta Dashboard, select **Directory** > **People** > **Add Person**.



2. Enter the user's **First Name**, **Last Name**, and **Username**. The username must match the **Primary email**, which populates automatically, and the username entered on the firewall. You can optionally enter an alternate email address for the user as the **Secondary Email**.



3. Enter the name of the group or **Groups** to associate with this user. When you start typing, the group name populates automatically.
4. Check **Send user activation email now**, then **Save** to add a single user or **Save and Add Another** to continue adding users.

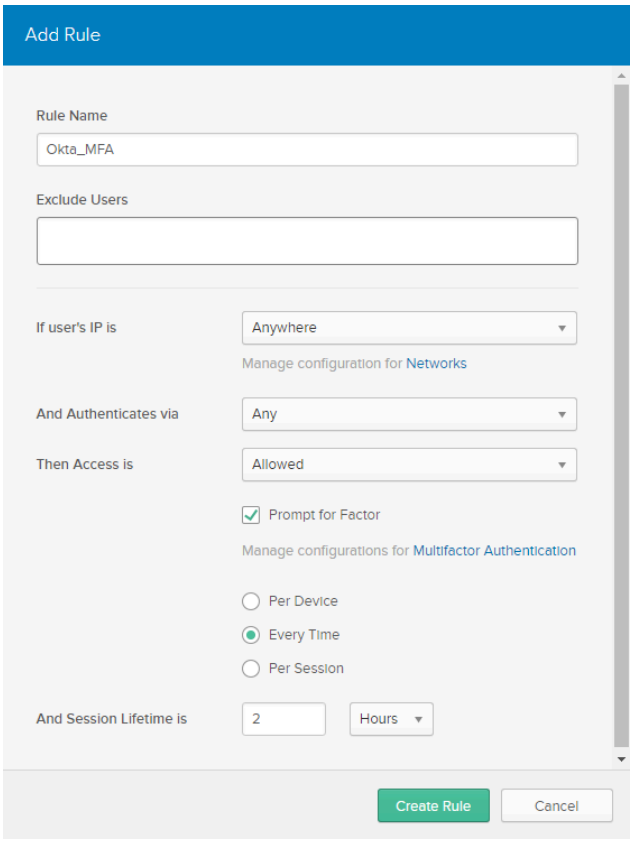
https://t.me/learningnets

STEP 5 | Assign a test policy to users.

1. Select **Security** > **Authentication** > **Sign On**.

   There is a **Default Policy** with a **Default Rule** that does not prompt users to log in with MFA.

2. Enter the **Rule Name** and check **Prompt for Factor** to enforce the MFA prompt, and select the type of prompt (**Per Device**, **Every Time**, or **Per Session**), then **Create Rule**.

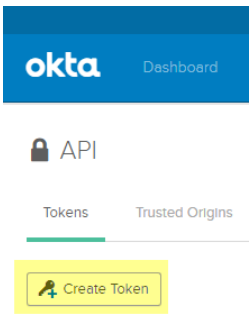


 ©2022 Palo Alto Networks, Inc.

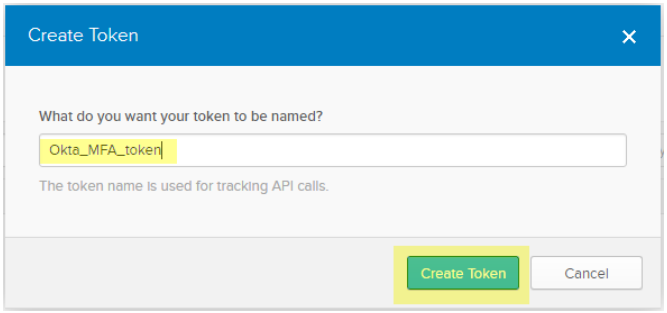
https://t.me/learningnets

STEP 6 | Record the Okta authentication token information in a safe place because it is only displayed once.

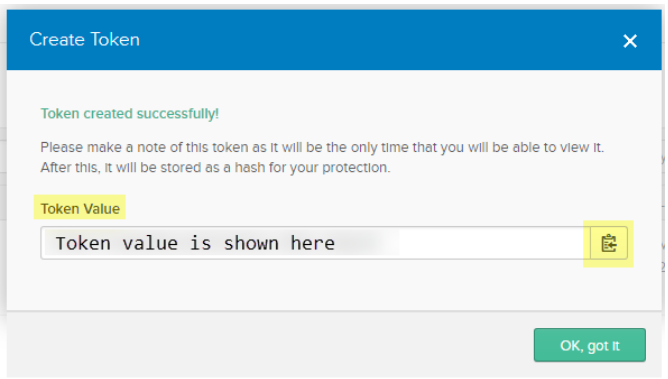
1. Select **Security** > **API** > **Tokens**.
2. Select **Create Token**.



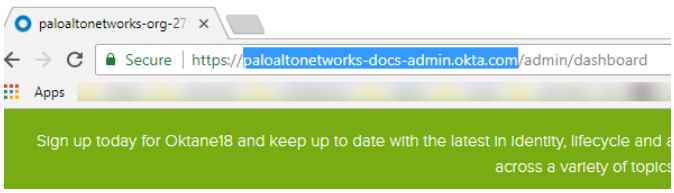
3. Enter a name for the token, then **Create Token**.



4. Copy the **Token Value**.
   You can click the **Copy to clipboard** button to copy the Token Value to your clipboard.

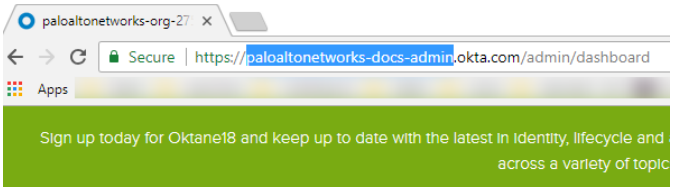


5. In the URL for the Okta Admin Dashboard, copy the portion of the URL after **`https://`** up to **`/admin`** to use as the **API host**.



6. Omit the domain **`okta.com`** from this URL to use as the **Organization**.

https://t.me/learningnets



For example, in the example Okta Admin Dashboard URL above, **`https://paloaltonetworks-doc-admin.okta.com/admin/dashboard`**:

- The API hostname is **`paloaltonetworks-doc-admin.okta.com`**.
- The Organization is **`paloaltonetworks-doc-admin`**.

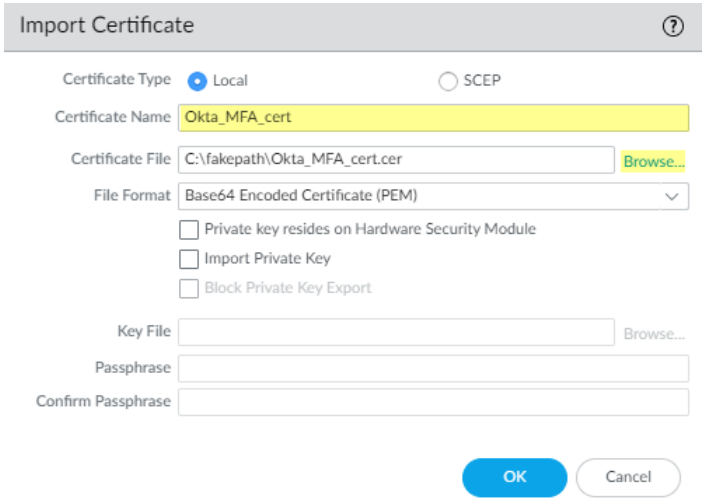
STEP 7 | Export all certificates in the certificate chain using Base-64 encoding:

1. Depending on your browser, use one of the following methods to export all certificates in the chain.
   - **Chrome**—Press **F12**, then select **Security** > **View Certificate** > **Details** > **Copy to File**.
   - **Firefox**—Select **Options** > **Privacy & Security** > **View Certificates** > **Export**.
   - **Internet Explorer**—Select **Settings** > **Internet Options** > **Content** > **Certificates** > **Export**.
2. Use the Certificate Export Wizard to export all certificates in the chain and select **Base-64 encoded X.509** as the format.

Configure the firewall to integrate with Okta

As a prerequisite, confirm that you have mapped all users that you want to authenticate using Okta.

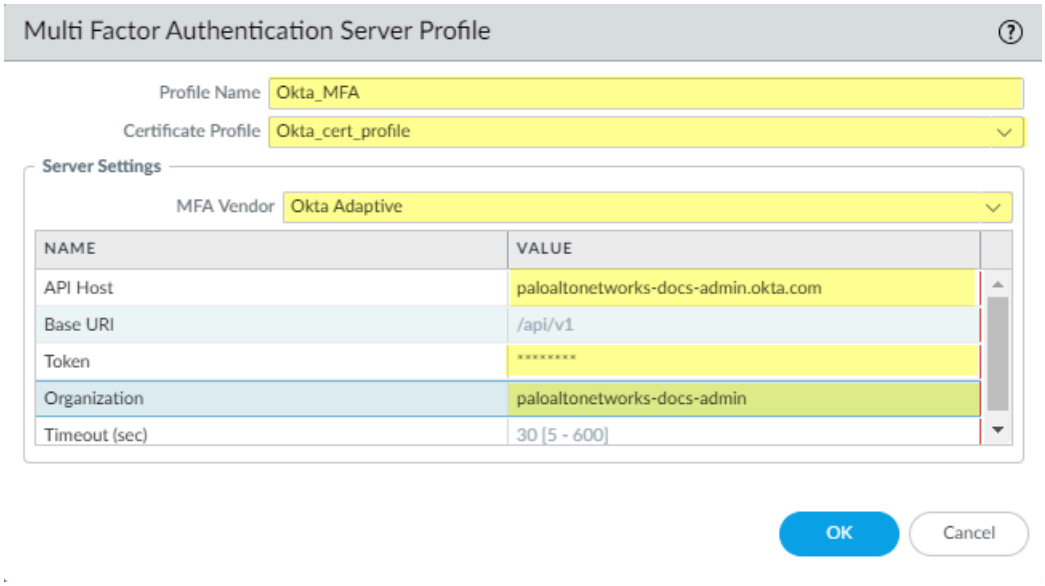
STEP 1 | Import all certificates in the certificate chain on the firewall and add the imported CA certificates (root and intermediate) to a Certificate Profile.



https://t.me/learningnets

STEP 2 | Add a **Multi Factor Authentication Server Profile** for Okta.

1. Select **Device** > **Server Profiles** > **Multi Factor Authentication**.
2. **Add** an MFA server profile.

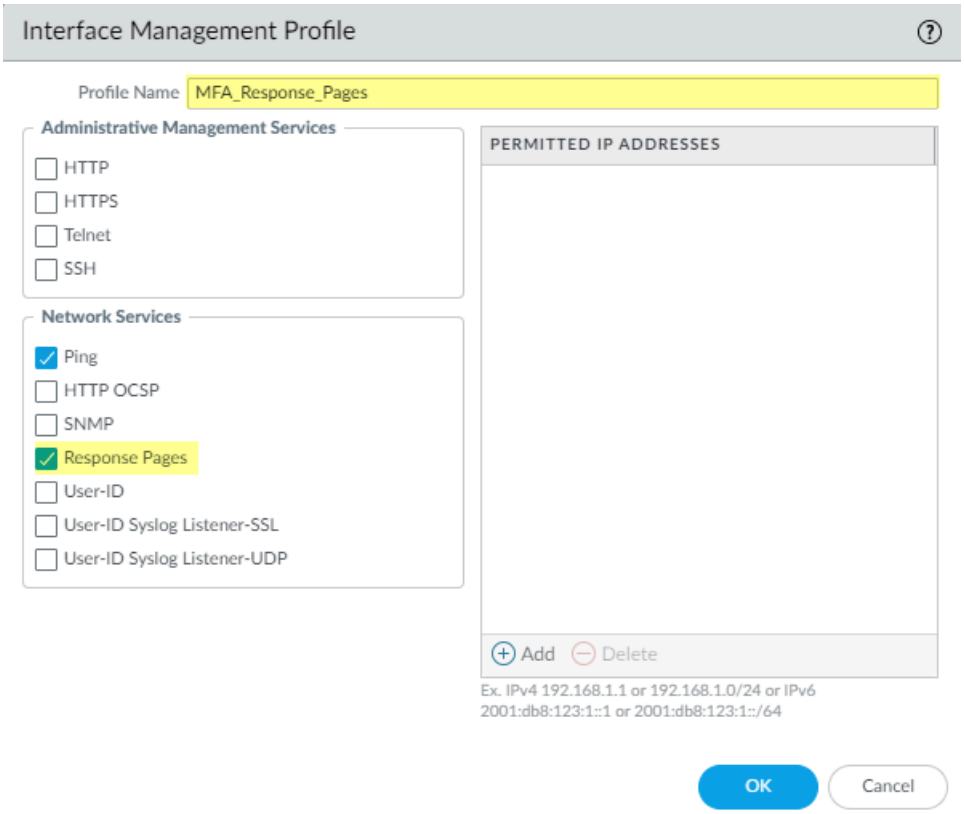


3. Enter a **Profile Name**.
4. Select the **Certificate Profile** you created in Step 1 in Configure the firewall to integrate with Okta.
5. Select **Okta Adaptive** as the **MFA Vendor**.
6. Enter the **API Host**, **Token**, and **Organization** from Step 4 in Configure the firewall to integrate with Okta.

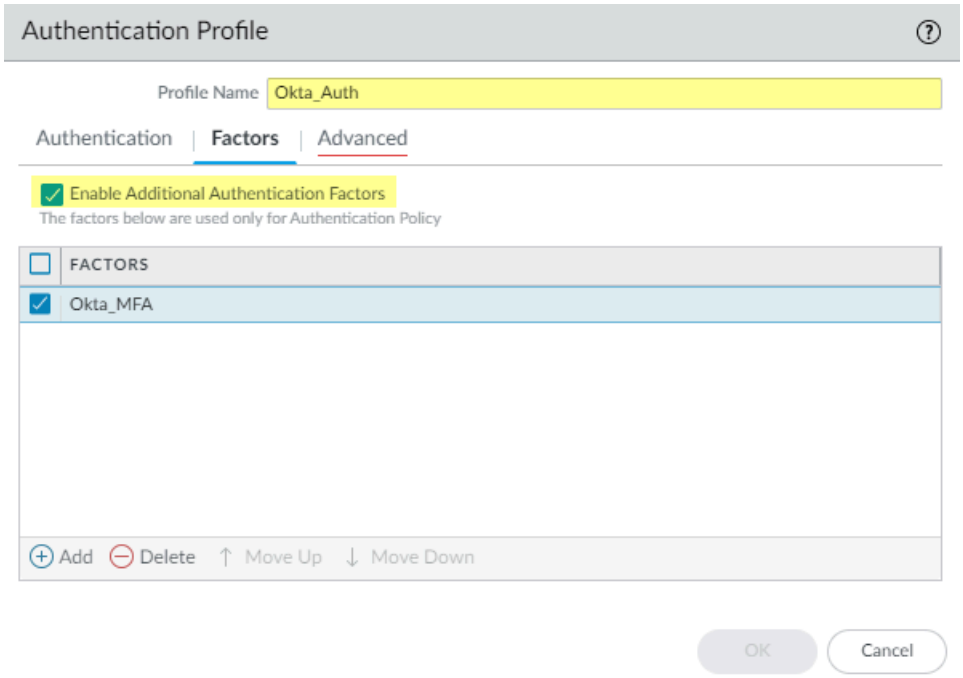
STEP 3 | Configure Authentication Portal using **Redirect Mode** to redirect users to the MFA vendor's challenge.

https://t.me/learningnets

STEP 4 | Enable response pages on the Interface Management Profile to redirect users to the response page challenge.



STEP 5 | Create an Authentication Profile and add the MFA vendor as a **Factor** (see Configure Multi-Factor Authentication, Step 3.)



https://t.me/learningnets

STEP 6 | Enable User-ID on the source zone to require identified users to respond to the challenge using your MFA vendor.

STEP 7 | Create an Authentication Enforcement Object to use the MFA vendor and create an Authentication policy rule (see Configure Authentication Policy, Steps 4 and 5).

STEP 8 | **Commit** your changes.

Verify MFA with Okta

STEP 1 | Verify your users received their enrollment emails, have activated their accounts, and have downloaded the Okta Verify app on their devices.

STEP 2 | Go to a website that will prompt the response page challenge.

> *If you are using a self-signed certificate instead of a PKI-assigned certificate from your organization, a security warning displays that users must click through to access the challenge.*

STEP 3 | Log in to the response page using your Okta credentials.

STEP 4 | Confirm the device receives the challenge push notification.

STEP 5 | Confirm users can successfully access the page after authenticating the challenge by accepting the push notification on their devices.

Configure MFA Between Duo and the Firewall

Multi-factor authentication (MFA) allows you to protect company assets by using multiple factors to verify the identity of users before allowing them to access network resources. There are multiple ways to use the Duo identity management service to authenticate with the firewall:

- Two-factor authentication for VPN logins using the GlobalProtect Gateway and a RADIUS server profile (supported on PAN-OS 7.0 and later).
- API-based integration using Authentication Portal and an MFA server profile (does not require a Duo Authentication Proxy or SAML IdP - supported on PAN-OS 8.0 and later).
- SAML integration for on-premise servers (supported on PAN-OS 8.0 and later).

To enable SAML MFA between the firewall and Duo to secure administrative access to the firewall:

- Configure Duo for SAML MFA with Duo Access Gateway
- Configure the Firewall to Integrate with Duo
- Verify MFA with Duo

Configure Duo for SAML MFA with Duo Access Gateway

Before you begin, verify that you have deployed the DuoAccessGateway (DAG) on an on-premise server in your DMZ zone.

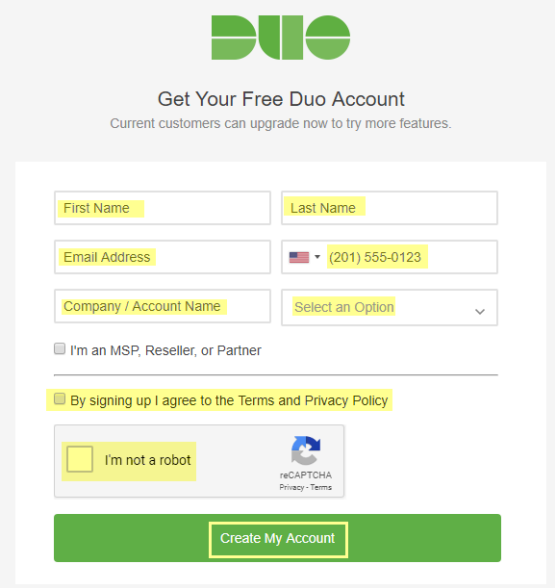
Create your Duo administrator account and configure the Duo Access Gateway to authenticate your users before they can access resources.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

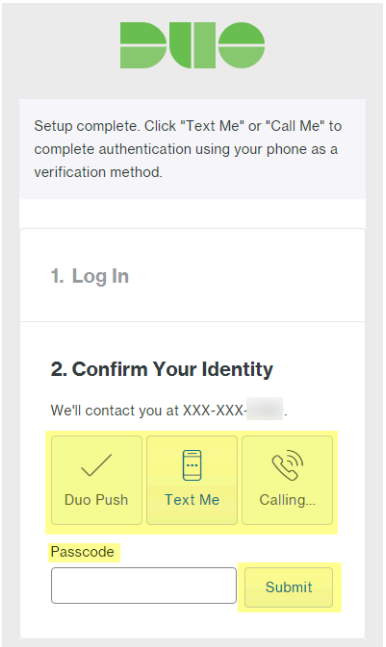
STEP 1 | Create your Duo administrator account.

1. On the Duo account creation page, enter your **First Name**, **Last Name**, **Email Address**, **Cell Phone Number**, **Company / Account Name**, and select the number of employees in the organization.
2. Agree to the Terms and Privacy Policy and respond to the reCAPTCHA challenge to **Create My Account**.



STEP 2 | Verify your Duo administrator account.

1. Select the authentication verification method (**Duo Push**, **Text Me**, or **Calling...**).
2. Enter the **Passcode** you receive and **Submit** it to verify your account.



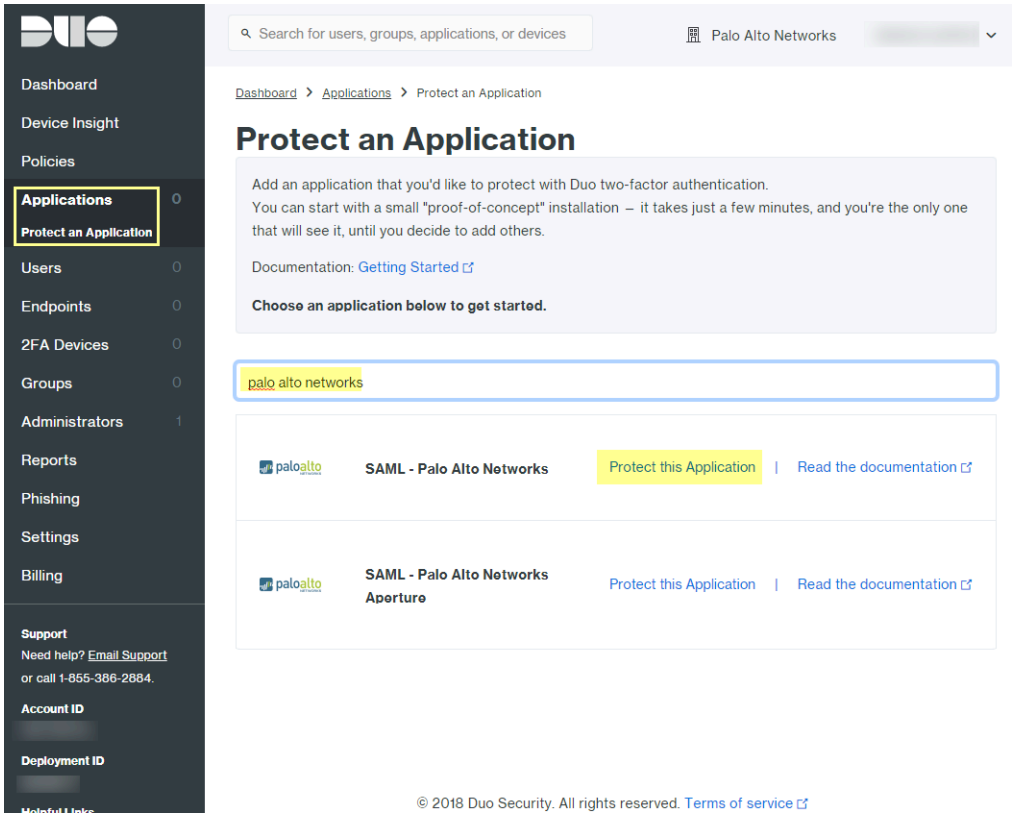
 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 3 | Configure your Duo service for SAML.

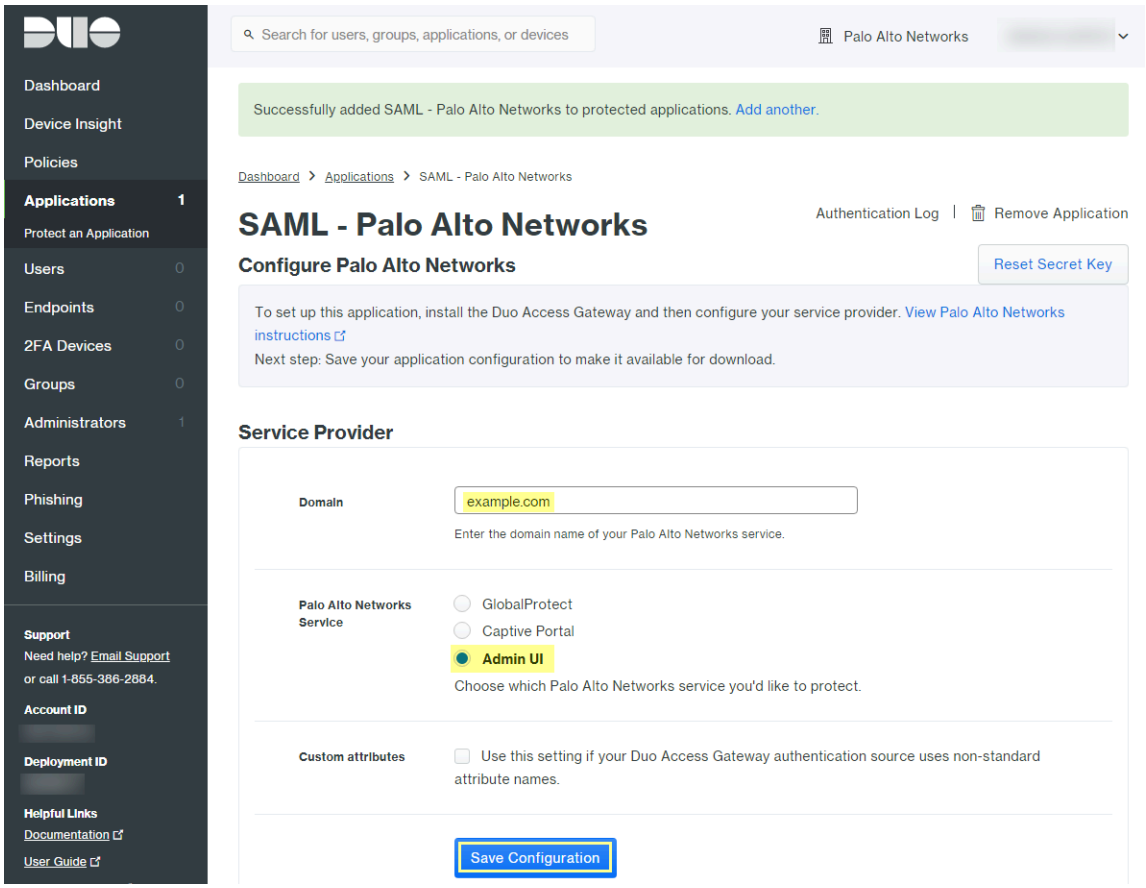
After creating your configuration, download the configuration file at the top of the page.

1. In the Duo Admin Panel, select **Applications** > **Protect an Application**.
2. Enter **`Palo Alto Networks`** to search the applications.
3. Locate **SAML - Palo Alto Networks** in the list of results, then **Protect this Application**.



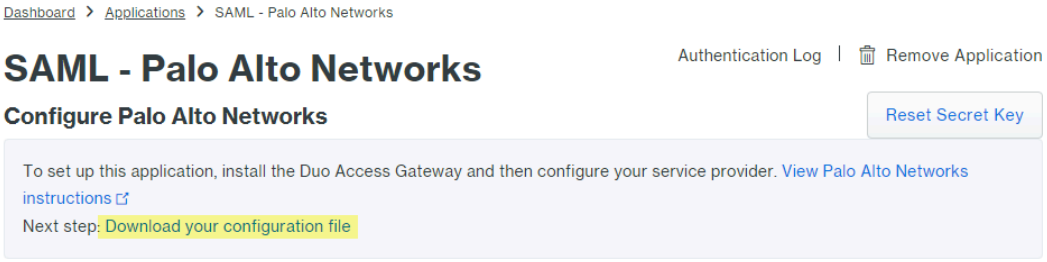
4. Enter the **Domain**.
5. Select **Admin UI** as the **Palo Alto Networks Service**.
6. Configure your **Policy** and other **Settings**, and **Save Configuration**.

https://t.me/learningnets



7. **Download your configuration file**.

   The link to download the file is at the top of the page.



STEP 4 | Upload the configuration file to the Duo Access Gateway (DAG).

1. In the DAG admin console, select **Applications**.
2. Click **Choose File** and select the configuration file you downloaded, then **Upload** it.
3. In **Settings** > **Session Management**, disable **User agent binding**, then **Save Settings**.

https://t.me/learningnets

STEP 5 | In the DAG admin console, configure your Active Directory or OpenLDAP server as the authentication source and download the metadata file.

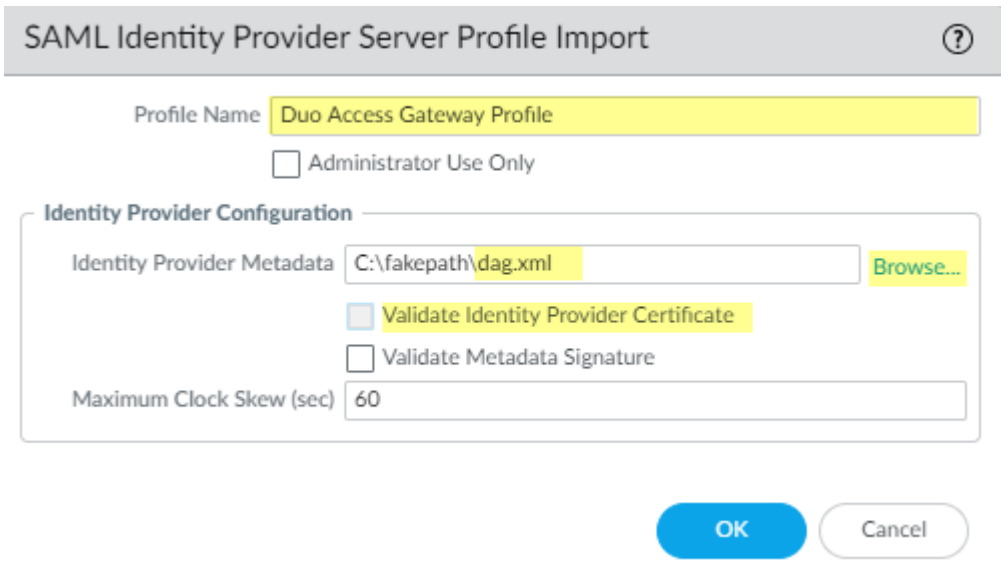
1. Log in to the DAG admin console.
2. In **Authentication Source** > **Set Active Source**, select your **Source type** (Active Directory or OpenLDAP) and **Set Active Source**.
3. In **Configure Sources**, enter the **Attributes**.
   - For Active Directory, enter **`mail,sAMAccountName,userPrincipalName,objectGUID`**.
   - For OpenLDAP, enter **`mail,uid`**.
   - For any custom attributes, append them to the end of the list and separate each attribute with a comma. Do not delete any existing attributes.
4. **Save Settings** to save the configuration.
5. Select **Applications** > **Metadata**, then click **Download XML metadata** to download the XML metadata you will need to import into the firewall.

   The file will be named dag.xml. Because this file includes sensitive information to authenticate your Duo account with the firewall, make sure to keep the file in a secure location to avoid the risk of compromising this information.

Configure the Firewall to Integrate with Duo

STEP 1 | Import the Duo metadata.

1. Log on to the firewall web interface.
2. On the firewall, select **Device** > **Server Profiles** > **SAML Identity Provider** > **Import**.
3. Enter the **Profile Name**.
4. **Browse** to the **Identity Provider Metadata** file (**dag.xml**).
5. If the Duo Access Gateway provides a self-signed certificate as the signing certificate for the IdP, you cannot **Validate Identity Provider Certificate**. In this case, ensure that you are using PAN-OS 10.1 to mitigate exposure to CVE-2020-2021.



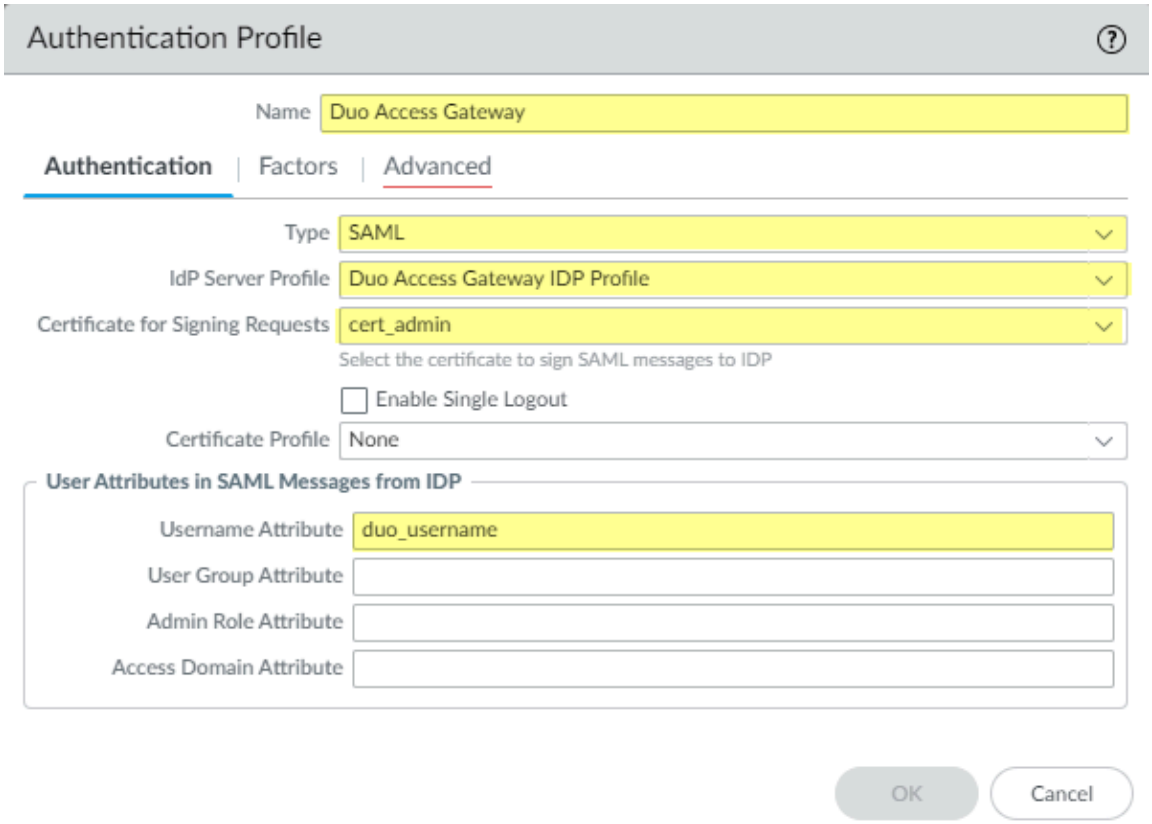
 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 2 | Add an authentication profile.

The authentication profile allows Duo as the identity provider that validates administrator login credentials.

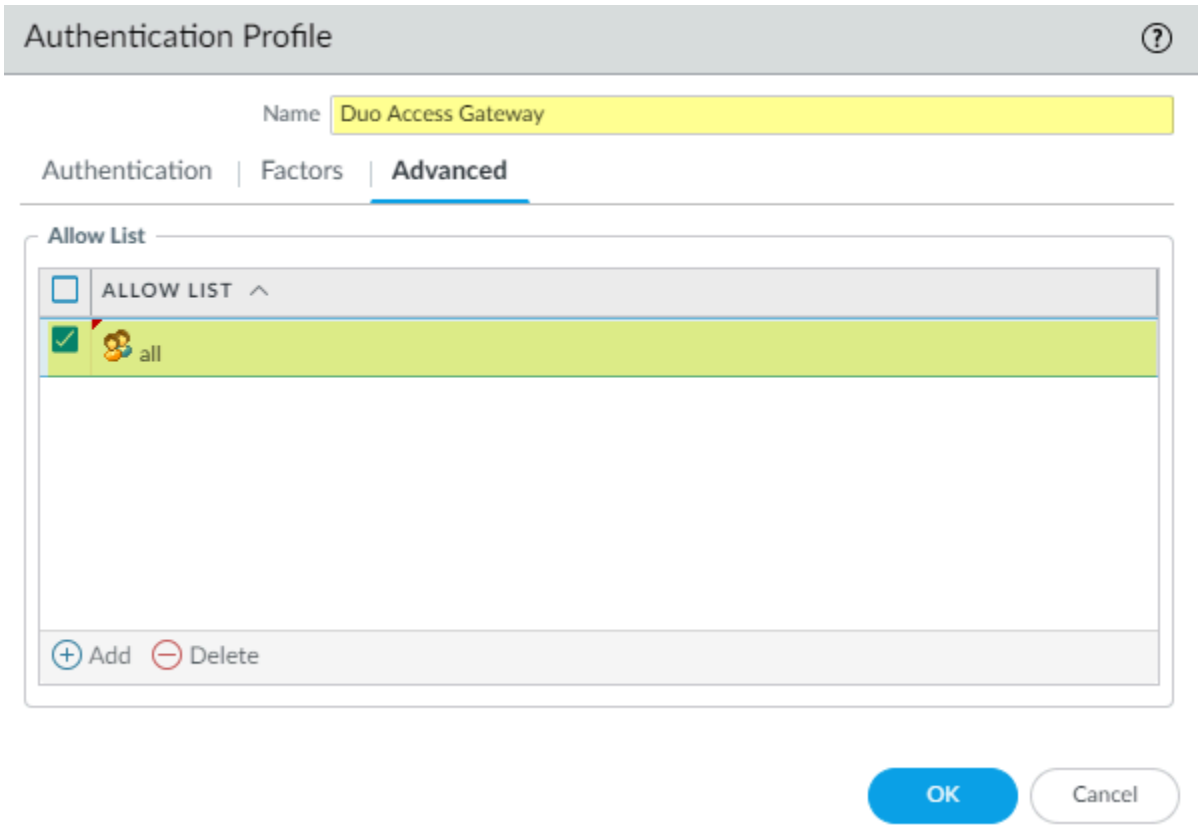
1. **Add** an **Authentication Profile**.
2. Enter the profile **Name**.
3. Select **SAML** as the authentication **Type**.
4. Select **Duo Access Gateway Profile** as the **IdP Server Profile**.
5. Select the certificate you want to use for SAML communication with the Duo Access Gateway for the **Certificate for Signing Requests**.
6. Enter **`duo_username`** as the **Username Attribute**.



7. Select **Advanced** to **Add** an allow list.
8. Select **all**, then click **OK**.
9. **Commit** the changes.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

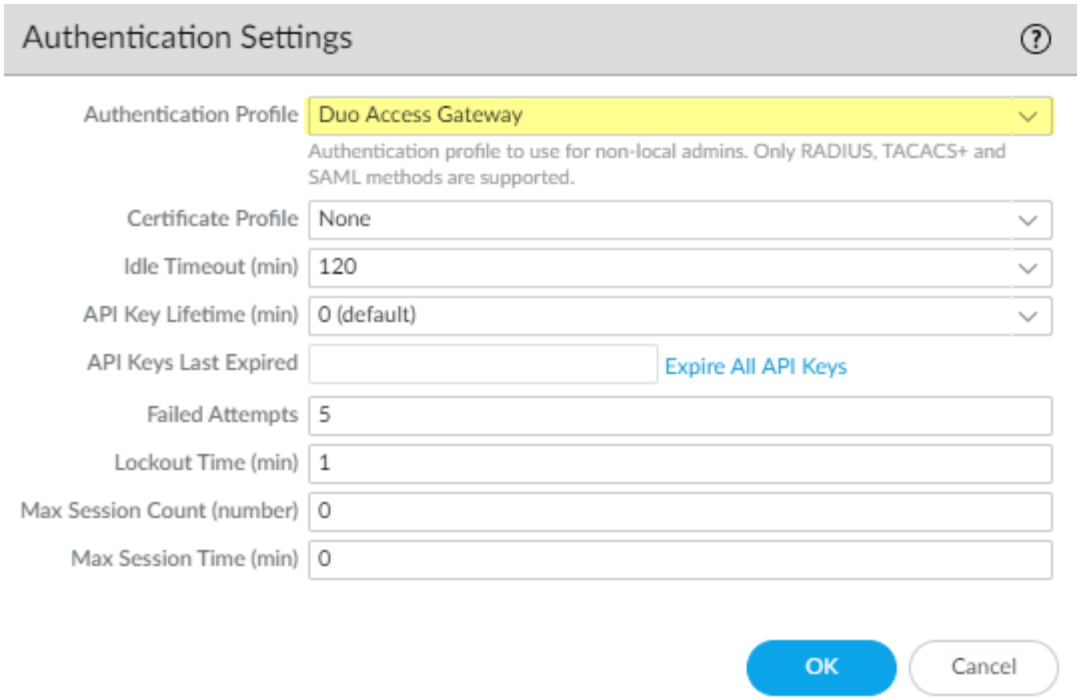


 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 3 | Specify the authentication settings that the firewall uses for SAML authentication with Duo.

1. Select **Device** > **Setup** > **Management** and edit the **Authentication Settings**.
2. Select **Duo Access Gateway** as the **Authentication Profile**, then click **OK**.



3. **Commit** your changes.

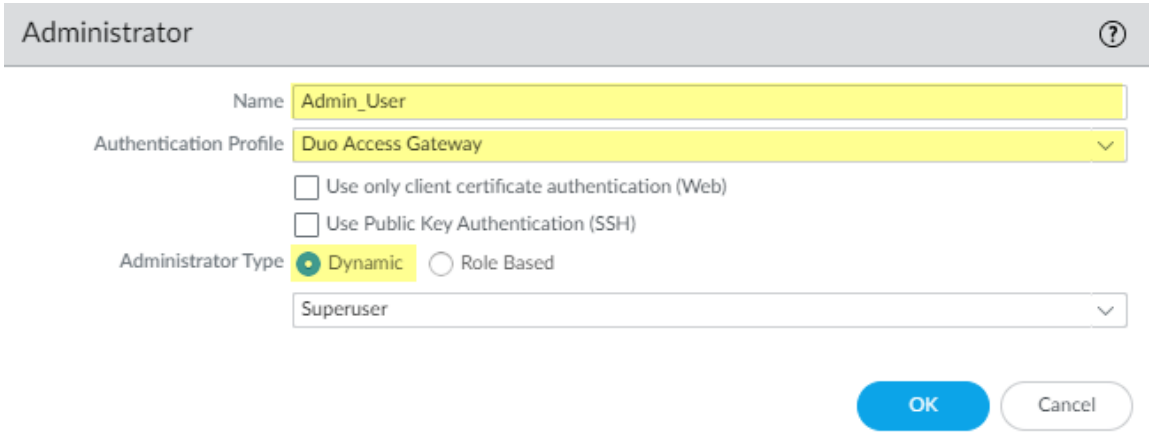
 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 4 | Add accounts for administrators who will authenticate to the firewall using Duo.

1. Select **Device** > **Administrators** and **Add** an account.
2. Enter a user **Name**.
3. Select **Duo Access Gateway** as the **Authentication Profile**.
4. Select the **Administrator Type**, then click **OK**.

   Select **Role Based** if you want to use a custom role for the user. Otherwise, select **Dynamic**. To require administrators to log in using SSO with Duo, assign the authentication profile to all current administrators.



Verify MFA with Duo

STEP 1 | Log in to the web interface on the firewall.

STEP 2 | Select **Use Single Sign-On** and **Continue**.

STEP 3 | Enter your login credentials on the Duo Access Gateway login page.

STEP 4 | Select an authentication method (push notification, phone call, or passcode entry).

When you authenticate successfully, you will be redirected to the firewall web interface.

https://t.me/learningnets

Configure SAML Authentication

To configure SAML single sign-on (SSO) and single logout (SLO), you must register the firewall and the IdP with each other to enable communication between them. If the IdP provides a metadata file containing registration information, you can import it onto the firewall to register the IdP and to create an IdP server profile. The server profile defines how to connect to the IdP and specifies the certificate that the IdP uses to sign SAML messages. You can also use a certificate for the firewall to sign SAML messages. Using certificates is a requirement to secure communications between the firewall and the IdP.

Palo Alto Networks requires HTTPS to ensure the confidentiality of all SAML transactions instead of alternative approaches such as encrypted SAML assertions. To ensure the integrity of all messages processed in a SAML transaction, Palo Alto Networks requires digital certificates to cryptographically sign all messages.

The following procedure describes how to configure SAML authentication for end users and firewall administrators. You can also configure SAML authentication for Panorama administrators.

SSO is available to administrators and to GlobalProtect and Authentication Portal end users. SLO is available to administrators and GlobalProtect end users, but not to Authentication Portal end users.

Administrators can use SAML to authenticate to the firewall web interface, but not to the CLI.

STEP 1 | Obtain the certificates that the IdP and firewall will use to sign SAML messages.

If the certificates don't specify key usage attributes, all usages are allowed by default, including signing messages. In this case, you can Obtain Certificates by any method.

If the certificates do specify key usage attributes, one of the attributes must be Digital Signature, which is not available on certificates that you generate on the firewall or Panorama. In this case, you must import the certificates:

- **Certificate the firewall uses to sign SAML messages**—Import the certificate from your enterprise certificate authority (CA) or a third-party CA.
- **Certificate the IdP uses to sign SAML messages** (Required for all deployments)—Import a metadata file containing the certificate from the IdP (see the next step). The IdP certificate is limited to the following algorithms:

  Public key algorithms—RSA (1,024 bits or larger) and ECDSA (all sizes). A firewall in FIPS/CC mode supports RSA (2,048 bits or larger) and ECDSA (all sizes).

  Signature algorithms—SHA1, SHA256, SHA384, and SHA512. A firewall in FIPS/CC mode supports SHA256, SHA384, and SHA512.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 2 | Add a SAML IdP server profile.

The server profile registers the IdP with the firewall and defines how they connect.

In this example, you import a SAML metadata file from the IdP so that the firewall can automatically create a server profile and populate the connection, registration, and IdP certificate information.

*If the IdP doesn't provide a metadata file, select **Device** > **Server Profiles** > **SAML Identity Provider**, **Add** the server profile, and manually enter the information (consult your IdP administrator for the values).*

1. Export the SAML metadata file from the IdP to a client system from which you can upload the metadata to the firewall.

   The certificate specified in the file must meet the requirements listed in the preceding step. Refer to your IdP documentation for instructions on exporting the file.
2. Select **Device** > **Server Profiles** > **SAML Identity Provider** or **Panorama** > **Server Profiles** > **SAML Identity Provider** on Panorama™ and **Import** the metadata file onto the firewall.
3. Enter a **Profile Name** to identify the server profile.
4. **Browse** to the **Identity Provider Metadata** file.
5. Select **Validate Identity Provider Certificate** (default) to validate the chain of trust and optionally the revocation status of the IdP certificate.

   To enable this option, a Certificate Authority (CA) must issue your IdP's signing certificate. You must create a Certificate Profile that has the CA that issued the IdP's signing certificate. In the Authentication Profile, select the SAML Server profile and Certificate Profile to validate the IdP certificate.

   If your IdP signing certificate is a self-signed certificate, there is no chain of trust; as a result, you cannot enable this option. The firewall always validates the signature of the SAML Responses or Assertions against the Identity Provider certificate that you configure whether or not you enable the **Validate Identity Provider Certificate** option. If your IdP provides a self-signed certificate, ensure that you are using PAN-OS 10.1 to mitigate exposure to CVE-2020-2021.

   Validate the certificate to ensure it hasn't been compromised and to improve security.
6. Enter the **Maximum Clock Skew**, which is the allowed difference in seconds between the system times of the IdP and the firewall at the moment when the firewall validates IdP messages (default is 60; range is 1 to 900). If the difference exceeds this value, authentication fails.
7. Click **OK** to save the server profile.
8. Click the server profile Name to display the profile settings. Verify that the imported information is correct and edit it if necessary.
9. Whether you import the IdP metadata or manually enter the IdP information, always ensure that the signing certificate of your SAML identity provider is the **Identity Provider Certificate** for your server profile and your IdP sends signed SAML Responses, Assertions, or both.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 3 | Configure an authentication profile.

The profile defines authentication settings that are common to a set of users.

1. Select **Device** > **Authentication Profile** and **Add** a profile.
2. Enter a **Name** to identify the profile.
3. Set the **Type** to **SAML**.
4. Select the **IdP Server Profile** you configured.
5. Select the **Certificate for Signing Requests**.

   The firewall uses this certificate to sign messages it sends to the IdP. You can import a certificate generated by your enterprise CA or you can generate a certificate using the root CA that was generated on the firewall or Panorama.
6. (Optional) **Enable Single Logout** (disabled by default).
7. Select the **Certificate Profile** that the firewall will use to validate the **Identity Provider Certificate**.
8. Enter the **Username Attribute** that IdP messages use to identify users (default **`username`**).

   *When you predefine dynamic administrator roles for users, use lower-case to specify the role (for example, enter **`superreader`**, not **`SuperReader`**). If you manage administrator authorization in the IdP identity store, specify the **Admin Role Attribute** and **Access Domain Attribute** also.*
9. Select **Advanced** and **Add** the users and user groups that are allowed to authenticate with this authentication profile.
10. Click **OK** to save the authentication profile.

STEP 4 | Assign the authentication profile to firewall applications that require authentication.

1. Assign the authentication profile to:
   - Administrator accounts that you manage locally on the firewall. In this example, Configure a Firewall Administrator Account before you verify the SAML configuration later in this procedure.
   - Administrator accounts that you manage externally in the IdP identity store. Select **Device** > **Setup** > **Management**, edit the Authentication Settings, and select the **Authentication Profile** you configured.
   - Authentication policy rules that secure the services and applications that end users access through Authentication Portal. See Configure Authentication Policy.
   - GlobalProtect portals and gateways that end users access.
2. **Commit** your changes.

   The firewall validates the **Identity Provider Certificate** that you assigned to the SAML IdP server profile.

https://t.me/learningnets

STEP 5 | Create a SAML metadata file to register the firewall application (management access, Authentication Portal, or GlobalProtect) on the IdP.

1. Select **Device** > **Authentication Profile** and, in the Authentication column for the authentication profile you configured, click **Metadata**.
2. In the **Service** drop-down, select the application you want to register:
   - **management** (default)—Administrative access to the web interface.
   - **authentication-portal**—End user access to services and applications through Authentication Portal.
   - **global-protect**—End user access to services and applications through GlobalProtect.
3. (Authentication Portal or GlobalProtect only) for the **Vsysname Combo**, select the virtual system in which the Authentication Portal settings or GlobalProtect portal are defined.
4. Enter the interface, IP address, or hostname based on the application you will register:
   - **management**—For the **Management Choice**, select **Interface** (default) and select an interface that is enabled for management access to the web interface. The default selection is the IP address of the MGT interface.
   - **authentication-portal**—For the **IP Hostname**, enter the IP address or hostname of the **Redirect Host** (see **Device** > **User Identification** > **Authentication Portal Settings**).
   - **global-protect**—For the **IP Hostname**, enter the hostname or IP address of the GlobalProtect portal or gateway.
5. Click **OK** and save the metadata file to your client system.
6. Import the metadata file into the IdP server to register the firewall application. Refer to your IdP documentation for instructions.

 ©2022 Palo Alto Networks, Inc.

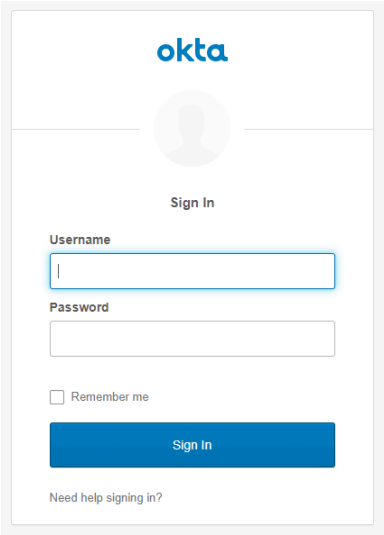
https://t.me/learningnets

STEP 6 | Verify that users can authenticate using SAML SSO.

For example, to verify that SAML is working for access to the web interface using a local administrator account:

1. Go to the URL of the firewall web interface.
2. Click **Use Single Sign-On**.
3. Enter the username of the administrator.
4. Click **Continue**.

   The firewall redirects you to authenticate to the IdP, which displays a login page. For example:



5. Log in using your SSO username and password.

   After you successfully authenticate on the IdP, it redirects you back to the firewall, which displays the web interface.
6. Use your firewall administrator account to request access to another SSO application.

   Successful access indicates SAML SSO authentication succeeded.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Configure Kerberos Single Sign-On

Palo Alto Networks firewalls and Panorama support Kerberos V5 single sign-on (SSO) to authenticate administrators to the web interface and end users to Authentication Portal. With Kerberos SSO enabled, the user needs to log in only for initial access to your network (such as logging in to Microsoft Windows). After this initial login, the user can access any browser-based service in the network (such as the firewall web interface) without having to log in again until the SSO session expires.

STEP 1 | Create a Kerberos keytab.

The keytab is a file that contains the principal name and password of the firewall, and is required for the SSO process. When you configure Kerberos in your Authentication Profile and Sequence, the firewall first checks for a Kerberos SSO hostname. If you provide a hostname, the firewall searches the keytabs for a service principal name that matches the hostname and uses only that keytab for decryption. If you do not provide a hostname, the firewall tries

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

each keytab in the authentication sequence until it is able to successfully authenticate using Kerberos.

If the Kerberos SSO hostname is included in the request sent to the firewall, then the hostname must match the service principal name of the keytab; otherwise, the Kerberos authentication request is not sent.

1. Log in to the Active Directory server and open a command prompt.
2. Enter the following command to register the service principal name (SPN) for GlobalProtect or Authentication Portal, where <*portal_fqdn*> and <*service_account_username*> are variables.

   setspn -s HTTP/<*portal_fqdn*> <*service_account_username*>
3. Create Kerberos account for the firewall. Refer to your Kerberos documentation for the steps.
4. Log in to the KDC and open a command prompt.
5. Enter the following command, where <*portal_fqdn*>, <*kerberos_realm*>, <*netbios_name*>, <*service_account_username*>, <*password*>, <*filename*>, and <*algorithm*> are variables.

   ktpass /princ HTTP <*portal_fqdn*>@<*kerberos_realm*> /mapuser <*netbios_name*>\<*service_account_username*> /pass <*password*>/out <*filename*>.keytab /ptype KRB5_NT_PRINCIPAL /crypto <*algorithm*>

   *The <kerberos_realm> value must be in all uppercase characters (for example, enter **AD1.EXAMPLE.COM**, not **ad1.example.com**).*

   *If the firewall is in FIPS/CC mode, the algorithm must be **aes128-cts-hmac-sha1-96** or **aes256-cts-hmac-sha1-96**. Otherwise, you can also use **des3-cbc-sha1 or arcfour-hmac**. To use an Advanced Encryption Standard (AES) algorithm, the functional level of the KDC must be Windows Server 2012 or later and you must enable AES encryption for the firewall account.*

   The algorithm in the keytab must match the algorithm in the service ticket that the TGS issues to clients. Your Kerberos administrator determines which algorithms the service tickets use.

STEP 2 | Configure an Authentication Profile and Sequence to define Kerberos settings and other authentication options that are common to a set of users.

- Enter the **Kerberos Realm** (usually the DNS domain of the users, except that the realm is uppercase).
- **Import** the **Kerberos Keytab** that you created for the firewall.

STEP 3 | Assign the authentication profile to the firewall application that requires authentication.

- Administrative access to the web interface—Configure a Firewall Administrator Account and assign the authentication profile you configured.
- End user access to services and applications—Assign the authentication profile you configured to an authentication enforcement object. When configuring the object, set the **Authentication Method** to **browser-challenge**. Assign the object to Authentication

https://t.me/learningnets

policy rules. For the full procedure to configure authentication for end users, see Configure Authentication Policy.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Configure Kerberos Server Authentication

You can use Kerberos to natively authenticate end users and firewall or Panorama administrators to an Active Directory domain controller or a Kerberos V5-compliant authentication server. This authentication method is interactive, requiring users to enter usernames and passwords.

To use a Kerberos server for authentication, the server must be accessible over an IPv4 address. IPv6 addresses are not supported.

STEP 1 | Add a Kerberos server profile.

The profile defines how the firewall connects to the Kerberos server.

1. Select **Device** > **Server Profiles** > **Kerberos** or **Panorama** > **Server Profiles** > **Kerberos** on Panorama™ and **Add** a server profile.
2. Enter a **Profile Name** to identify the server profile.
3. **Add** each server and specify a **Name** (to identify the server), IPv4 address or FQDN of the **Kerberos Server**, and optional **Port** number for communication with the server (default 88).

   If you use an FQDN address object to identify the server and you subsequently change the address, you must commit the change in order for the new server address to take effect.
4. Click **OK** to save your changes to the profile.

STEP 2 | Assign the server profile to an Configure an Authentication Profile and Sequence.

The authentication profile defines authentication settings that are common to a set of users.

STEP 3 | Assign the authentication profile to the firewall application that requires authentication.

- Administrative access to the web interface—Configure a Firewall Administrator Account and assign the authentication profile you configured.
- End user access to services and applications—Assign the authentication profile you configured to an authentication enforcement object and assign the object to Authentication policy rules. For the full procedure to configure authentication for end users, see Configure Authentication Policy.

STEP 4 | Verify that the firewall can Test Authentication Server Connectivity to authenticate users.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Configure TACACS+ Authentication

You can configure TACACS+ authentication for end users and firewall or Panorama administrators. You can also use a TACACS+ server to manage administrator authorization (role and access domain assignments) by defining Vendor-Specific Attributes (VSAs). For all users, you must configure a TACACS+ server profile that defines how the firewall or Panorama connects to the server. You then assign the server profile to an authentication profile for each set of users who require common authentication settings. What you do with the authentication profile depends on which users the TACACS+ server authenticates:

- **End users**—Assign the authentication profile to an authentication enforcement object and assign the object to Authentication policy rules. For the full procedure, see Configure Authentication Policy.
- **Administrative accounts with authorization managed locally on the firewall or Panorama**—Assign the authentication profile to firewall administrator or Panorama administrator accounts.
- **Administrative accounts with authorization managed on the TACACS+ server**—The following procedure describes how to configure TACACS+ authentication and authorization for firewall administrators. For Panorama administrators, refer to Configure TACACS+ Authentication for PanoramaAdministrators.

STEP 1 | Add a TACACS+ server profile.

The profile defines how the firewall connects to the TACACS+ server.

1. Select **Device** > **Server Profiles** > **TACACS+** or **Panorama** > **Server Profiles** > **TACAS+** on Panorama™ and **Add** a profile.
2. Enter a **Profile Name** to identify the server profile.
3. (Optional) Select **Administrator Use Only** to restrict access to administrators.
4. Enter a **Timeout** interval in seconds after which an authentication request times out (default is 3; range is 1–20).
5. Select the **Authentication Protocol** (default is **CHAP**) that the firewall uses to authenticate to the TACACS+ server.

   *Select **CHAP** if the TACACS+ server supports that protocol; it is more secure than **PAP**.*
6. **Add** each TACACS+ server and enter the following:
   - **Name** to identify the server
   - **TACACS+ Server** IP address or FQDN. If you use an FQDN address object to identify the server and you subsequently change the address, you must commit the change for the new server address to take effect.
   - **Secret**/**Confirm Secret** (a key to encrypt usernames and passwords)
   - Server **Port** for authentication requests (default is 49)
7. Click **OK** to save the server profile.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 2 | Assign the TACACS+ server profile to an authentication profile.

The authentication profile defines authentication settings that are common to a set of users.

1. Select **Device** > **Authentication Profile** and **Add** a profile.
2. Enter a **Name** to identify the profile.
3. Set the **Type** to **TACACS+**.
4. Select the **Server Profile** you configured.
5. Select **Retrieve user group from TACACS+** to collect user group information from VSAs defined on the TACACS+ server.

   The firewall matches the group information against the groups you specify in the Allow List of the authentication profile.
6. Select **Advanced** and, in the Allow List, **Add** the users and groups that are allowed to authenticate with this authentication profile.
7. Click **OK** to save the authentication profile.

STEP 3 | Configure the firewall to use the authentication profile for all administrators.

1. Select **Device** > **Setup** > **Management** and edit the Authentication Settings.
2. Select the **Authentication Profile** you configured and click **OK**.

STEP 4 | Configure the roles and access domains that define authorization settings for administrators.

If you already defined TACACS+ VSAs on the TACACS+ server, the names you specify for roles and access domains on the firewall must match the VSA values.

1. Configure an Admin Role Profile if the administrator will use a custom role instead of a predefined (dynamic) role.
2. Configure an access domain if the firewall has more than one virtual system—Select **Device** > **Access Domain**, **Add** an access domain, enter a **Name** to identify the access domain, and **Add** each virtual system that the administrator will access, and then click **OK**.

STEP 5 | **Commit** your changes to activate them on the firewall.

STEP 6 | Configure the TACACS+ server to authenticate and authorize administrators.

Refer to your TACACS+ server documentation for the specific instructions to perform these steps:

1. Add the firewall IP address or hostname as the TACACS+ client.
2. Add the administrator accounts.

   *If you selected **CHAP** as the **Authentication Protocol**, you must define accounts with reversibly encrypted passwords. Otherwise, CHAP authentication will fail.*
3. Define TACACS+ VSAs for the role, access domain, and user group of each administrator.

   *When you predefine dynamic administrator roles for users, use lower-case to specify the role (for example, enter **`superuser`**, not **`SuperUser`**).*

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 7 | Verify that the TACACS+ server performs authentication and authorization for administrators.

1. Log in the firewall web interface using an administrator account that you added to the TACACS+ server.
2. Verify that you can access only the web interface pages that are allowed for the role you associated with the administrator.
3. In the **Monitor**, **Policies**, and **Objects** tabs, verify that you can access only the virtual systems that are allowed for the access domain you associated with the administrator.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Configure RADIUS Authentication

You can configure RADIUS authentication for end users and firewall or Panorama administrators. For administrators, you can use RADIUS to manage authorization (role and access domain assignments) by defining Vendor-Specific Attributes (VSAs). You can also use RADIUS to implement Multi-Factor Authentication (MFA) for administrators and end users. To enable RADIUS authentication, you must configure a RADIUS server profile that defines how the firewall or Panorama connects to the server (see Step 1 below). You then assign the server profile to an authentication profile for each set of users who require common authentication settings (see Step 5 below). What you do with the authentication profile depends on which users the RADIUS server authenticates:

- **End users**—Assign the authentication profile to an authentication enforcement object and assign the object to Authentication policy rules. For the full procedure, see Configure Authentication Policy.



You can also configure client systems to send RADIUS Vendor-Specific Attributes (VSAs) to the RADIUS server by assigning the authentication profile to a GlobalProtect portal or gateway. RADIUS administrators can then perform administrative tasks based on those VSAs.

- **Administrative accounts with authorization managed locally on the firewall or Panorama**—Assign the authentication profile to firewall administrator or Panorama administrator accounts.
- **Administrative accounts with authorization managed on the RADIUS server**—The following procedure describes how to configure RADIUS authentication and authorization for firewall administrators. For Panorama administrators, refer to Configure RADIUS Authentication for PanoramaAdministrators.

https://t.me/learningnets

STEP 1 | Add a RADIUS server profile.

The profile defines how the firewall connects to the RADIUS server.

1. Select **Device** > **Server Profiles** > **RADIUS** or **Panorama** > **Server Profiles** > **RADIUS** on Panorama™and **Add** a profile.
2. Enter a **Profile Name** to identify the server profile.
3. (Optional) Select **Administrator Use Only** to restrict access to administrators.
4. Enter a **Timeout** interval in seconds after which an authentication request times out (default is 3; range is 1–120).

   If you use the server profile to integrate the firewall with an MFA service, enter an interval that gives users enough time to authenticate. For example, if the MFA service prompts for a one-time password (OTP), users need time to see the OTP on their endpoint device and then enter the OTP in the MFA login page.
5. Enter the number of **Retries**.
6. Select the **Authentication Protocol** (default is **PEAP-MSCHAPv2**) that the firewall uses to authenticate to the RADIUS server.

   Depending on which factors you want to use to authenticate users within your multi-factor authentication (MFA) environment, select the appropriate authentication protocol:

   - **Username, password, and push (an automatically triggered out-of-band request)**: Supported with all authentication protocols
   - **Push, password, token, and PIN (when password or token or PIN are provided together)**: Supported with PAP, PEAP with GTC, and EAP-TTLS with PAP
   - **Username, password, token, and PIN, and challenge-response (when password or token or PIN are provided together)**: Supported with PAP and PEAP with GTC

   If you select an EAP authentication method (PEAP-MSCHAPv2, PEAP with GTC, or EAP-TTLS with PAP), confirm that your RADIUS server supports Transport Layer Security (TLS) 1.1 or higher and that the root and intermediate certificate authorities (CAs) for your RADIUS server are included in the certificate profile associated with the RADIUS server profile. If you select an EAP method and you do not associate a correctly configured certificate profile with the RADIUS profile, authentication fails.
7. **Add** each RADIUS server and enter the following:
   - **Name** to identify the server
   - **RADIUS Server** IP address or FQDN. If you use an FQDN to identify the server and you subsequently change the address, you must commit the change for the new server address to take effect.
   - **Secret**/**Confirm Secret** is a key to encrypt passwords and can be up to 64 characters in length.
   - Server **Port** for authentication requests (default is 1812)
8. Click **OK** to save the server profile.

For redundancy, add multiple RADIUS servers in the sequence you want the firewall to use. If you have selected an EAP method, configure an authentication sequence to ensure that users will be able to successfully respond to the authentication challenge. There is no alternate authentication method with EAP: if the user fails the authentication challenge and you have

https://t.me/learningnets

not configured an authentication sequence that allows another authentication method, authentication fails.

STEP 2 | If you are using PEAP-MSCHAPv2 with GlobalProtect, select **Allow users to change passwords after expiry** to allow GlobalProtect users to changed expired passwords to log in.

STEP 3 | (PEAP-MSCHAPv2, PEAP with GTC, or EAP-TTLS with PAP only) To anonymize the user's identity in the outer tunnel that is created after authenticating with the server, select **Make Outer Identity Anonymous**.

You must configure the RADIUS server so that the entire chain allows access for anonymous users. Some RADIUS server configurations may not support anonymous outer IDs, and you may need to clear the option. When cleared, the RADIUS server transmits usernames in cleartext.

STEP 4 | If you select an EAP authentication method, select a Certificate Profile.

STEP 5 | Assign the RADIUS server profile to an authentication profile.

The authentication profile defines authentication settings that are common to a set of users.

1. Select **Device** > **Authentication Profile** and **Add** a profile.
2. Enter a **Name** to identify the authentication profile.
3. Set the **Type** to **RADIUS**.
4. Select the **Server Profile** you configured.
5. Select **Retrieve user group from RADIUS** to collect user group information from VSAs defined on the RADIUS server.

   The firewall matches the group information against the groups you specify in the Allow List of the authentication profile.
6. Select **Advanced** and, in the Allow List, **Add** the users and groups that are allowed to authenticate with this authentication profile.
7. Click **OK** to save the authentication profile.

STEP 6 | Configure the firewall to use the authentication profile for all administrators.

1. Select **Device** > **Setup** > **Management** and edit the Authentication Settings.
2. Select the **Authentication Profile** you configured and click **OK**.

STEP 7 | Configure the roles and access domains that define authorization settings for administrators.

If you already defined RADIUS VSAs on the RADIUS server, the names you specify for roles and access domains on the firewall must match the VSA values.

1. Configure an Admin Role Profile if the administrator uses a custom role instead of a predefined (dynamic) role.
2. Configure an access domain if the firewall has more than one virtual system:
   1. Select **Device** > **Access Domain**, **Add** an access domain, and enter a **Name** to identify the access domain.
   2. **Add** each virtual system that the administrator will access, and then click **OK**.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 8 | **Commit** your changes to activate them on the firewall.

STEP 9 | Configure the RADIUS server to authenticate and authorize administrators.

Refer to your RADIUS server documentation for the specific instructions to perform these steps:

1. Add the firewall IP address or hostname as the RADIUS client.
2. Add the administrator accounts.

   *If the RADIUS server profile specifies **CHAP** as the **Authentication Protocol**, you must define accounts with reversibly encrypted passwords. Otherwise, CHAP authentication will fail.*
3. Define the vendor code for the firewall (25461) and define the RADIUS VSAs for the role, access domain, and user group of each administrator.

   When you predefine dynamic administrator roles for users, use lower-case to specify the role (for example, enter **`superuser`**, not **`SuperUser`**).

   *When configuring the advanced vendor options on the ACS, you must set both the **Vendor Length Field Size** and **Vendor Type Field Size** to **1**. Otherwise, authentication will fail.*
4. If you have selected an EAP method, the firewall validates the server but not the client. To ensure client validity, restrict clients by IP address or subdomain.

STEP 10 | Verify that the RADIUS server performs authentication and authorization for administrators.

1. Log in the firewall web interface using an administrator account that you added to the RADIUS server.
2. Verify that you can access only the web interface pages that are allowed for the role you associated with the administrator.
3. In the **Monitor**, **Policies**, and **Objects** tabs, verify that you can access only the virtual systems that are allowed for the access domain you associated with the administrator.
4. In **Monitor** > **Authentication**, verify the **Authentication Protocol**.
5. Test the connection and the validity of the certificate profile using the following CLI command:

```
admin@PA-220 > test authentication authentication-profile
 auth-profile username <username> password <password>
```

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Configure LDAP Authentication

You can use LDAP to authenticate end users who access applications or services through Authentication Portal and authenticate firewall or Panorama administrators who access the web interface.

You can also connect to an LDAP server to define policy rules based on user groups. For details, see Map Users to Groups.

STEP 1 | Add an LDAP server profile.

The profile defines how the firewall connects to the LDAP server.

1. Select **Device** > **Server Profiles** > **LDAP** or **Panorama** > **Server Profiles** > **LDAP** on Panorama™ and **Add** a server profile.
2. Enter a **Profile Name** to identify the server profile.
3. (Multi-vsys only) Select the **Location** in which the profile is available.
4. (Optional) Select **Administrator Use Only** to restrict access to administrators.
5. **Add** the LDAP servers (up to four). For each server, enter a **Name** (to identify the server), **LDAP Server** IP address or FQDN, and server **Port** (default 389).

   If you use an FQDN address object to identify the server and you subsequently change the address, you must commit the change for the new server address to take effect.
6. Select the server **Type**.
7. Select the **Base DN**.

   To identify the Base DN of your directory, open the **Active Directory Domains and Trusts** Microsoft Management Console snap-in and use the name of the top-level domain.
8. Enter the **Bind DN** and **Password** to enable the authentication service to authenticate the firewall.

   The Bind DN account must have permission to read the LDAP directory.
9. Enter the **Bind Timeout** and **Search Timeout** in seconds (default is 30 for both).
10. Enter the **Retry Interval** in seconds (default is 60).
11. Enable the option to **Require SSL/TLS secured connection** (enabled by default). The protocol that the endpoint uses depends on the server port:
    - 389 (default)—TLS (Specifically, the device uses the StartTLS operation, which upgrades the initial plaintext connection to TLS.)
    - 636—SSL
    - Any other port—The device first attempts to use TLS. If the directory server doesn't support TLS, the device falls back to SSL.
12. (Optional) For additional security, enable to the option to **Verify Server Certificate for SSL sessions** so that the endpoint verifies the certificate that the directory server

https://t.me/learningnets

presents for SSL/TLS connections. To enable verification, you must also enable the option to **Require SSL/TLS secured connection**. For verification to succeed, the certificate must meet one of the following conditions:

- It is in the list of device certificates: **Device** > **Certificate Management** > **Certificates** > **Device Certificates.** If necessary, import the certificate into the device.
- The certificate signer is in the list of trusted certificate authorities: **Device** > **Certificate Management** > **Certificates** > **Default Trusted Certificate Authorities**.

13. Click **OK** to save the server profile.

STEP 2 | Assign the server profile to Configure an Authentication Profile and Sequence to define various authentication settings.

STEP 3 | Assign the authentication profile to the firewall application that requires authentication.

- **Administrative access to the web interface**—Configure a Firewall Administrator Account and assign the authentication profile you configured.
- **End user access to services and applications**—For the full procedure to configure authentication for end users, see Configure Authentication Policy.

STEP 4 | Verify that the firewall can Test Authentication Server Connectivity to authenticate users.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Connection Timeouts for Authentication Servers

You can configure the firewall to use External Authentication Services for authenticating administrators who access the firewall or Panorama and end users who access services or applications through Authentication Portal. To ensure that the firewall does not waste resources by continuously trying to reach an authentication server that is unreachable, you can set a timeout interval after which the firewall stops trying to connect. You set the timeout in the server profiles that define how the firewall connects to the authentication servers. When choosing timeout values, your goal is to strike a balance between the need to conserve firewall resources and to account for normal network delays that affect how quickly authentication servers respond to the firewall.

- Guidelines for Setting Authentication Server Timeouts
- Modify the PAN-OS Web Server Timeout
- Modify the Authentication Portal Session Timeout

Guidelines for Setting Authentication Server Timeouts

The following are some guidelines for setting the timeouts for firewall attempts to connect with External Authentication Services.

- In addition to the timeouts you set in server profiles for specific servers, the firewall has a global PAN-OS web server timeout. This global timeout applies when the firewall connects to any external server for authenticating administrative access to the firewall web interface or PAN-OS XML API and end user access to applications or services through Authentication Portal. The global timeout is 30 seconds by default (range is 3 to 125). It must be the same as or greater than the total time that any server profile allows for connection attempts. The total time in a server profile is the timeout value multiplied by the number of retries and the number of servers. For example, if a RADIUS server profile specifies a 3-second timeout, 3 retries, and 4 servers, the total time that the profile allows for connection attempts is 36 seconds (3 x 3 x 4). Modify the PAN-OS Web Server Timeout if necessary.

  Do not change the PAN-OS web server timeout unless you see authentication failures. Setting the timeout too high could degrade the performance of the firewall or cause it to drop authentication requests. You can review authentication failures in Authentication logs.

- The firewall applies an Authentication Portal session timeout that defines how long end users can take to respond to the authentication challenge in a Authentication Portal web form. The web form displays when users request services or applications that match an Authentication policy rule. The session timeout is 30 seconds by default (range is 1 to 1,599,999). It must be the same as or greater than the PAN-OS web server timeout. Modify the Authentication Portal Session Timeout if necessary. Keep in mind that increasing the PAN-OS web server and Authentication Portal session timeouts might degrade the performance of the firewall or cause it to drop authentication requests.

  The Authentication Portal session timeout is not related to the timers that determine how long the firewall retains IP address-to-username mappings.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

- Timeouts are cumulative for authentication sequences. For example, consider the case of an authentication sequence with two authentication profiles. One authentication profile specifies a RADIUS server profile with a 3-second timeout, 3 retries, and 4 servers. The other authentication profile specifies a TACACS+ server profile with a 3-second timeout and 2 servers. The longest possible period in which the firewall can try to authenticate user accounts with that authentication sequence is 42 seconds: 36 seconds for the RADIUS server profile plus 6 seconds for the TACACS+ server profile.
- The non-configurable timeout for Kerberos servers is 17 seconds for each server specified in the Kerberos server profile.
- To configure the timeouts and related settings for other server types, see:
  - Add an MFA server profile.
  - Add a SAML IdP server profile.
  - Add a TACACS+ server profile.
  - Add a RADIUS server profile.
  - Add an LDAP server profile.

Modify the PAN-OS Web Server Timeout

The PAN-OS web server timeout must be the same as or greater than the timeout in any authentication server profile multiplied by the number of retries and the number of servers in that profile.

Do not change the PAN-OS web server timeout unless you see authentication failures. Setting the timeout too high could degrade the performance of the firewall or cause it to drop authentication requests. You can review authentication failures in Authentication logs.

STEP 1 | Access the firewall CLI.

STEP 2 | Set the PAN-OS web server timeout by entering the following commands, where *<value>* is the number of seconds (default is 30; range is 3 to 125).

```
> configure
# set deviceconfig setting l3-service timeout <value>
# commit
```

Modify the Authentication Portal Session Timeout

The Authentication Portal session timeout must be the same as or greater than the PAN-OS web server timeout. For details, see Connection Timeouts for Authentication Servers.

The more you raise the PAN-OS web server and Authentication Portal session timeouts, the slower Authentication Portal will respond to users.

STEP 1 | Select **Device** > **Setup** > **Session** and edit the Session Timeouts.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 2 | Enter a new **Authentication Portal** value in seconds (default is 30; range is 1 to 1,599,999) and click **OK**.

STEP 3 | **Commit** your changes.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Configure Local Database Authentication

You can configure a user database that is local to the firewall to authenticate administrators who access the firewall web interface and to authenticate end users who access applications through Authentication Portal or GlobalProtect. Perform the following steps to configure Local Authentication with a local database.



External Authentication Services *are usually preferable to local authentication because they provide the benefit of central account management.*

You can also configure local authentication without a database, but only for firewall *or* Panorama *administrators.*

STEP 1 | Add the user account to the local database.

1. Select **Device** > **Local User Database** > **Users** and click **Add**.
2. Enter a user **Name** for the administrator.
3. Enter a **Password** and **Confirm Password** or enter a **Password Hash**.
4. **Enable** the account (enabled by default) and click **OK**.

STEP 2 | Add the user group to the local database.

Required if your users require group membership.

1. Select **Device** > **Local User Database** > **User Groups** and click **Add**.
2. Enter a **Name** to identify the group.
3. **Add** each user who is a member of the group and click **OK**.

STEP 3 | Configure an authentication profile.

The authentication profile defines authentication settings that are common to a set of users. Set the authentication **Type** to **Local Database**.

STEP 4 | Assign the authentication profile to an administrator account or to an Authentication policy rule for end users.

- **Administrators**—Configure a Firewall Administrator Account:

  Specify the **Name** of a user you defined earlier in this procedure.

  Assign the **Authentication Profile** that you configured for the account.

- **End users**—For the full procedure to configure authentication for end users, see Configure Authentication Policy.

STEP 5 | Verify that the firewall can Test Authentication Server Connectivity to authenticate users.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Configure an Authentication Profile and Sequence

An authentication profile defines the authentication service that validates the login credentials of administrators who access the firewall web interface and end users who access applications through Authentication Portal or GlobalProtect. The service can be Local Authentication that the firewall provides or External Authentication Services. The authentication profile also defines options such as Kerberos single sign-on (SSO).

Some networks have multiple databases (such as TACACS+ and LDAP) for different users and user groups. To authenticate users in such cases, configure an *authentication sequence*—a ranked order of authentication profiles that the firewall matches a user against during login. The firewall checks against each profile in sequence until one successfully authenticates the user. A user is denied access only if authentication fails for all the profiles in the sequence. The sequence can specify authentication profiles that are based on any authentication service that the firewall supports excepts Multi-Factor Authentication (MFA) and SAML.

STEP 1 | (External service only) Enable the firewall to connect to an external server for authenticating users:

1. Set up the external server. Refer to your server documentation for instructions.
2. Configure a server profile for the type of authentication service you use.
   - Add a RADIUS server profile.

     If the firewall integrates with an MFA service through RADIUS, you must add a RADIUS server profile. In this case, the MFA service provides all the authentication factors. If the firewall integrates with an MFA service through a vendor API, you can still use a RADIUS server profile for the first factor but MFA server profiles are required for additional factors.

   - Add an MFA server profile.
   - Add a SAML IdP server profile.
   - Add a Kerberos server profile.
   - Add a TACACS+ server profile.
   - Add an LDAP server profile.

STEP 2 | (Local database authentication only) Configure a user database that is local to the firewall.

Perform these steps for each user and user group for which you want to configure Local Authentication based on a user identity store that is local to the firewall:

1. Add the user account to the local database.
2. (Optional) Add the user group to the local database.

STEP 3 | (Kerberos SSO only) Create a Kerberos keytab for the firewall if Kerberos single sign-on (SSO) is the primary authentication service.

Create a Kerberos keytab. A keytab is a file that contains Kerberos account information for the firewall. To support Kerberos SSO, your network must have a Kerberos infrastructure.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 4 | Configure an authentication profile.

Define one or both of the following:

- **Kerberos SSO**—The firewall first tries SSO authentication. If that fails, it falls back to the specified authentication **Type**.
- **External authentication or local database authentication**—The firewall prompts the user to enter login credentials, and uses an external service or local database to authenticate the user.

1. Select **Device** > **Authentication Profile** and **Add** the authentication profile.
2. Enter a **Name** to identify the authentication profile.
3. Select the **Type** of authentication service.
   - If you use Multi-Factor Authentication, the selected type applies only to the first authentication factor. You select services for additional MFA factors in the **Factors** tab.
   - If you select **RADIUS**, **TACACS+**, **LDAP**, or **Kerberos**, select the **Server Profile**.
   - If you select **LDAP**, select the **Server Profile** and define the **Login Attribute**. For Active Directory, enter **sAMAccountName** as the value.
   - If you select **SAML**, select the **IdP Server Profile**.
   - If you select **Cloud Authentication Service**, configure a Cloud Identity Engine instance to communicate with the firewall. For more information on the Cloud Identity Engine, see the Cloud Identity Engine Getting Started guide.
4. If you want to enable Kerberos SSO, enter the **Kerberos Realm** (usually the DNS domain of the users, except that the realm is UPPERCASE) and **Import** the **Kerberos Keytab** that you created for the firewall or Panorama.
5. (MFA only) Select **Factors**, **Enable Additional Authentication Factors**, and **Add** the MFA server profiles you configured.

   The firewall will invoke each MFA service in the listed order, from top to bottom.
6. Select **Advanced** and **Add** the users and groups that can authenticate with this profile.

   You can select users and groups from the local database or, if you configured the firewall to Map Users to Groups, from an LDAP-based directory service such as Active Directory. By default, the list is empty, meaning no users can authenticate.

   You can also select custom groups defined in a group mapping configuration.
7. (Optional) To modify the user information before the firewall sends the authentication request to the server, configure a **Username Modifier**.
   - **%USERDOMAIN%\%USERINPUT%**—If the source does not include the domain (for example, it uses the sAMAccountName), the firewall adds the **User Domain** you specify before the username. If the source includes the domain, the firewall replaces that domain with the **User Domain**. If the **User Domain** is empty, the firewall removes

https://t.me/learningnets

the domain from the user information that the firewall receives from source before the firewall sends the request to the authentication server.

Because LDAP servers do not support backslashes in the sAMAccountName, do not use this option to authenticate with an LDAP server.

- **%USERINPUT%**—(Default) The firewall sends the user information to the authentication server in the format it receives from the source.
- **%USERINPUT%@%USERDOMAIN%**—If the source does not include the domain, the firewall adds the **User Domain** value after the username. If the source includes domain, the firewall replaces that domain with the **User Domain** value. If the **User Domain** is empty, the firewall removes the domain from the user information that the firewall receives from the source before the firewall sends the request to the authentication server.
- **None**—If you manually enter **None**:
  - For LDAP and Kerberos server profiles, the firewall uses the domain it receives from the source to select the appropriate authentication profile, then removes the domain when it sends the authentication request to the server. This allows you to include the **User Domain** during the authentication sequence but remove the domain before the firewall sends the authentication request to the server. For example, if you are using an LDAP server profile and the samAccountName as the attribute, use this option so that the firewall does not send the domain to the authentication server that expects only a username and not a domain.
  - For RADIUS server profiles:
    - If the source sends the user information in **domain\username** format, the firewall sends the user information to the server in the same format.
    - If the source sends the user information in **username@domain** format, the firewall normalizes the user information to the **domain\username** format before sending it to the server.
    - If the source sends only the username, the firewall adds the **User Domain** you specify before sending the information to the server in **domain\username** format.
  - For local databases, TACACS+, and SAML, the firewall sends the user information to the authentication server in the format it receives from the source.

8. Click **OK** to save the authentication profile.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 5 | Configure an authentication sequence.

Required if you want the firewall to try multiple authentication profiles to authenticate users. The firewall evaluates the profiles in top-to-bottom order until one profile successfully authenticates the user.

1. Select **Device** > **Authentication Sequence** and **Add** the authentication sequence.
2. Enter a **Name** to identify the authentication sequence.

   *To expedite the authentication process, **Use domain to determine authentication profile**: the firewall matches the domain name that a user enters during login with the **User Domain** or **Kerberos Realm** of an authentication profile in the sequence, and then uses that profile to authenticate the user. If the firewall does not find a match, or if you disable the option, the firewall tries the profiles in the top-to-bottom sequence.*
3. **Add** each authentication profile. To change the evaluation order of the profiles, select a profile and **Move Up** or **Move Down**.
4. Click **OK** to save the authentication sequence.

STEP 6 | Assign the authentication profile or sequence to an administrative account for firewall administrators or to Authentication policy for end users.

- **Administrators**—Assign the authentication profile based on how you manager administrator authorization:

  Authorization managed locally on the firewall—Configure a Firewall Administrator Account.

  Authorization managed on a SAML, TACACS+, or RADIUS server—Select **Device** > **Setup** > **Management**, edit the Authentication Settings, and select the **Authentication Profile**.
- **End users**—For the full procedure to configure authentication for end users, see Configure Authentication Policy.

STEP 7 | Verify that the firewall can Test Authentication Server Connectivity to authenticate users.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Test Authentication Server Connectivity

The test authentication feature enables you to verify whether the firewall or Panorama can communicate with the authentication server specified in an authentication profile and whether an authentication request succeeds for a specific user. You can test authentication profiles that authenticate administrators who access the web interface or that authenticate end users who access applications through GlobalProtect or Authentication Portal. You can perform authentication tests on the candidate configuration to verify the configuration is correct before committing.

STEP 1 | Configure an authentication profile. You do not need to commit the authentication profile or server profile configuration before testing.

STEP 2 | Log into the firewall CLI.

STEP 3 | (Firewalls with multiple virtual systems) Define the target virtual system that the test command will access.

This is required on firewalls with multiple virtual systems so that the test authentication command can locate the user you will test.

Define the target virtual system by entering:

```
admin@PA-325060> set system setting target-vsys <vsys-name>
```

For example, if the user is defined in vsys2, enter:

```
admin@PA-3250> set system setting target-vsys vsys2
```

*The **target-vsys** option is per login session; the firewall clears the option when you log off.*

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 4 | Test the authentication profile by entering the following command:

```
admin@PA-3250> test authentication authentication-
profile <authentication-profile-name> username <username> password
```

For example, to test an authentication profile named **my-profile** for a user named **bsimpson**, enter:

```
admin@PA-3250> test authentication authentication-profile my-
profile username bsimpson password
```

*When running the **test** command, the names of authentication profiles and server profiles are case sensitive. Also, if an authentication profile has a username modifier defined, you must enter the modifier with the username. For example, if you add the username modifier **%USERINPUT%@%USERDOMAIN%** for a user named **bsimpson** and the domain name is **mydomain.com**, enter **bsimpson@mydomain.com** as the username. This ensures that the firewall sends the correct credentials to the authentication server. In this example, mydomain.com is the domain that you define in the **User Domain** field in the authentication profile.*

STEP 5 | View the test output.

If the authentication profile is configured correctly, the output displays `Authentication succeeded`. If there is a configuration issue, the output displays information to help you troubleshoot the configuration.

The output results vary based on several factors related to the authentication type that you are testing as well as the type of issue. For example, RADIUS and TACACS+ use different underlying libraries, so the same issue that exists for both of these types will produce different errors. Also, if there is a network problem, such as using an incorrect port or IP address in the authentication server profile, the output error is not specific. This is because the test command cannot perform the initial handshake between the firewall and the authentication server to determine details about the issue.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Authentication Policy

Authentication policy enables you to authenticate end users before they can access services and applications. Whenever a user requests a service or application (such as by visiting a web page), the firewall evaluates Authentication policy. Based on the matching Authentication policy rule, the firewall then prompts the user to authenticate using one or more methods (factors), such as login and password, Voice, SMS, Push, or One-time Password (OTP) authentication. For the first factor, users authenticate through a Authentication Portal web form. For any additional factors, users authenticate through a Multi-Factor Authentication (MFA) login page.



To implement Authentication policy for GlobalProtect, refer to Configure GlobalProtect to facilitate multi-factor authentication notifications.

After the user authenticates for all factors, the firewall evaluates Security Policy to determine whether to allow access to the service or application.

To reduce the frequency of authentication challenges that interrupt the user workflow, you can specify a timeout period during which a user authenticates only for initial access to services and applications, not for subsequent access. Authentication policy integrates with Authentication Portal to record the timestamps used to evaluate the timeout and to enable user-based policies and reports.

Based on user information that the firewall collects during authentication, User-ID creates a new IP address-to-username mapping or updates the existing mapping for that user (if the mapping information has changed). The firewall generates User-ID logs to record the additions and updates. The firewall also generates an Authentication log for each request that matches an Authentication rule. If you favor centralized monitoring, you can configure reports based on User-ID or Authentication logs and forward the logs to Panorama or external services as you would for any other log types.

- Authentication Timestamps
- Configure Authentication Policy

Authentication Timestamps

When configuring an Authentication policy rule, you can specify a timeout period during which a user authenticates only for initial access to services and applications, not for subsequent access. Your goal is to specify a timeout that strikes a balance between the need to secure services and applications and the need to minimize interruptions to the user workflow. When a user authenticates, the firewall records a timestamp for the first authentication challenge (factor) and a timestamp for any additional Multi-Factor Authentication (MFA) factors. When the user subsequently requests services and applications that match an Authentication rule, the firewall evaluates the timeout specified in the rule relative to each timestamp. This means the firewall reissues authentication challenges on a per-factor basis when timeouts expire. If you Redistribute User Mappings and Authentication Timestamps, all your firewalls will enforce Authentication policy timeouts consistently for all users.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

The firewall records a separate timestamp for each MFA vendor. For example, if you use Duo v2 and PingID servers to issue challenges for MFA factors, the firewall records one timestamp for the response to the Duo factor and one timestamp for the response to the PingID factor.

Within the timeout period, a user who successfully authenticates for one Authentication rule can access services or applications that other rules protect. However, this portability applies only to rules that trigger the same authentication factors. For example, a user who successfully authenticates for a rule that triggers TACACS+ authentication must authenticate again for a rule that triggers SAML authentication, even if the access requests are within the timeout period for both rules.

When evaluating the timeout in each Authentication rule and the global timer defined in the Authentication Portal settings (see Configure Authentication Portal), the firewall prompts the user to re-authenticate for whichever setting expires first. Upon re-authenticating, the firewall records new authentication timestamps for the rules and resets the time count for the Authentication Portal timer. Therefore, to enable different timeout periods for different Authentication rules, set the Authentication Portal timer to a value that is the same as or higher than the timeout in any rule.

Configure Authentication Policy

Perform the following steps to configure Authentication policy for end users who access services through Authentication Portal. Before starting, ensure that your Security Policy allows users to access the services and URL categories that require authentication.

STEP 1 | Configure Authentication Portal. If you use Multi-Factor Authentication (MFA) services to authenticate users, you must set the **Mode** to **Redirect**.

STEP 2 | Configure the firewall to use one of the following services to authenticate users.

- External Authentication Services—Configure a server profile to define how the firewall connects to the service.
- Local database authentication—Add each user account to the local user database on the firewall.
- Kerberos single sign-on (SSO)—Create a Kerberos keytab for the firewall. Optionally, you can configure the firewall to use Kerberos SSO as the primary authentication service and, if SSO failures occur, fall back to an external service or local database authentication.

STEP 3 | Configure an Authentication Profile and Sequence for each set of users and Authentication policy rules that require the same authentication services and settings.

Select the **Type** of authentication service and related settings:

- **External service**—Select the **Type** of external server and select the **Server Profile** you created for it.
- **Local database authentication**—Set the **Type** to **Local Database**. In the **Advanced** settings, **Add** the Authentication Portal users and user groups you created.
- **Kerberos SSO**—Specify the **Kerberos Realm** and **Import** the **Kerberos Keytab**.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 4 | Configure an authentication enforcement object.

The object associates each authentication profile with an Authentication Portal method. The method determines whether the first authentication challenge (factor) is transparent or requires a user response.

1. Select **Objects** > **Authentication** and **Add** an object.
2. Enter a **Name** to identify the object.
3. Select an **Authentication Method** for the authentication service **Type** you specified in the authentication profile:
   - **browser-challenge**—Select this method if you want the client browser to respond to the first authentication factor instead of having the user enter login credentials. For this method, you must configure Kerberos SSO in the authentication profile. If the browser challenge fails, the firewall falls back to the **web-form** method.
   - **web-form**—Select this method if you want the firewall to display a Authentication Portal web form for users to enter login credentials.
4. Select the **Authentication Profile** you configured.
5. Enter the **Message** that the Authentication Portal web form will display to tell users how to authenticate for the first authentication factor.
6. Click **OK** to save the object.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 5 | Configure an Authentication policy rule.

Create a rule for each set of users, services, and URL categories that require the same authentication services and settings.

*The firewall does not apply the Authentication Portal timeout if your authentication policy uses default authentication enforcement objects (for example, **default-browser-challenge**).To require users to re-authenticate after the Authentication Portal timeout, clone the rule for the default authentication object and move it before the existing rule for the default authentication object.*

1. Select **Policies** > **Authentication** and **Add** a rule.
2. Enter a **Name** to identify the rule.
3. Select **Source** and **Add** specific zones and IP addresses or select **Any** zones or IP addresses.

   The rule applies only to traffic coming from the specified IP addresses or from interfaces in the specified zones.
4. Select **User** and select or **Add** the source users and user groups to which the rule applies (default is **any**).
5. Select or **Add** the Host Information Profiles to which the rule applies (default is **any**).
6. Select **Destination** and **Add** specific zones and IP addresses or select **any** zones or IP addresses.

   The IP addresses can be resources (such as servers) for which you want to control access.
7. Select **Service/URL Category** and select or **Add** the services and service groups for which the rule controls access (default is **service-http**).
8. Select or **Add** the URL Categories for which the rule controls access (default is **any**). For example, you can create a custom URL category that specifies your most sensitive internal sites.
9. Select **Actions** and select the **Authentication Enforcement** object you created.
10. Specify the **Timeout** period in minutes (default 60) during which the firewall prompts the user to authenticate only once for repeated access to services and applications.

    ***Timeout** is a tradeoff between tighter security (less time between authentication prompts) and the user experience (more time between authentication prompts). More frequent authentication is often the right choice for access to critical systems and sensitive areas such as a data center. Less frequent authentication is often the right choice at the network perimeter and for businesses for which the user experience is key.*
11. Click **OK** to save the rule.

STEP 6 | (MFA only) Customize the MFA login page.

The firewall displays this page so that users can authenticate for any additional MFA factors.

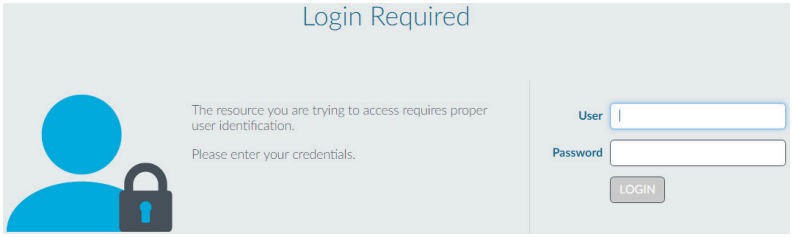
 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 7 | Verify that the firewall enforces Authentication policy.

1. Log in to your network as one of the source users specified in an Authentication policy rule.
2. Request a service or URL category that matches one specified in the rule.

   The firewall displays the Authentication Portal web form for the first authentication factor. For example:



   If you configured the firewall to use one or more MFA services, authenticate for the additional authentication factors.
3. End the session for the service or URL you just accessed.
4. Start a new session for the same service or application. Be sure to perform this step within the **Timeout** period you configured in the Authentication rule.

   The firewall allows access without re-authenticating.
5. Wait until the **Timeout** period expires and request the same service or application.

   The firewall prompts you to re-authenticate.

STEP 8 | (Optional) Redistribute Data and Authentication Timestamps to other firewalls that enforce Authentication policy to ensure they all apply the timeouts consistently for all users.

https://t.me/learningnets

Troubleshoot Authentication Issues

When users fail to authenticate to a Palo Alto Networks firewall or Panorama, or the Authentication process takes longer than expected, analyzing authentication-related information can help you determine whether the failure or delay resulted from:

- **User behavior**—For example, users are locked out after entering the wrong credentials or a high volume of users are simultaneously attempting access.
- **System or network issues**—For example, an authentication server is inaccessible.
- **Configuration issues**—For example, the Allow List of an authentication profile doesn't have all the users it should have.

The following CLI commands display information that can help you troubleshoot these issues:

Task	Command

Display the number of locked user accounts associated with the authentication profile (**auth-profile**), authentication sequence (**is-seq**), or virtual system (**vsys**).
 	
To unlock users, use the following operational command:
 	
`> request authentication [unlock-admin | unlock-user]`
	`PA-220> show authentication locked-users`
	`{`
	`vsys <value> |`
	`auth-profile <value> |`
	`is-seq`
	`{yes | no}`
	`{auth-profile | vsys} <value>`
	`}`

Use the **debug authentication** command to troubleshoot authentication events.
 	
Use the **show** options to display authentication request statistics and the current debugging level:
• **show** displays the current debugging level for the authentication service (authd).
• **show-active-requests** displays the number of active checks for authentication requests, allow lists, locked user accounts, and Multi-Factor Authentication (MFA) requests.
• **show-pending-requests** displays the number of pending checks for authentication requests, allow lists, locked user accounts, and MFA requests.
	`PA-220> debug authentication`
	`{`
	`on {debug | dump | error | info | warn} |`
	`show |`
	`show-active-requests |`
	`show-pending-requests |`
	`connection-show |`
	`{`
	`connection-id |`
	`protocol-type`
	`{`
	`Kerberos connection-id <value> |`
	`LDAP connection-id <value> |`
	`RADIUS connection-id <value> |`
	`TACACS+ connection-id <value> |`
	`}`

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Task	Command
• **connection-show** displays authentication request and response statistics for all authentication servers or for a specific protocol type. Use the **connection-debug** options to enable or disable authentication debugging: • Use the **on** option to enable or the **off** option to disable debugging for authd. • Use the **connection-debug-on** option to enable or the **connection-debug-off** option to disable debugging for all authentication servers or for a specific protocol type.	**connection-debug-on |** **{** **connection-id |** **debug-prefix |** **protocol-type** **{** **Kerberos connection-id <value> |** **LDAP connection-id <value> |** **RADIUS connection-id <value> |** **TACACS+ connection-id <value> |** **}** **connection-debug-off |** **{** **connection-id |** **protocol-type** **{** **Kerberos connection-id <value> |** **LDAP connection-id <value> |** **RADIUS connection-id <value> |** **TACACS+ connection-id <value> |** **}** **connection-debug-on** **}**
Test the connection and validity of the certificate profile.	PA-220> test authentication authentication-profile auth-profile username <*username*>password <*password*>
Troubleshoot a specific authentication using the **Authentication ID** displayed in **Monitor** > **Logs** > **Authentication**.	PA-220> grep <*Authentication ID*>

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Certificate Management

The following topics describe the different keys and certificates that Palo Alto Networks® firewalls and Panorama use, and how to obtain and manage them:

- Keys and Certificates
- Default Trusted Certificate Authorities (CAs)
- Certificate Revocation
- Certificate Deployment
- Set Up Verification for Certificate Revocation Status
- Configure the Master Key
- Master Key Encryption
- Obtain Certificates
- Export a Certificate and Private Key
- Configure a Certificate Profile
- Configure an SSL/TLS Service Profile
- Configure an SSH Service Profile
- Replace the Certificate for Inbound Management Traffic
- Configure the Key Size for SSL Forward Proxy Server Certificates
- Revoke and Renew Certificates
- Secure Keys with a Hardware Security Module

https://t.me/learningnets

Keys and Certificates

To ensure trust between parties in a secure communication session, Palo Alto Networks firewalls and Panorama use digital certificates. Each certificate contains a cryptographic key to encrypt plaintext or decrypt ciphertext. Each certificate also includes a digital signature to authenticate the identity of the issuer. The issuer must be in the list of trusted certificate authorities (CAs) of the authenticating party. Optionally, the authenticating party verifies the issuer did not revoke the certificate (see Certificate Revocation).

Palo Alto Networks firewalls and Panorama use certificates in the following applications:

- User authentication for Authentication Portal, multi-factor authentication (MFA), and web interface access to a firewall or Panorama.
- Device authentication for GlobalProtect VPN (remote user-to-site or large scale).
- Device authentication for IPSec site-to-site VPN with Internet Key Exchange (IKE).
- External dynamic list (EDL) validation.
- User-ID agent and TS agent access.
- Decrypting inbound and outbound SSL traffic.

  A firewall decrypts the traffic to apply policy rules, then re-encrypts it before forwarding the traffic to the final destination. For outbound traffic, the firewall acts as a forward proxy server, establishing an SSL/TLS connection to the destination server. To secure a connection between itself and the client, the firewall uses a *signing certificate* to automatically generate a copy of the destination server certificate.

The following table describes the keys and certificates that Palo Alto Networks firewalls and Panorama use. As a best practice, use different keys and certificates for each usage.

Table 1: Palo Alto Networks Device Keys/Certificates

Key/Certificate Usage	Description
Administrative Access	Secure access to firewall or Panorama administration interfaces (HTTPS access to the web interface) requires a server certificate for the MGT interface (or a designated interface on the dataplane if the firewall or Panorama does not use MGT) and, optionally, a certificate to authenticate the administrator.
Authentication Portal	In deployments where Authentication policy identifies users who access HTTPS resources, designate a server certificate for the Authentication Portal interface. If you configure Authentication Portal to use certificates for identifying users (instead of, or in addition to, interactive authentication), deploy client certificates also. For more information on Authentication Portal, see Map IP Addresses to Usernames Using Authentication Portal.

https://t.me/learningnets

Key/Certificate Usage	Description
Forward Trust	For outbound SSL/TLS traffic, if a firewall acting as a forward proxy trusts the CA that signed the certificate of the destination server, the firewall uses the forward trust CA certificate to generate a copy of the destination server certificate to present to the client. To set the private key size, see Configure the Key Size for SSL Forward Proxy Server Certificates. For added security, store the key on a hardware security module (for details, see Secure Keys with a Hardware Security Module).
Forward Untrust	For outbound SSL/TLS traffic, if a firewall acting as a forward proxy does not trust the CA that signed the certificate of the destination server, the firewall uses the forward untrust CA certificate to generate a copy of the destination server certificate to present to the client.
SSL Inbound Inspection	The keys that decrypt inbound SSL/TLS traffic for inspection and policy enforcement. For this application, import onto the firewall a private key for each server that is subject to SSL/TLS inbound inspection. See Configure SSL Inbound Inspection. *Beginning in PAN-OS 8.0, firewalls use the Elliptic-Curve Diffie-Hellman Ephemeral (ECDHE) algorithm to perform strict certificate checking. This means that if the firewall uses an intermediate certificate, you must reimport the certificate from your web server to the firewall after you upgrade to a PAN-OS 8.0 or later release and combine the server certificate with the intermediate certificate (install a chained certificate). Otherwise, SSL Inbound Inspection sessions that have an intermediate certificate in the chain will fail. To install a chained certificate:* 1. *Open each certificate (.cer) file in a plain-text editor such as Notepad.* 2. *Paste each certificate end-to-end with the Server Certificate at the top with each signer included below.* 3. *Save the file as a text (.txt) or certificate (.cer) file (the name of the file cannot contain blank spaces).* 4. *Import the combined (chained) certificate into the firewall.*
SSL Exclude Certificate	Certificates for servers to exclude from SSL/TLS decryption. For example, if you enable SSL decryption but your network includes servers for which the firewall should not decrypt traffic (for example, web services for your HR systems), import the corresponding certificates onto the firewall and configure them as SSL Exclude Certificates. See Decryption Exclusions.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Key/Certificate Usage	Description
GlobalProtect	All interaction among GlobalProtect components occurs over SSL/TLS connections. Therefore, as part of the GlobalProtect deployment, deploy server certificates for all GlobalProtect portals, gateways, and Mobile Security Managers. Optionally, deploy certificates for authenticating users also.  *The GlobalProtect Large Scale VPN (LSVPN) feature requires a CA signing certificate.*
Site-to-Site VPNs (IKE)	In a site-to-site IPSec VPN deployment, peer devices use Internet Key Exchange (IKE) gateways to establish a secure channel. IKE gateways use certificates or preshared keys to authenticate the peers to each other. You configure and assign the certificates or keys when defining an IKE gateway on a firewall. See Site-to-Site VPN Overview.
Master Key	The firewall uses a master key to encrypt all private keys and passwords. If your network requires a secure location for storing private keys, you can use an encryption (wrapping) key stored on a hardware security module (HSM) to encrypt the master key. For details, see Encrypt a Master Key Using an HSM.
Secure Syslog	The certificate to enable secure connections between the firewall and a syslog server. See Syslog Field Descriptions.
Trusted Root CA	The designation for a root certificate issued by a CA that the firewall trusts. The firewall can use a self-signed root CA certificate to automatically issue certificates for other applications (for example, SSL Forward Proxy). Also, if a firewall must establish secure connections with other firewalls, the root CA that issues their certificates must be in the list of trusted root CAs on the firewall.
Inter-Device Communication	By default, Panorama, firewalls, and Log Collectors use a set of predefined certificates for the SSL/TLS connections used for management and log forwarding. However, you can enhance these connection by deploying custom certificates to the devices in your deployment. These certificates can also be used to secure the SSL/TLS connection between Panorama HA peers.

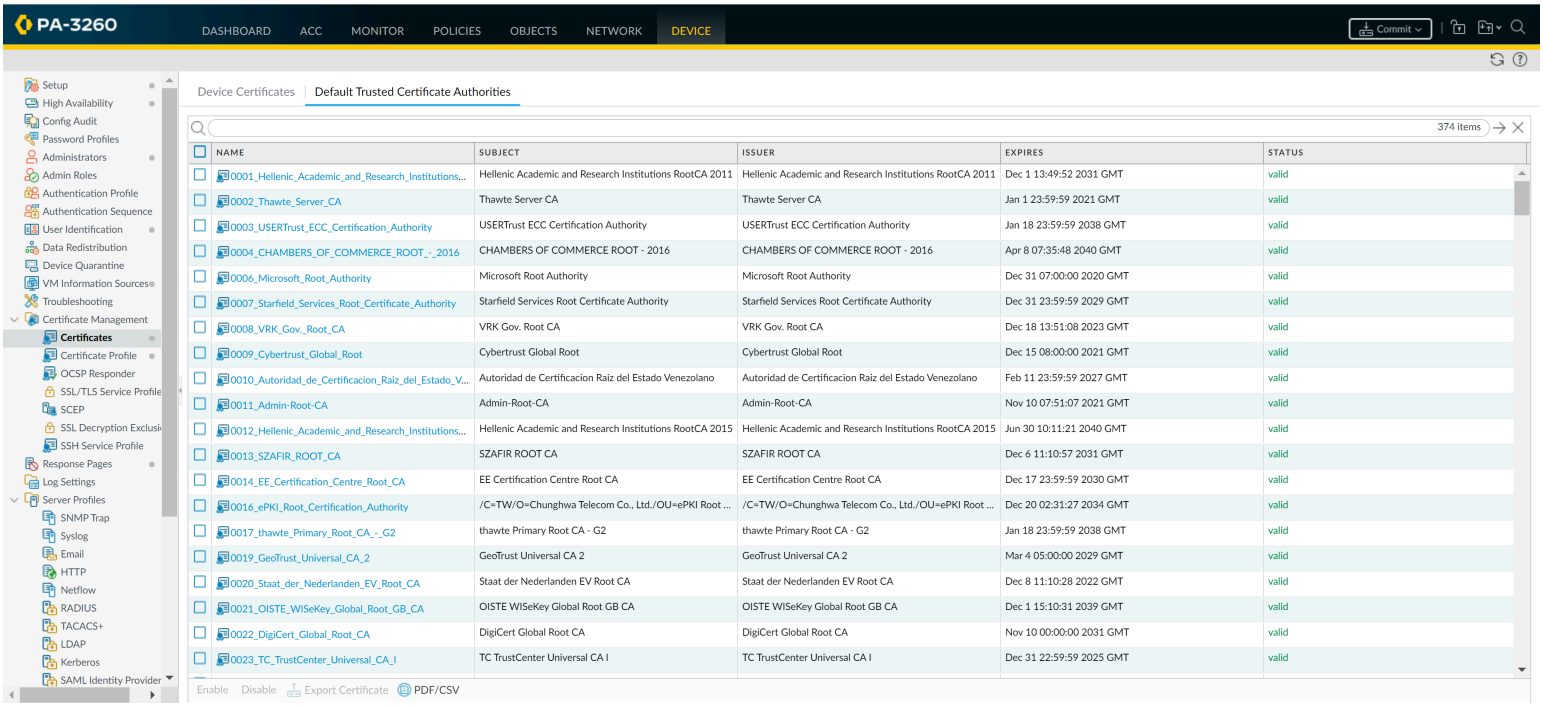
 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Default Trusted Certificate Authorities (CAs)

The firewall trusts the most common and trusted authorities (CAs) by default. These trusted certificate providers are responsible for issuing the certificates the firewall requires to secure connections to the internet.

To view and manage the list of CAs that the firewall trusts by default, select **Device** > **Certificate Management** > **Certificates** > **Default Trusted Certificate Authorities**:



The only additional CAs you might want to add are trusted enterprise CAs that your organization requires—see Obtain Certificates.

https://t.me/learningnets

Certificate Revocation

Palo Alto Networks firewalls and Panorama use digital certificates to ensure trust between parties in a secure communication session. Configuring a firewall or Panorama to check the revocation status of certificates provides additional security. A party that presents a revoked certificate is not trustworthy. When a certificate is part of a chain, the firewall or Panorama checks the status of every certificate in the chain except the root CA certificate, for which it cannot verify revocation status.

Various circumstances can invalidate a certificate before the expiration date. Some examples are a change of name, change of association between subject and certificate authority (for example, an employee terminates employment), and compromise (known or suspected) of the private key. Under such circumstances, the certificate authority that issued the certificate must revoke it.

The firewall and Panorama support the following methods for verifying certificate revocation status. If you configure both methods, the firewall or Panorama first tries the OCSP method; if the OCSP server is unavailable, it uses the CRL method.

- Certificate Revocation List (CRL)
- Online Certificate Status Protocol (OCSP)



In PAN-OS, certificate revocation status verification is an optional feature. It is a best practice to enable it for certificate profiles, which define user and device authentication for Authentication Portal, GlobalProtect, site-to-site IPSec VPN, and web interface access to the firewall or Panorama, to verify that the certificate hasn't been revoked.

Certificate Revocation List (CRL)

Each certificate authority (CA) periodically issues a certificate revocation list (CRL) to a public repository. The CRL identifies revoked certificates by serial number. After the CA revokes a certificate, the next CRL update will include the serial number of that certificate. The firewall supports CRLs in Distinguished Encoding Rules (DER) and Privacy Enhanced Mail (PEM) formats.

The Palo Alto Networks firewall downloads and caches the last-issued CRL for every CA listed in the trusted CA list of the firewall. Caching only applies to validated certificates; if a firewall never validated a certificate, the firewall cache does not store the CRL for the issuing CA. Also, the cache only stores a CRL until it expires.



If you configure multiple CRL distribution points (CDPs) and the firewall cannot reach the first CDP, the firewall does not check the remaining CDPs. To redirect invalid CRL requests, configure a DNS proxy *as an alternate server.*

To use CRLs for verifying the revocation status of certificates used for the decryption of inbound and outbound SSL/TLS traffic, see Configure Revocation Status Verification of Certificates Used for SSL/TLS Decryption.

To use CRLs for verifying the revocation status of certificates that authenticate users and devices, configure a certificate profile and assign it to the interfaces that are specific to the application: Authentication Portal, GlobalProtect (remote user-to-site or large scale), site-to-site IPSec VPN, or web interface access to Palo Alto Networks firewalls or Panorama. For details, see Configure Revocation Status Verification of Certificates.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Online Certificate Status Protocol (OCSP)

When establishing an SSL/TLS session, clients can use Online Certificate Status Protocol (OCSP) to check the revocation status of the authentication certificate. The authenticating client sends a request containing the serial number of the certificate to the OCSP responder (server). The responder searches the database of the certificate authority (CA) that issued the certificate and returns a response containing the status (good, revoked or unknown) to the client. The advantage of the OCSP method is that it can verify status in real-time, instead of depending on the issue frequency (hourly, daily, or weekly) of CRLs.

The Palo Alto Networks firewall downloads and caches OCSP status information for every CA listed in the trusted CA list of the firewall. Caching only applies to validated certificates; if a firewall never validated a certificate, the firewall cache does not store the OCSP information for the issuing CA. If your enterprise has its own public key infrastructure (PKI), you can configure the firewall as an OCSP responder (see Configure an OCSP Responder).

To use OCSP for verifying the revocation status of certificates when the firewall functions as an SSL forward proxy, perform the steps under Configure Revocation Status Verification of Certificates Used for SSL/TLS Decryption.

The following applications use certificates to authenticate users and/or devices: Authentication Portal, GlobalProtect (remote user-to-site or large scale), site-to-site IPSec VPN, and web interface access to Palo Alto Networks firewalls or Panorama. To use OCSP for verifying the revocation status of the certificates:

- Configure an OCSP responder (if you are configuring the firewall as an OCSP responder).
- Enable the HTTP OCSP service on the firewall (if you are configuring the firewall as an OCSP responder).
- Create or obtain a certificate for each application.
- Configure a certificate profile for each application.
- Assign the certificate profile to the relevant application.

To cover situations where the OCSP responder is unavailable, configure CRL as a fall-back method. For details, see Configure Revocation Status Verification of Certificates.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Certificate Deployment

The basic approaches to deploy certificates for Palo Alto Networks firewalls or Panorama are:

- **Obtain certificates from a trusted third-party CA**—The benefit of obtaining a certificate from a trusted third-party certificate authority (CA) such as VeriSign or GoDaddy is that end clients will already trust the certificate because common browsers include root CA certificates from well-known CAs in their trusted root certificate stores. Therefore, for applications that require end clients to establish secure connections with the firewall or Panorama, purchase a certificate from a CA that the end clients trust to avoid having to pre-deploy root CA certificates to the end clients. (Some such applications are a GlobalProtect portal or GlobalProtect Mobile Security Manager.) However, most third-party CAs cannot issue signing certificates. Therefore, this type of certificate is not appropriate for applications (for example, SSL/TLS decryption and large-scale VPN) that require the firewall to issue certificates. See Obtain a Certificate from an External CA.
- **Obtain certificates from an enterprise CA**—Enterprises that have their own internal CA can use it to issue certificates for firewall applications and import them onto the firewall. The benefit is that end clients probably already trust the enterprise CA. You can either generate the needed certificates and import them onto the firewall, or generate a certificate signing request (CSR) on the firewall and send it to the enterprise CA for signing. The benefit of this method is that the private key does not leave the firewall. An enterprise CA can also issue a signing certificate, which the firewall uses to automatically generate certificates (for example, for GlobalProtect large-scale VPN or sites requiring SSL/TLS decryption). See Import a Certificate and Private Key.
- **Generate self-signed certificates**—You can Create a Self-Signed Root CA Certificate on the firewall and use it to automatically issue certificates for other firewall applications.

  If you use this method to generate certificates for an application that requires an end client to trust the certificate, end users will see a certificate error because the root CA certificate is not in their trusted root certificate store. To prevent this, deploy the self-signed root CA certificate to all end user systems. You can deploy the certificates manually or use a centralized deployment method such as an Active Directory Group Policy Object (GPO).

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Set Up Verification for Certificate Revocation Status

To verify the revocation status of certificates, the firewall uses Online Certificate Status Protocol (OCSP) and/or certificate revocation lists (CRLs). For details on these methods, see Certificate Revocation If you configure both methods, the firewall first tries OCSP and only falls back to the CRL method if the OCSP responder is unavailable. If your enterprise has its own public key infrastructure (PKI), you can configure the firewall to function as the OCSP responder.

The following topics describe how to configure the firewall to verify certificate revocation status:

- Configure an OCSP Responder
- Configure Revocation Status Verification of Certificates
- Configure Revocation Status Verification of Certificates Used for SSL/TLS Decryption

Configure an OCSP Responder

To use Online Certificate Status Protocol (OCSP) for verifying the revocation status of certificates, you must configure the firewall to access an OCSP responder (server). The entity that manages the OCSP responder can be a third-party certificate authority (CA). If your enterprise has its own public key infrastructure (PKI), you can use external OCSP responders or you can configure the firewall itself as an OCSP responder. For details on OCSP, see Certificate Revocation.

Configure an OCSP responder Certificate Profile *only when you generate a new certificate (**Device** > **Certificate Management** > **Certificates**). Specify the **OCSP Responder** when you generate a new certificate so that the firewall populates the Authority Information Access (AIA) field with the appropriate URL and then specify the new certificate in the Certificate Profile. Configuring a Certificate Profile does not override the Certificate Profile for existing certificates or Root CAs.*

You can enable OCSP validation or override the AIA field of certificate in the Certificate Profile*. The Certificate Profile configuration determines which certificate validation mechanisms are used on certificates that authenticate to services hosted on the firewall, such as GlobalProtect.*

STEP 1 | Define an external OCSP responder or configure the firewall itself as an OCSP responder.

1. Select **Device** > **Certificate Management** > **OCSP Responder** and click **Add**.
2. Enter a **Name** to identify the responder (up to 31 characters). The name is case-sensitive. It must be unique and use only letters, numbers, spaces, hyphens, and underscores.
3. If the firewall has more than one virtual system (vsys), select a **Location** (vsys or **Shared**) for the certificate.
4. In the **Host Name** field, enter the host name (recommended) or IP address of the OCSP responder. You can enter an IPv4 or IPv6 address. From this value, PAN-OS automatically derives a URL and adds it to the certificate being verified.

   If you configure the firewall itself as an OCSP responder, the host name must resolve to an IP address in the interface that the firewall uses for OCSP services.
5. Click **OK**.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 2 | If you want the firewall to use the management interface for the OCSP responder interface, enable OCSP communication on the firewall. Otherwise, continue to the next step to configure an alternate interface.

1. Select **Device** > **Setup** > **Interfaces** > **Management**.
2. In the Network Services section, select the **HTTP OCSP** check box, then click **OK**.

STEP 3 | To use an alternate interface as the OCSP responder interface, add an Interface Management Profile to the interface used for OCSP services.

1. Select **Network** > **Network Profiles** > **Interface Mgmt**.
2. Click **Add** to create a new profile or click the name of an existing profile.
3. Select the **HTTP OCSP** check box and click **OK**.
4. Select **Network** > **Interfaces** and click the name of the interface that the firewall will use for OCSP services. The OCSP **Host Name** specified in Step 1 must resolve to an IP address in this interface.
5. Select **Advanced** > **Other info** and select the Interface Management Profile you configured.
6. Click **OK** and **Commit**.

Configure Revocation Status Verification of Certificates

The firewall and Panorama use certificates to authenticate users and devices for such applications as Authentication Portal, GlobalProtect, site-to-site IPSec VPN, and web interface access to the firewall/Panorama. To improve security, it is a best practice to configure the firewall or Panorama to verify the revocation status of certificates that it uses for device/user authentication.

STEP 1 | Configure a Certificate Profile for each application.

Assign one or more root CA certificates to the profile and select how the firewall verifies certificate revocation status.

For details on the certificates that various applications use, see Keys and Certificates

STEP 2 | Assign the certificate profiles to the relevant applications.

The steps to assign a certificate profile depend on the application that requires it.

Configure Revocation Status Verification of Certificates Used for SSL/TLS Decryption

The firewall decrypts inbound and outbound SSL/TLS traffic to inspect the traffic for threats. When you create a Security policy rule that allows traffic and apply Security profiles to the rule, create an analogous Decryption policy rule to decrypt that traffic. If you don't decrypt the traffic, the firewall can't use the Security profiles to inspect the traffic (you can't inspect what you can't see). The firewall re-encrypts the traffic before forwarding it. (See SSL Inbound Inspection and SSL Forward Proxy.) You can configure the firewall to verify the revocation status of certificates used for decryption as follows.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Enabling revocation status verification for SSL/TLS decryption certificates will add time to the process of establishing the session. The first attempt to access a site might fail if the verification does not finish before the session times out. For these reasons, verification is disabled by default.

STEP 1 | Define the service-specific timeout intervals for revocation status requests.

1. Select **Device** > **Setup** > **Session** and, in the Session Features section, select **Decryption Certificate Revocation Settings**.
2. Perform one or both of the following steps, depending on whether the firewall will use Online Certificate Status Protocol (OCSP) or the Certificate Revocation List (CRL) method to verify the revocation status of certificates. If the firewall will use both, it first tries OCSP; if the OCSP responder is unavailable, the firewall then tries the CRL method.
   - In the CRL section, select the **Enable** check box and enter the **Receive Timeout**. This is the interval (1-60 seconds) after which the firewall stops waiting for a response from the CRL service.
   - In the OCSP section, select the **Enable** check box and enter the **Receive Timeout**. This is the interval (1-60 seconds) after which the firewall stops waiting for a response from the OCSP responder.

   Depending on the **Certificate Status Timeout** value you specify in Step 2, the firewall might register a timeout before either or both of the **Receive Timeout** intervals pass.

STEP 2 | Define the total timeout interval for revocation status requests.

Enter the **Certificate Status Timeout**. This is the interval (1-60 seconds) after which the firewall stops waiting for a response from any certificate status service and applies the session-blocking logic you optionally define in Step 3. The **Certificate Status Timeout** relates to the OCSP/CRL **Receive Timeout** as follows:

- If you enable both OCSP and CRL—The firewall registers a request timeout after the lesser of two intervals passes: the **Certificate Status Timeout** value or the aggregate of the two **Receive Timeout** values.
- If you enable only OCSP—The firewall registers a request timeout after the lesser of two intervals passes: the **Certificate Status Timeout** value or the OCSP **Receive Timeout** value.
- If you enable only CRL—The firewall registers a request timeout after the lesser of two intervals passes: the **Certificate Status Timeout** value or the CRL **Receive Timeout** value.

STEP 3 | Define the blocking behavior for unknown certificate status or a revocation status request timeout.

If you want the firewall to block SSL/TLS sessions when the OCSP or CRL service returns a certificate revocation status of unknown, select the **Block Session With Unknown Certificate Status** check box. Otherwise, the firewall proceeds with the session.

If you want the firewall to block SSL/TLS sessions after it registers a request timeout, select the **Block Session On Certificate Status Check Timeout** check box. Otherwise, the firewall proceeds with the session.

STEP 4 | Click **OK** and **Commit**.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Configure the Master Key

Every firewall and Panorama management server has a default master key that encrypts all the private keys and passwords in the configuration to secure them (such as the private key used for SSL Forward Proxy Decryption).

Change the default master key as soon as possible to ensure that you use a unique master key for encryption.

In a high availability (HA) configuration, you must use the same master key on both firewalls because the master key is not synchronized across HA peers. Otherwise, HA synchronization will not work properly.

If you are using Panorama to manage your firewalls, you can configure the same master key on Panorama and all managed firewalls or configure a unique master key for each managed firewall. For managed firewalls in an HA configuration, you must configure the same master key for each HA peer. See Manage the Master Key from Panorama if the firewall is managed by a Panorama™ management server.

Be sure to store the master key in a safe location. You cannot recover the master key and the only way to restore the default master key is to Reset the Firewall to Factory Default Settings.

STEP 1 | Backup the configuration.

STEP 2 | (HA only) Disable Config Sync.

This step is required before deploying a new master key to any firewall HA pair

Before you deploy a new master key to any firewall HA pair, you must disable Config Sync. For Panorama-managed firewalls, if you do not disable Config Sync before deploying a new master key, Panorama loses connectivity to the primary firewall.

1. Select **Device** > **High Availability** > **General** and edit the **Setup**.
2. Disable (clear) **Enable Config Sync** and then click **OK**.
3. **Commit** your configuration changes.

STEP 3 | Select **Device** > **Master Key and Diagnostics** and edit the Master Key section.

STEP 4 | Enter the **Current Master Key** if one exists.

STEP 5 | Define a new **New Master Key** and then **Confirm New Master Key**. The key must contain exactly 16 characters.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 6 | To specify the master key **Lifetime**, enter the number of **Days** and/or **Hours** after which the key will expire.

You must configure a new master key before the current key expires. If the master key expires, the firewall or Panorama automatically reboots in Maintenance mode. You must then Reset the Firewall to Factory Default Settings.



*Set the **Lifetime** to two years or less, depending on how many encryptions the device performs. The more encryptions a device performs, the shorter the **Lifetime** you should set. The critical consideration is to not run out of unique encryptions before you change the master key. Each master key can provide up to 2^{32} unique encryptions based on the master key value and the Initialization Vector (IV) value. After 2^{32} unique encryptions, encryptions repeat (are no longer unique), which is a security risk.*

*Set a **Time for Reminder** value (see next step) for the master key and when the reminder notification occurs, change the master key.*

STEP 7 | Enter a **Time for Reminder** that specifies the number of **Days** and **Hours** before the master key expires when the firewall generates an expiration alarm. The firewall automatically opens the System Alarms dialog to display the alarm.



*Set the reminder so that it gives you plenty of time to configure a new master key before it expires in a scheduled maintenance window. When the **Time for Reminder** expires and the firewall or Panorama sends a notification log, change the master key, don't wait for the **Lifetime** to expire. For grouped devices, track every device (e.g., firewalls that Panorama manages and firewall HA pairs) and when the reminder value expires for the any device in the group, change the master key.*

*To ensure the expiration alarm displays, select **Device** > **Log Settings**, edit the Alarm Settings, and **Enable Alarms**.*

STEP 8 | Enable **Auto Renew Master Key** to configure the firewall to automatically renew the master key. To configure **Auto Renew With Same Master Key**, specify the number of **Days** and/or **Hours** to renew the same master key. The key extension allows the firewall to remain operational and continue securing your network; it is not a replacement for configuring a new key if the existing master key lifetime expires soon.

Automatically renewing the master key has benefits and risks. The benefit is that extending the master key **Lifetime** protects against failure to change the master key before its lifetime expires. The risk is that encryptions will repeat and cause a security risk if the number

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

of encryptions the device performs with the master key exceeds the number of unique encryptions the master key can generate (2^{32} unique encryptions).

If the Master Key expires (you do not automatically renew it and you do not replace it in a timely manner), the device goes into maintenance mode.

*If you enable **Auto Renew Master Key**, set it so that the total time (lifetime plus the auto renew time) does not cause the device to run out of unique encryptions. For example, if you believe the device will consume the master key's number of unique encryptions in two and a half years, you could set the **Lifetime** for two years, set the **Time for Reminder** to 60 days, and set the **Auto Renew Master Key** for 60-90 days to provide the extra time to configure a new master key before the **Lifetime** expires. However, the best practice is still to change the master key before the lifetime expires to ensure that no device repeats encryptions.*

Consider the number of days until your next available maintenance window when configuring the master key to automatically renew after the lifetime of the key expires.

STEP 9 | (Optional) For added security, select whether to use an **HSM** to encrypt the master key. For details, see Encrypt a Master Key Using an HSM.

STEP 10 | Click **OK** and **Commit**.

STEP 11 | (HA only) Re-enable Config Sync.

1. Select **Device** > **High Availability** > **General** and edit the **Setup**.
2. Enable (check) **Enable Config Sync** and then click **OK**.
3. **Commit** your configuration changes.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Master Key Encryption

On physical and virtual Palo Alto Networks devices, you can configure the master key to use the AES-256-CBC or the AES-256-GCM (introduced in PAN-OS 10.0) encryption algorithm to encrypt data such as keys and passwords. AES-256-GCM provides stronger encryption than AES-256-CBC and improves your security posture. It also includes a built-in integrity check. The master key uses the configured encryption algorithm to encrypt sensitive data stored on the firewall and on Panorama. When you set the encryption algorithm to AES-256-GCM, you can still use an HSM to encrypt the master key with an encryption key that is stored on the HSM.

The default encryption algorithm that the master key uses to encrypt data is AES-256-CBC—the same algorithm that the master key used prior to PAN-OS 10.0. AES-256-CBC is the default encryption level because when you manage firewalls with Panorama, the managed firewalls may be on different PAN-OS releases, and firewalls on PAN-OS releases earlier than PAN-OS 10.1 do not support AES-256-GCM. This is why Panorama must use the lowest level of encryption that its managed devices can use. For example, if some managed devices run PAN-OS 10.1 and some run earlier versions, Panorama must use AES-256-CBC. However, if all managed devices run PAN-OS 10.1 or later, then Panorama and all of its managed devices can use AES-256-GCM.



Use the same encryption level on Panorama and its managed devices and use the same encryption level on firewall pairs. Upgrade devices to use the strongest possible encryption algorithm. If all Panorama-managed devices run PAN-OS 10.0, use AES-256-GCM on all devices. The configuration of managed or paired devices that use different encryption levels may become out of sync.

When you change the encryption algorithm to AES-256-GCM, devices use it instead of AES-256-CBC to encrypt sensitive data. When you change from one algorithm to another, you can also specify whether to:

- Re-encrypt existing encrypted data with the new algorithm.
- Leave existing data encrypted with the old encryption algorithm and use the new algorithm only for new (future) encryptions.



By default, when you change the encryption algorithm, the device uses the new algorithm to re-encrypt existing encrypted data as well as to encrypt new data. If you manage devices with Panorama, they may be on different versions of PAN-OS and may not support the newest encryption algorithms. Be sure you understand which encryption algorithms Panorama and its managed devices support before you change the encryption algorithm or re-encrypt data that has already been encrypted.

- Configure Master Key Encryption Level
- Master Key Encryption on a Firewall HA Pair
- Master Key Encryption Logs
- Unique Master Key Encryptions for AES-256-GCM

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Configure Master Key Encryption Level

You configure the master key encryption algorithm level and whether to re-encrypt all currently encrypted data with a new encryption algorithm level using the CLI. Depending on the order of the keywords, you can change the encryption level or you can change the encryption level and also specify whether to re-encrypt previously encrypted data.

The following operational CLI command changes the encryption level and automatically re-encrypts all currently encrypted data with the specified encryption level:

```
admin@PA-NGFW>request encryption-level level <0|1|2>
```

The following operational CLI command changes the encryption level and specifies whether to re-encrypt all currently encrypted data with the new encryption level:

```
admin@PA-NGFW>request encryption-level re-encrypt <yes|no> level <0|
1|2>
```

Keyword	Options
level	**0** = Use the default algorithm (AES-256-CBC) to encrypt data **1** = Use the AES-256-CBC algorithm to encrypt data **2** = Use the AES-256-GCM algorithm to encrypt data The firewall re-encrypts all currently encrypted data and encrypts new sensitive data using the specified algorithm. If you don't want to re-encrypt existing encrypted data with the new algorithm, specify **re-encrypt no** in the command string. This prevents the firewall from automatically re-encrypting data that the firewall has already encrypted. *Only use AES-256-GCM when Panorama and all of its managed devices (or both devices in an HA pair) run PAN-OS 10.1 or greater and configure all of the devices to use AES-256-GCM. Managed or paired devices that use different encryption levels may become out of sync.*
re-encrypt	**no** = Do not re-encrypt currently encrypted data. The firewall does not re-encrypt currently encrypted data. Currently encrypted data remains encrypted with whichever algorithm the firewall originally used to encrypt the data. The firewall uses the specified algorithm only to encrypt sensitive data in the future.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Keyword	Options
	yes = Re-encrypt currently encrypted data with the specified algorithm and use that algorithm to encrypt sensitive data in the future.

Use the operational CLI command **show system masterkey-properties** to verify the encryption algorithm (level) currently configured on the device, for example:

```
admin@PA-NGFW>show system masterkey-properties
```

```
Master key expires at: unspecified
Reminders will begin at: unspecified
Master key on hsm: no
Automatically renew master key lifetime: 0
Encryption Level: 1
```

The output shows that the current encryption level is 1, which is AES-256-CBC.

If you downgrade to an earlier version of PAN-OS, the device automatically reverts the encryption algorithm to a level that the downgraded PAN-OS version supports and automatically re-encrypts encrypted data using that level so that the device can decrypt and use the data as needed. For example, if your device is on PAN-OS 10.1 and uses AES-256-GCM as the encryption algorithm (which is not supported on earlier versions of PAN-OS), and you downgrade to PAN-OS 9.1, then the device re-encrypts the encrypted data to AES-256-CBC, which is supported in PAN-OS 9.1.

Master Key Encryption on a Firewall HA Pair

To use the AES-256-GCM encryption level on a firewall high availability (HA) pair, both firewalls must run PAN-OS 10.0 so that both firewalls support AES-256-GCM. If either firewall in the HA pair runs an earlier version than PAN-OS 10.0, you can't use AES-256-GCM. When both firewalls are on PAN-OS 10.0, both firewalls can decode AES-256-CBC or AES-256-GCM encryption keys, so they can use the either encryption level. However, both firewalls should use the same encryption level to avoid the possibility of becoming out of sync.

Use AES-256-GCM encryption on both firewalls in the HA pair. Whether you use AES-256-GCM or AES-256-CBC, use the same algorithm on both firewalls.

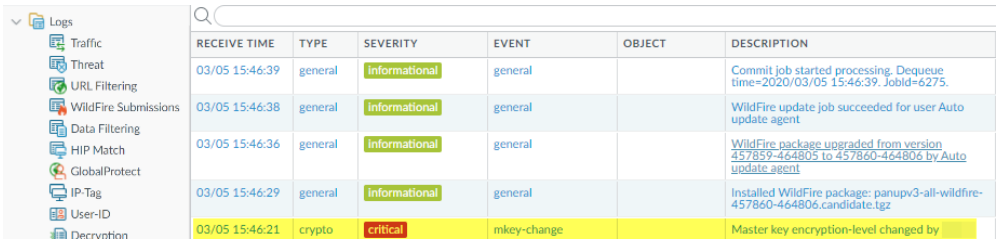
You do not need to disable HA to change the encryption level on a firewall in an HA pair in which both firewalls run PAN-OS 10.0.

Master Key Encryption Logs

The firewall generates a System Log (**Monitor** > **Logs** > **System**) when you change the master key encryption algorithm (level).

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets



To view all of the System Logs for master key encryption, create a filter that shows all logs of the **Type** crypto: **`(subtype eq crypto)`**.

Unique Master Key Encryptions for AES-256-GCM

The master key can only generate a finite number of unique encryptions before it runs out of unique combinations and must repeat encryptions. The firewall creates unique encryptions using the AES-256-GCM encryption algorithm with an Initialization Vector (IV). An IV is an arbitrary number that should only be used one time to create an encryption to ensure that each encryption is unique.

Each encryption using the master key and IV must be unique to prevent forgery attacks. The firewall meets the uniqueness requirement that the probability that the authenticated encryption is ever created with the same IV and the same key on two or more distinct sets of input data is no greater than 2^{32}.

When the IV runs through all of its unique values, the IV value repeats. When the IV value repeats, using the same master key and the repeated IV value to encrypt data means that the encryption is the same as an encryption previously used on other data. Change the Master Key before the system runs out of unique encryptions to prevent the firewall from using the same encryption (master key and IV value combination) on more than one piece of sensitive data. Unique encryption combinations should never be repeated or reused.

To track when you need to change the master key, set the master key **Lifetime** and **Reminder** values on each appliance (**Device** > **Master Key and Diagnostics** and edit the master key). Set the values conservatively, based on the expected volume of master key encryptions, to ensure that all encryptions are unique and no encryption combinations are repeated or reused.

https://t.me/learningnets

Obtain Certificates

- Create a Self-Signed Root CA Certificate
- Generate a Certificate
- Import a Certificate and Private Key
- Obtain a Certificate from an External CA
- Install a Device Certificate
- Deploy Certificates Using SCEP

Create a Self-Signed Root CA Certificate

A self-signed root certificate authority (CA) certificate is the top-most certificate in a certificate chain. A firewall can use this certificate to automatically issue certificates for other uses. For example, the firewall issues certificates for SSL/TLS decryption and for satellites in a GlobalProtect large-scale VPN.

When establishing a secure connection with the firewall, the remote client must trust the root CA that issued the certificate. Otherwise, the client browser will display a warning that the certificate is invalid and might (depending on security settings) block the connection. To prevent this, after generating the self-signed root CA certificate, import it into the client systems.

On a Palo Alto Networks firewall or Panorama, you can generate self-signed certificates only if they are CA certificates.

STEP 1 | Select **Device** > **Certificate Management** > **Certificates** > **Device Certificates**.

STEP 2 | If the firewall has more than one virtual system (vsys), select a **Location** (vsys or **Shared**) for the certificate.

STEP 3 | Click **Generate**.

STEP 4 | Enter a **Certificate Name**, such as **`GlobalProtect_CA`**. The name is case-sensitive and can have up to 63 characters on the firewall or up to 31 characters on Panorama. It must be unique and use only letters, numbers, hyphens, and underscores.

STEP 5 | In the **Common Name** field, enter the FQDN (recommended) or IP address of the interface where you will configure the service that will use this certificate.

STEP 6 | If the firewall has more than one vsys and you want the certificate to be available to every vsys, select the **Shared** check box.

STEP 7 | Leave the **Signed By** field blank to designate the certificate as self-signed.

STEP 8 | (Required) Select the **Certificate Authority** check box.

STEP 9 | Leave the **OCSP Responder** field blank; revocation status verification doesn't apply to root CA certificates.

STEP 10 | Click **Generate** and **Commit**.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Generate a Certificate

Palo Alto Networks firewalls and Panorama use certificates to authenticate clients, servers, users, and devices in several applications, including SSL/TLS decryption, Authentication Portal, GlobalProtect, site-to-site IPSec VPN, and web interface access to the firewall/Panorama. Generate certificates for each usage: for details, see Keys and Certificates.

To generate a certificate, you must first Create a Self-Signed Root CA Certificate or import one (Import a Certificate and Private Key) to sign it. To use Online Certificate Status Protocol (OCSP) for verifying certificate revocation status, Configure an OCSP Responder before generating the certificate.

STEP 1 | Select **Device** > **Certificate Management** > **Certificates** > **Device Certificates**.

STEP 2 | If the firewall has more than one virtual system (vsys), select a **Location** (vsys or **Shared**) for the certificate.

STEP 3 | Click **Generate**.

STEP 4 | Select **Local** (default) as the **Certificate Type** unless you want to deploy SCEP certificates to GlobalProtect endpoints.

STEP 5 | Enter a **Certificate Name**. The name is case-sensitive and can have up to 63 characters on the firewall or up to 31 characters on Panorama. It must be unique and use only letters, numbers, hyphens, and underscores.

STEP 6 | In the **Common Name** field, enter the FQDN (recommended) or IP address of the interface where you will configure the service that will use this certificate.

STEP 7 | If the firewall has more than one vsys and you want the certificate to be available to every vsys, select the **Shared** check box.

STEP 8 | In the **Signed By** field, select the root CA certificate that will issue the certificate.

STEP 9 | (Optional) Select an **OCSP Responder**.

STEP 10 | For the key generation **Algorithm**, select **RSA** (default) or **Elliptical Curve DSA** (ECDSA). ECDSA is recommended for client browsers and operating systems that support it.

> *Firewalls that run PAN-OS 6.1 and earlier releases will delete any ECDSA certificates that you push from Panorama™, and any RSA certificates signed by an ECDSA certificate authority (CA) will be invalid on those firewalls.*

You cannot use a hardware security module (HSM) to store ECDSA keys used for SSL/TLS Decryption.

STEP 11 | Select the **Number of Bits** to define the certificate key length. Higher numbers are more secure but require more processing time.

https://t.me/learningnets

STEP 12 | Select the **Digest** algorithm. From most to least secure, the options are: **sha512**, **sha384**, **sha256** (default), **sha1**, and **md5**.

*Client certificates that are used when requesting firewall services that rely on TLSv1.2 (such as administrator access to the web interface) cannot have **sha512** as a digest algorithm. The client certificates must use a lower digest algorithm (such as **sha384**) or you must limit the **Max Version** to **TLSv1.1** when you Configure an SSL/TLS Service Profile for the firewall services.*

STEP 13 | For the **Expiration**, enter the number of days (default is 365) for which the certificate is valid.

STEP 14 | (Optional) **Add** the **Certificate Attributes** to uniquely identify the firewall and the service that will use the certificate.

*If you add a **Host Name** (DNS name) attribute, it is a best practice for it to match the **Common Name**, because the host name populates the Subject Alternate Name (SAN) field of the certificate and some browsers require the SAN to specify the domains the certificate protects; in addition, the **Host Name** matching the **Common Name** is mandatory for GlobalProtect.*

STEP 15 | Click **Generate** and, in the Device Certificates page, click the certificate Name.

Regardless of the time zone on the firewall, it always displays the corresponding Greenwich Mean Time (GMT) for certificate validity and expiration dates/times.

STEP 16 | Select the check boxes that correspond to the intended use of the certificate on the firewall.

For example, if the firewall will use this certificate to secure forwarding of syslogs to an external syslog server, select the **Certificate for Secure Syslog** check box.

STEP 17 | Click **OK** and **Commit**.

Import a Certificate and Private Key

If your enterprise has its own public key infrastructure (PKI), you can import a certificate and private key into the firewall from your enterprise certificate authority (CA). Enterprise CA certificates (unlike most certificates purchased from a trusted, third-party CA) can automatically issue CA certificates for applications such as SSL/TLS decryption or large-scale VPN.

On a Palo Alto Networks firewall or Panorama, you can import self-signed certificates only if they are CA certificates.

Instead of importing a self-signed root CA certificate into all the client systems, it is a best practice to import a certificate from the enterprise CA because the clients will already have a trust relationship with the enterprise CA, which simplifies the deployment.

If the certificate you will import is part of a certificate chain, it is a best practice to import the entire chain.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 1 | From the enterprise CA, export the certificate and private key that the firewall will use for authentication.

When exporting a private key, you must enter a passphrase to encrypt the key for transport. Ensure the management system can access the certificate and key files. When importing the key onto the firewall, you must enter the same passphrase to decrypt it.

STEP 2 | Select **Device** > **Certificate Management** > **Certificates** > **Device Certificates**.

STEP 3 | If the firewall has more than one virtual system (vsys), select a **Location** (vsys or **Shared**) for the certificate.

STEP 4 | Click **Import** and enter a **Certificate Name**. The name is case-sensitive and can have up to 63 characters on the firewall or up to 31 characters on Panorama. It must be unique and use only letters, numbers, hyphens, and underscores.

STEP 5 | To make the certificate available to all virtual systems, select the **Shared** check box. This check box appears only if the firewall supports multiple virtual systems.

STEP 6 | Enter the path and name of the **Certificate File** received from the CA, or **Browse** to find the file.

STEP 7 | Select a **File Format**:

- **Encrypted Private Key and Certificate (PKCS12)**—This is the default and most common format, in which the key and certificate are in a single container (**Certificate File**). If a hardware security module (HSM) will store the private key for this certificate, select the **Private key resides on Hardware Security Module** check box.
- **Base64 Encoded Certificate (PEM)**—You must import the key separately from the certificate. If a hardware security module (HSM) stores the private key for this certificate, select the **Private key resides on Hardware Security Module** check box and skip the next step. Otherwise, select the **Import Private Key** check box, enter the **Key File** or **Browse** to it, then continue to the next step.

STEP 8 | Enter and re-enter (confirm) the **Passphrase** used to encrypt the private key.

STEP 9 | Click **OK**. The Device Certificates page displays the imported certificate.

Obtain a Certificate from an External CA

The advantage of obtaining a certificate from an external certificate authority (CA) is that the private key does not leave the firewall. To obtain a certificate from an external CA, generate a certificate signing request (CSR) and submit it to the CA. After the CA issues a certificate with the specified attributes, import it onto the firewall. The CA can be a well-known, public CA or an enterprise CA.

To use Online Certificate Status Protocol (OCSP) for verifying the revocation status of the certificate, Configure an OCSP Responder before generating the CSR.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 1 | Request the certificate from an external CA.

1. Select **Device** > **Certificate Management** > **Certificates** > **Device Certificates**.
2. If the firewall has more than one virtual system (vsys), select a **Location** (vsys or **Shared**) for the certificate.
3. Click **Generate**.
4. Enter a **Certificate Name**. The name is case-sensitive and can have up to 63 characters on the firewall or up to 31 characters on Panorama. It must be unique and use only letters, numbers, hyphens, and underscores.
5. In the **Common Name** field, enter the FQDN (recommended) or IP address of the interface where you will configure the service that will use this certificate.
6. If the firewall has more than one vsys and you want the certificate to be available to every vsys, select the **Shared** check box.
7. In the **Signed By** field, select **External Authority (CSR)**.
8. If applicable, select an **OCSP Responder**.
9. (Optional) **Add** the **Certificate Attributes** to uniquely identify the firewall and the service that will use the certificate.

   *If you add a **Host Name** attribute, it should match the **Common Name** (this is mandatory for GlobalProtect). The host name populates the Subject Alternative Name field of the certificate.*
10. Click **Generate**. The **Device Certificates** tab displays the CSR with a Status of `pending`.

STEP 2 | Submit the CSR to the CA.

1. Select the CSR and click **Export** to save the .csr file to a local computer.
2. Upload the .csr file to the CA.

STEP 3 | Import the certificate.

1. After the CA sends a signed certificate in response to the CSR, return to the **Device Certificates** tab and click **Import**.
2. Enter the **Certificate Name** used to generate the CSR.
3. Enter the path and name of the PEM **Certificate File** that the CA sent, or **Browse** to it.
4. Click **OK**. The **Device Certificates** tab displays the certificate with a Status of `valid`.

STEP 4 | Configure the certificate.

1. Click the certificate **Name**.
2. Select the check boxes that correspond to the intended use of the certificate on the firewall. For example, if the firewall will use this certificate to secure forwarding of syslogs to an external syslog server, select the **Certificate for Secure Syslog** check box.
3. Click **OK** and **Commit**.

Install a Device Certificate

Your next-generation firewall can leverage cloud services such as Device Telemetry and IoT. To do this, you must install a device certificate to successfully authenticate the firewall with the Palo Alto Networks Customer Support Portal (CSP) to leverage these cloud services. The circumstances

https://t.me/learningnets

under which a device certificate is required will differ from feature to feature, so install a device certificate only if the feature's setup documentation tells you that this needs to be done.

You only need to install a device certificate once. Every feature that uses device certificates will use the certificate installed on your firewall if it already exists.

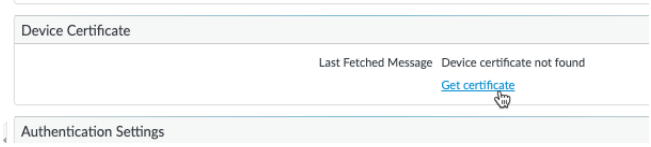
You can install a device certificate to firewalls that are managed by Panorama. If you want to install a device certificate directly to a single next-generation firewall (that is, you are *not* using Panorama):

STEP 1 | Generate the One Time Password (OTP).

1. Log in to the Customer Support Portal.
2. Select **Assets** > **Device Certificates** and **Generate OTP**.
3. For the **Device Type**, select **Generate OTP for Next-Gen Firewalls**.
4. Select your **PAN OS Device** serial number.
5. **Generate OTP** and copy the OTP.

STEP 2 | Log in to your next-generation firewall as an admin user.

STEP 3 | Select **Device** > **Setup** > **Management** > **Device Certificate** and **Get certificate**.



STEP 4 | Paste the **One-time Password** you generated and click **OK**.

STEP 5 | Your next-generation firewall successfully retrieves and installs the certificate.

Deploy Certificates Using SCEP

If you have a Simple Certificate Enrollment Protocol (SCEP) server in your enterprise PKI, you can configure a SCEP profile to automate the generation and distribution of unique client certificates. SCEP operation is dynamic in that the enterprise PKI generates a user-specific certificate when the SCEP client requests it and sends the certificate to the SCEP client. The SCEP client then transparently deploys the certificate to the client device.

You can use a SCEP profile with GlobalProtect to assign user-specific client certificates to each GlobalProtect user. In this use case, the GlobalProtect portal acts as a SCEP client to the SCEP server in your enterprise PKI. Additionally, you can use a SCEP profile to assign client certificates to Palo Alto Networks devices for mutual authentication with other Palo Alto Networks devices for management access and inter-device communication.

STEP 1 | Create a SCEP profile.

1. Select **Device** > **Certificate Management** > **SCEP** and then **Add** a new profile.
2. Enter a **Name** to identify the SCEP profile.
3. If this profile is for a firewall with multiple virtual systems capability, select a virtual system or **Shared** as the **Location** where the profile is available.

https://t.me/learningnets

STEP 2 | (**Optional**) To make the SCEP-based certificate generation more secure, configure a SCEP challenge-response mechanism between the PKI and portal for each certificate request.

After you configure this mechanism, its operation is invisible, and no further input from you is necessary.

To comply with the U.S. Federal Information Processing Standard (FIPS), use a **Dynamic** SCEP challenge and specify a **Server URL** that uses HTTPS.

Select one of the following options:

- **None**—(**Default**) The SCEP server does not challenge the portal before it issues a certificate.
- **Fixed**—Obtain the enrollment challenge password from the SCEP server in the PKI infrastructure and then enter the password into the Password field.
- **Dynamic**—Enter a username and password of your choice (possibly the credentials of the PKI administrator) and the SCEP **Server URL** where the portal-client submits these credentials. The uses the credentials to authenticate with the SCEP server which transparently generates an OTP password for the portal upon each certificate request. (You can see this OTP change after a screen refresh in `The enrollment challengepassword is` field after each certificate request.) The PKI transparently passes each new password to the portal, which then uses the password for its certificate request.

STEP 3 | Specify the settings for the connection between the SCEP server and the portal to enable the portal to request and receive client certificates.

You can include additional information about the client device or user by specifying tokens in the **Subject** name of the certificate.

The portal includes the token value and host ID in the CSR request to the SCEP server.

1. Configure the **Server URL** that the portal uses to reach the SCEP server in the PKI (for example, `http://10.200.101.1/certsrv/mscep/`).
2. Enter a string (up to 255 characters in length) in the **CA-IDENT Name** field to identify the SCEP server.
3. Enter the **Subject** name to use in the certificates generated by the SCEP server. The subject must be a distinguished name in the <**`attribute`**>=<**`value`**> format and must include a common name (CN) attribute (`CN=`<**`variable`**>). The CN supports the following dynamic tokens:
   - `$USERNAME`—Use this token to enable the portal to request certificates for a specific user. To use this variable with GlobalProtect, you must also Enable Group Mapping. The username entered by the user must match the name in the user-group mapping table.
   - `$EMAILADDRESS`—Use this token to request certificates associated with a specific email address. To use this variable, you must also Enable Group Mapping and configure the **Mail Attributes** in the Mail Domains section of the Server Profile. If GlobalProtect cannot identify an email address for the user, it generates a unique ID and populates the CN with that value.
   - `$HOSTID`—To request certificates for the device only, specify the host ID token. When a user attempts to log in to the portal, the endpoint sends identifying information that includes its host ID value. The host ID value varies by device type,

https://t.me/learningnets

either GUID (Windows) MAC address of the interface (Mac), Android ID (Android devices), UDID (iOS devices), or a unique name that GlobalProtect assigns (Chrome).

- `$UDID`—Use the UDID common name attribute to request certificates based on the client's device UDID for GlobalProtect or device serial number for mutual authentication between Palo Alto Networks devices.

When the GlobalProtect portal pushes the SCEP settings to the agent, the CN portion of the subject name is replaced with the actual value (username, host ID, or email address) of the certificate owner (for example, **`O=acme,CN=johndoe`**).

4. Select the **Subject Alternative Name Type**:

   *Use static entries for the Subject Alternative Name Type. The firewall does not support dynamic tokens such as **$USERNAME**.*

   - **RFC 822 Name**—Enter the email name in a certificate's subject or Subject Alternative Name extension.
   - **DNS Name**—Enter the DNS name used to evaluate certificates.
   - **Uniform Resource Identifier**—Enter the name of the resource from which the client will obtain the certificate.
   - **None**—Do not specify attributes for the certificate.

STEP 4 | (**Optional**) Configure cryptographic settings for the certificate.

- Select the key length (**Number of Bits**) for the certificate.

  If the firewall is in FIPS-CC mode and the key generation algorithm is RSA. The RSA keys must be 2,048 bits or larger.
- Select the **Digest for CSR** which indicates the digest algorithm for the certificate signing request (CSR): sha1, sha256, or sha384.

STEP 5 | (**Optional**) Configure the permitted uses of the certificate, either for signing or encryption.

- To use this certificate for signing, select the **Use as digital signature** check box. This enables the endpoint use the private key in the certificate to validate a digital signature.
- To use this certificate for encryption, select the **Use for key encipherment** check box. This enables the client use the private key in the certificate to encrypt data exchanged over the HTTPS connection established with the certificates issued by the SCEP server.

STEP 6 | (**Optional**) To ensure that the portal is connecting to the correct SCEP server, enter the **CA Certificate Fingerprint**. Obtain this fingerprint from the SCEP server interface in the Thumbprint field.

1. Enter the URL for the SCEP server's administrative UI (for example, **`http://<hostname or IP>/CertSrv/mscep_admin/`**).
2. Copy the thumbprint and enter it in the **CA Certificate Fingerprint** field.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 7 | Enable mutual SSL authentication between the SCEP server and the firewall. This is required to comply with the U.S. Federal Information Processing Standard (FIPS).



FIPS-CC operation is indicated on the firewall login page and in its status bar.

Select the SCEP server's root **CA Certificate**. Optionally, you can enable mutual SSL authentication between the SCEP server and the firewall by selecting a **Client Certificate**.

STEP 8 | Save and commit the configuration.

1. Click **OK** to save the settings and close the SCEP configuration.
2. **Commit** the configuration.

The portal attempts to request a CA certificate using the settings in the SCEP profile and saves it to the firewall hosting the portal. If successful, the CA certificate is shown in **Device** > **Certificate Management** > **Certificates**.

STEP 9 | (**Optional**) If after saving the SCEP profile, the portal fails to obtain the certificate, you can manually generate a certificate signing request (CSR) from the portal.

1. Select **Device** > **Certificate Management** > **Certificates** > **Device Certificates** and then click **Generate**.
2. Enter a **Certificate Name**. This name cannot contain spaces.
3. Select the **SCEP Profile** to use to submit a CSR to your enterprise PKI.
4. Click **OK** to submit the request and generate the certificate.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Export a Certificate and Private Key

Palo Alto Networks recommends that you use your enterprise public key infrastructure (PKI) to distribute a certificate and private key in your organization. However, if necessary, you can also export a certificate and private key from the firewall or Panorama. You can use an exported certificate and private key in the following cases:

- Configure Certificate-Based Administrator Authentication to the Web Interface
- Enable SSL Between GlobalProtect LSVPN Components to configure GlobalProtect agent/app authentication to portals and gateways
- SSL Forward Proxy decryption
- Obtain a Certificate from an External CA

STEP 1 | Select **Device** > **Certificate Management** > **Certificates** > **Device Certificates**.

STEP 2 | If the firewall has more than one virtual system (vsys), select a **Location** (a specific vsys or **Shared**) for the certificate.

STEP 3 | Select the certificate, click **Export**, and select a **File Format**:

- **Base64 Encoded Certificate (PEM)**—This is the default format. It is the most common and has the broadest support on the Internet. If you want the exported file to include the private key, select the **Export Private Key** check box.
- **Encrypted Private Key and Certificate (PKCS12)**—This format is more secure than PEM but is not as common or as broadly supported. The exported file will automatically include the private key.
- **Binary Encoded Certificate (DER)**—More operating system types support this format than the others. You can export only the certificate, not the key: ignore the **Export Private Key** check box and passphrase fields.

STEP 4 | Enter a **Passphrase** and **Confirm Passphrase** to encrypt the private key if the **File Format** is PKCS12 or if it is PEM and you selected the **Export Private Key** check box. You will use this passphrase when importing the certificate and key into client systems.

STEP 5 | Click **OK** and save the certificate/key file to your computer.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Configure a Certificate Profile

Certificate profiles define user and device authentication for Authentication Portal, multi-factor authentication (MFA), GlobalProtect, site-to-site IPSec VPN, external dynamic list (EDL) validation, Dynamic DNS (DDNS), User-ID agent and TS agent access, and web interface access to Palo Alto Networks firewalls or Panorama. The profiles specify which certificates to use, how to verify certificate revocation status, and how that status constrains access. Configure a certificate profile for each application.

It is a best practice to enable Online Certificate Status Protocol (OCSP) and Certificate Revocation List (CRL) status verification for certificate profiles to verify that the certificate hasn't been revoked. Enable both OCSP and CRL so that if the OCSP server isn't available, the firewall uses CRL. For details on these methods, see Certificate Revocation.

STEP 1 | Obtain the certificate authority (CA) certificates you will assign.

Perform one of the following steps to obtain the CA certificates you will assign to the profile. You must assign at least one.

- Generate a Certificate.
- Export a certificate from your enterprise CA and then import it onto the firewall (see step to 3).

STEP 2 | Identify the certificate profile.

1. Select **Device** > **Certificate Management** > **Certificate Profile** and click **Add**.
2. Enter a **Name** to identify the profile. The name is case-sensitive, must be unique and can use up to 63 characters on the firewall or up to 31 characters on Panorama that include only letters, numbers, spaces, hyphens, and underscores.
3. If the firewall has more than one virtual system (vsys), select a **Location** (vsys or **Shared**) for the certificate.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 3 | Assign one or more certificates.

Perform the following steps for each CA certificate:

1. In the CA Certificates table, click **Add**.
2. Select a **CA Certificate**. Alternatively, to import a certificate, click **Import**, enter a **Certificate Name**, **Browse** to the **Certificate File** you exported from your enterprise CA, and click **OK**.
3. (**Optional**) If the firewall uses OCSP to verify certificate revocation status, configure the following fields to override the default behavior. For most deployments, these fields do not apply.
   - By default, the firewall uses the "Authority Information Access" (AIA) information from the certificate to extract the OCSP responder information. To override the AIA information, enter a **Default OCSP URL** (starting with **`http://`** or **`https://`**).
   - By default, the firewall uses the certificate selected in the **CA Certificate** field to validate OCSP responses. To use a different certificate for validation, select it in the **OCSP Verify CA Certificate** field.
4. Click **OK**. The CA Certificates table displays the assigned certificate.

STEP 4 | Define the methods for verifying certificate revocation status and the associated blocking behavior.

1. Select **Use CRL** and/or **Use OCSP**. If you select both, the firewall first tries OCSP and falls back to the CRL method only if the OCSP responder is unavailable.
2. Depending on the verification method, enter the **CRL Receive Timeout** and/or **OCSP Receive Timeout**. These are the intervals (1-60 seconds) after which the firewall stops waiting for a response from the CRL/OCSP service.
3. Enter the **Certificate Status Timeout**. This is the interval (1-60 seconds) after which the firewall stops waiting for a response from any certificate status service and applies any

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

session-blocking logic you define. The **Certificate Status Timeout** relates to the OCSP/CRL **Receive Timeout** as follows:

- If you enable both OCSP and CRL—The firewall registers a request timeout after the lesser of two intervals passes: the **Certificate Status Timeout** value or the aggregate of the two **Receive Timeout** values.
- If you enable only OCSP—The firewall registers a request timeout after the lesser of two intervals passes: the **Certificate Status Timeout** value or the OCSP **Receive Timeout** value.
- If you enable only CRL—The firewall registers a request timeout after the lesser of two intervals passes: the **Certificate Status Timeout** value or the CRL **Receive Timeout** value.

4. If you want the firewall to block sessions when the OCSP or CRL service returns a certificate revocation status of unknown, select **Block session if certificate status is unknown**. Otherwise, the firewall allows the sessions.
5. If you want the firewall to block sessions after it registers an OCSP or CRL request timeout, select **Block session if certificate status cannot be retrieved within timeout**. Otherwise, the firewall allows the sessions.
6. (**GlobalProtect only**) If you want the firewall to block sessions when the serial number attribute in the subject of the client certificate does not match the host ID that the GlobalProtect app reports for the endpoint, select **Block sessions if the certificate was not issued to the authenticating device**.

STEP 5 | Click **OK** and **Commit**

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Configure an SSL/TLS Service Profile

Palo Alto Networks firewalls and Panorama use SSL/TLS service profiles to specify a certificate and the allowed protocol versions for SSL/TLS services. The firewall and Panorama use SSL/TLS for Authentication Portal, GlobalProtect portals and gateways, inbound traffic on the management (MGT) interface, the URL Admin Override feature, and the User-ID™ syslog listening service. By defining the protocol versions, you can use a profile to restrict the cipher suites that are available for securing communication with the clients requesting the services. This improves network security by enabling the firewall or Panorama to avoid SSL/TLS versions that have known weaknesses. If a service request involves a protocol version that is outside the specified range, the firewall or Panorama downgrades or upgrades the connection to a supported version.

In the client systems that request firewall services, the certificate trust list (CTL) must include the certificate authority (CA) certificate that issued the certificate specified in the SSL/TLS service profile. Otherwise, users will see a certificate error when requesting firewall services. Most third-party CA certificates are present by default in client browsers. If an enterprise or firewall-generated CA certificate is the issuer, you must deploy that CA certificate to the CTL in client browsers.

STEP 1 | For each desired service, generate or import a certificate on the firewall (see Obtain Certificates).

Use only signed certificates, not CA certificates, in SSL/TLS service profiles.

STEP 2 | Select **Device** > **Certificate Management** > **SSL/TLS Service Profile**.

STEP 3 | If the firewall has more than one virtual system (vsys), select the **Location** (vsys or **Shared**) where the profile is available.

STEP 4 | Click **Add** and enter a **Name** to identify the profile.

STEP 5 | Select the **Certificate** you just obtained.

STEP 6 | Define the range of protocols that the service can use:

- For the **Min Version**, select the earliest allowed TLS version: **TLSv1.0** (default), **TLSv1.1**, or **TLSv1.2**.
- For the **Max Version**, select the latest allowed TLS version: **TLSv1.0**, **TLSv1.1**, **TLSv1.2**, or **Max** (latest available version). The default is **Max**.

*As a best practice, set the **Min Version** to **TLSv1.2** and the **Max Version** to **Max**.*

*On firewalls in FIPS/CC mode running PAN-OS 8.0 or a later release, **TLSv1.1** is the earliest supported TLS version; do not select **TLSv1.0**.*

*Client certificates that are used when requesting firewall services that rely on TLSv1.2 cannot have SHA512 as a digest algorithm. The client certificates must use a lower digest algorithm (such as SHA384) or you must limit the **Max Version** to **TLSv1.1** for the firewall services.*

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 7 | Click **OK** and **Commit**.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Configure an SSH Service Profile

SSH service profiles enable you to customize SSH parameters to enhance the security and integrity of SSH connections to your Palo Alto Networks management and high availability (HA) appliances. By default, SSH supports all ciphers, key exchange algorithms, and message authentication codes, which leaves your connection vulnerable to attack. With an SSH service profile, you can restrict the algorithms your SSH server supports. You can also generate a new host key and specify data volume, time, and packet-based thresholds for SSH session key regeneration and exchange.

Depending on the SSH server instance, configure either a management or HA SSH service profile. You can configure profiles from the firewall or Panorama™ web interface (if applying settings across multiple firewalls or appliances) or the CLI.

You can configure a maximum of four management and four HA server profiles.

*To use the same SSH connection settings for each Dedicated Log Collector (M-series or Panorama virtual appliance in Log Collector mode) in a Collector Group, configure an SSH service profile from the Panorama management server, **Commit** the changes to Panorama, and then **Push** the configuration to the Log Collectors. You can also perform these steps from the CLI using **`set log-collector-group <name> general-setting management ssh`** commands.*

- Create an SSH Management Profile
- Create an SSH HA Profile

Create an SSH Management Profile

You must create an SSH management profile to customize SSH settings for management connections.

You can configure or update an existing management profile from your CLI.

 ©2022 Palo Alto Networks, Inc.

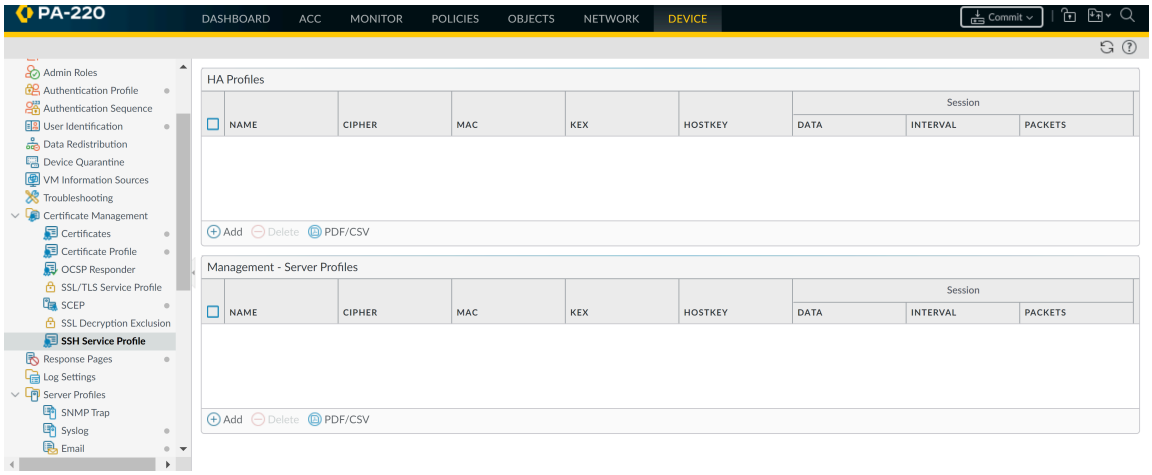
https://t.me/learningnets

STEP 1 | Create a Management - Server Profile.

1. Select **Device** > **Certification Management** > **SSH Service Profile**.
2. **Add** a Management - Server Profile.

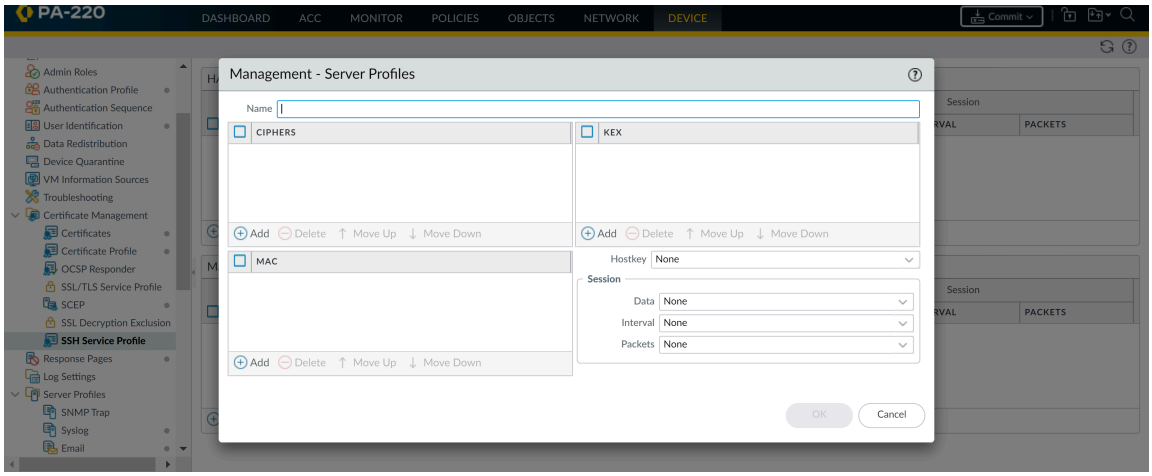
 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets



 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

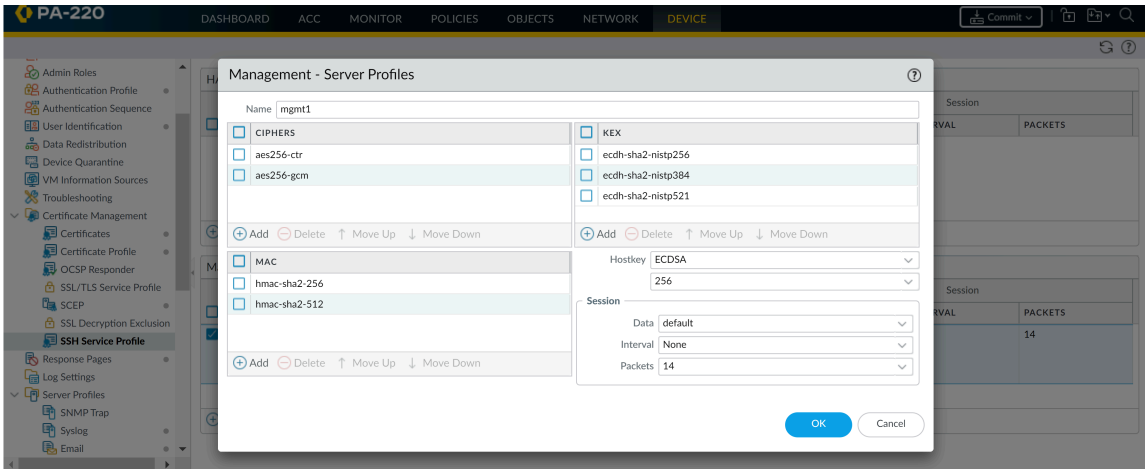


https://t.me/learningnets

3. Enter a **Name** to identify the profile.
4. (Optional) **Add** the ciphers, message authentication codes, or key exchange algorithms the profile will support.
5. (Optional) Select a **Hostkey** and key length.
6. (Optional) Enter values for the SSH session rekey parameters: **Data**, **Interval**, and **Packets**.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets



 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

7. Click **OK** and **Commit**.

 ©2022 Palo Alto Networks, Inc.

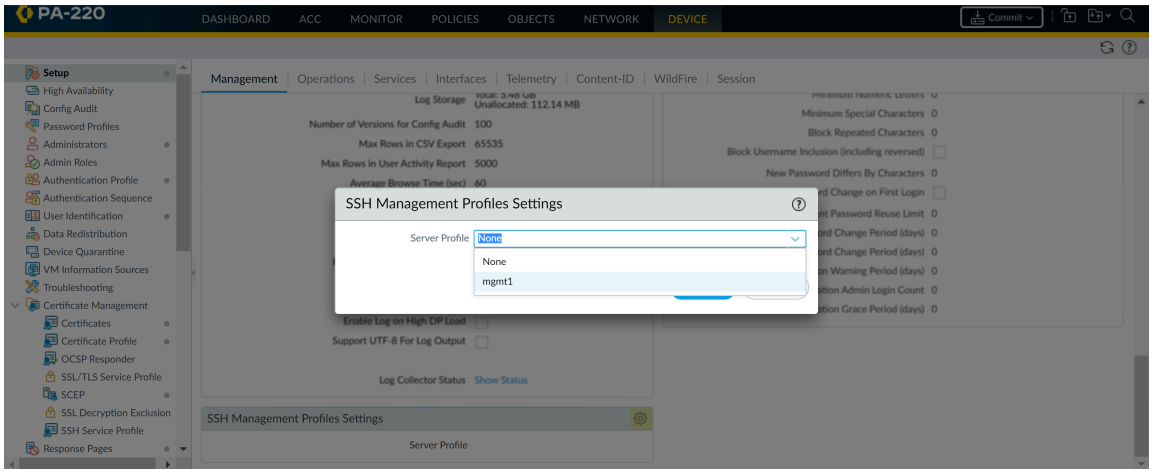
https://t.me/learningnets

STEP 2 | Select a management profile to apply.

1. Select **Device** > **Setup** > **Management**. Under SSH Management Profiles Settings, select an existing profile.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets



https://t.me/learningnets

2. Click **OK** and **Commit** the changes.

STEP 3 | Restart management SSH service from the CLI to apply the profile.

You must restart the connection each time you apply a new profile or make changes to a profile in use; this reboots the appliance. The new configurations will not affect active sessions. The profile will apply to subsequent connections (or sessions).

1. admin@PA-3260> **`set ssh service-restart mgmt`**

Create an SSH HA Profile

To secure SSH communications between appliances in an HA pair, you should create an SSH HA profile. Before you can create a profile, you must establish an HA connection between the appliances. If an HA connection has not been established, you must enable encryption on the control link connection, export the HA key to a network location, and import the HA key on the peer. (See Configure Active/Passive HA or Configure Active/Active HA.)

You can configure or update an existing HA profile from your CLI.

 ©2022 Palo Alto Networks, Inc.

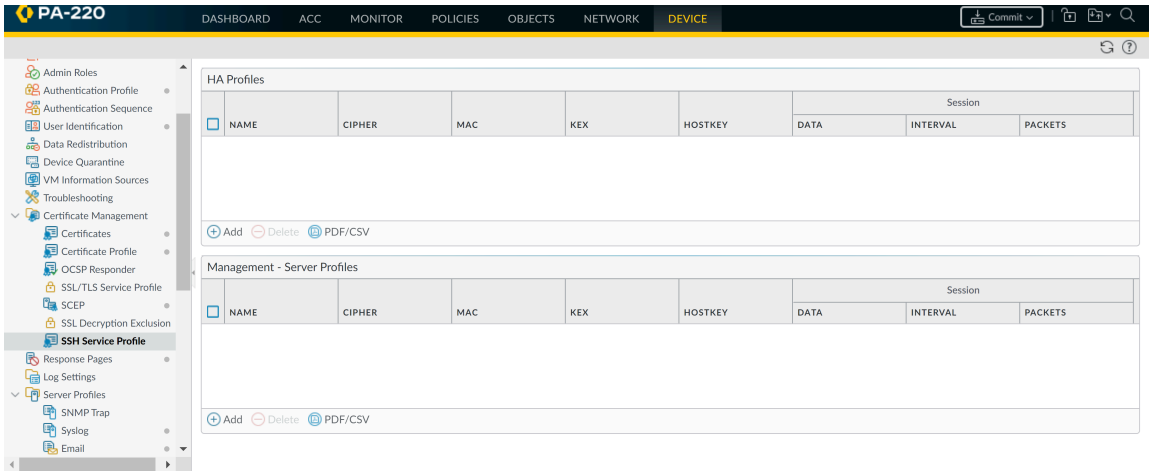
https://t.me/learningnets

STEP 1 | Create an HA Profile.

1. Select **Device** > **Certification Management** > **SSH Service Profile**.
2. **Add** an HA Profile.

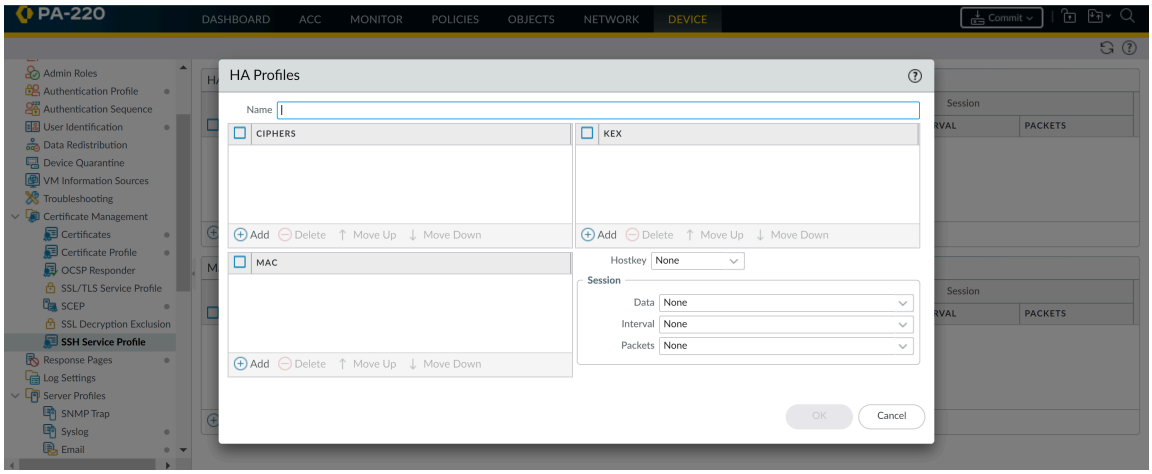
 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets



 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

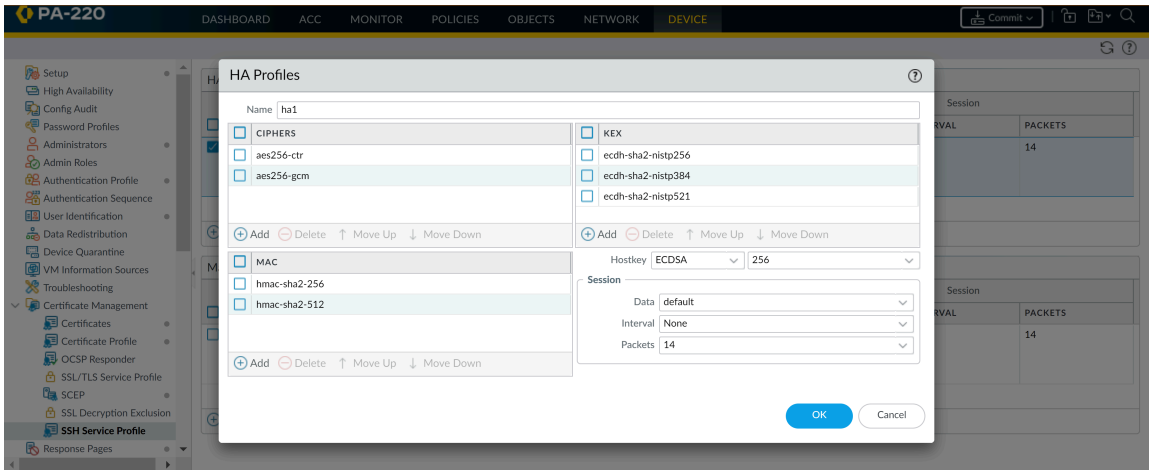


https://t.me/learningnets

3. Enter a **Name** to identify the profile.
4. (Optional) **Add** the ciphers, message authentication codes, or key exchange algorithms the profile will support.
5. (Optional) Select a **Hostkey** and key length.
6. (Optional) Enter values for the SSH session rekey parameters: **Data**, **Interval**, and **Packets**.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets



https://t.me/learningnets

7. Click **OK** and **Commit**.

 ©2022 Palo Alto Networks, Inc.

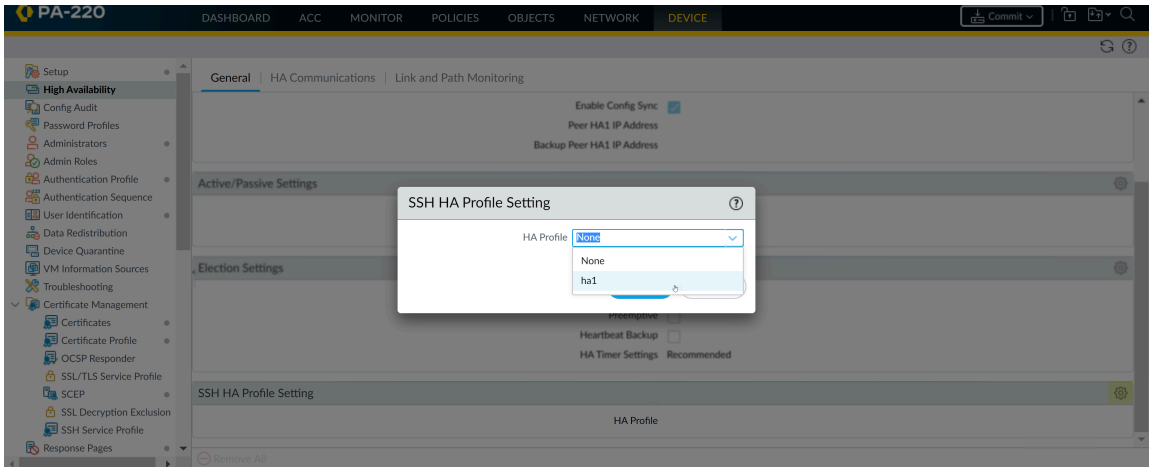
https://t.me/learningnets

STEP 2 | Select an HA Profile to apply.

1. Select **Device** > **High Availability** > **General**. Under SSH HA Profile Setting, select an existing profile.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets



 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

2. Click **OK** and **Commit** the changes.

STEP 3 | Restart HA1 SSH service from the CLI to apply the profile.

You must restart the connection each time you apply a new profile or make changes to a profile in use; this reboots the appliance. The new configuration will not affect active sessions. The profile will apply to subsequent connections (or sessions).

1. admin@PA-3260> **set ssh service-restart ha**



You can use the following commands if connection between the HA pair has been established and you'd like to minimize the downtime that accompanies an SSH service restart. If no HA connection has been established, you must restart SSH service.

- *(HA1 Backup is configured) admin@PA-3260>* ***request high-availability session-reestablish***
- *(No HA1 Backup is configured or HA1 Backup link is down) admin@PA-3260>* ***request high-availability session-reestablish force***

  You can force the firewall to reestablish HA1 sessions if there is no HA1 backup, which causes a brief split-brain condition between the HA peers. (Using the ***force*** *option when an HA1 backup is configured has no effect.)*

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Replace the Certificate for Inbound Management Traffic

When you first boot up the firewall or Panorama, it automatically generates a default certificate that enables HTTPS access to the web interface and XML API over the management (MGT) interface and (on the firewall only) over any other interface that supports HTTPS management traffic (for details, see Use Interface Management Profiles to Restrict Access). To improve the security of inbound management traffic, replace the default certificate with a new certificate issued specifically for your organization.

You cannot view, modify, or delete the default certificate.

To secure management traffic, you must also Configure Administrative Accounts and Authentication.

STEP 1 | Obtain the certificate that will authenticate the firewall or Panorama to the client systems of administrators.

You can simplify your Certificate Deployment by using a certificate that the client systems already trust. Therefore, we recommend that you Import a Certificate and Private Key from your enterprise certificate authority (CA) or Obtain a Certificate from an External CA; the trusted root certificate store of the client systems is likely to already have the associated root CA certificate that ensures trust.

If you Generate a Certificate *on the firewall or Panorama, administrators will see a certificate error because the root CA certificate is not in the trusted root certificate store of client systems. To prevent this, deploy the self-signed root CA certificate to all client systems.*

Regardless of how you obtain the certificate, we recommend a ***Digest*** *algorithm of* ***sha256*** *or higher for enhanced security.*

STEP 2 | Configure an SSL/TLS Service Profile.

Select the **Certificate** you just obtained.

For enhanced security, we recommend that you set the ***Min Version*** *(earliest allowed TLS version) to* ***TLSv1.2*** *for inbound management traffic. We also recommend that you use a different SSL/TLS Service Profile for each firewall or Panorama service instead of reusing this profile for all services.*

STEP 3 | Apply the SSL/TLS Service Profile to inbound management traffic.

1. Select **Device** > **Setup** > **Management** and edit the General Settings.
2. Select the **SSL/TLS Service Profile** you just configured.
3. Click **OK** and **Commit**.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Configure the Key Size for SSL Forward Proxy Server Certificates

When responding to a client in an SSL Forward Proxy session, the firewall creates a copy of the certificate that the destination server presents and uses the copy to establish a connection with the client. By default, the firewall generates certificates with the same key size as the certificate that the destination server presented. However, you can change the key size for the firewall-generated certificate as follows:

STEP 1 | Select **Device** > **Setup** > **Session** and, in the Decryption Settings section, click **SSL Forward Proxy Settings**.

STEP 2 | Select a **Key Size**:

- **Defined by destination host**—The firewall determines the key size and the hashing algorithm for the certificates it generates to establish SSL proxy sessions with clients based on the destination server certificate. If the destination server uses a 1,024-bit RSA key, the firewall generates a certificate with a 1,024-bit RSA key. If the destination server uses a key size larger than 1,024 bits (for example, 2,048 bits or 4,096 bits), the firewall generates a certificate that uses a 2,048-bit RSA key. If the destination server uses the SHA-1 hashing algorithm, the firewall generates a certificate with the SHA-1 hashing algorithm. If the destination server uses a hashing algorithm stronger than SHA-1, the firewall generates a certificate with the SHA-256 algorithm. This is the default setting.
- **1024-bit RSA**—The firewall generates certificates that use a 1,024-bit RSA key and SHA-256 hashing algorithm regardless of the key size of the destination server certificates. As of December 31, 2013, public certificate authorities (CAs) and popular browsers have limited support for X.509 certificates that use keys of fewer than 2,048 bits. In the future, depending on security settings, when presented with such keys the browser might warn the user or block the SSL/TLS session entirely.
- **2048-bit RSA**—The firewall generates certificates that use a 2,048-bit RSA key and SHA-256 hashing algorithm regardless of the key size of the destination server certificates. Public CAs and popular browsers support 2,048-bit keys, which provide better security than the 1,024-bit keys.

Changing the key size setting clears the current certificate cache.

STEP 3 | Click **OK** and **Commit**.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Revoke and Renew Certificates

- Revoke a Certificate
- Renew a Certificate

Revoke a Certificate

Various circumstances can invalidate a certificate before the expiration date. Some examples are a change of name, change of association between subject and certificate authority (for example, an employee terminates employment), and compromise (known or suspected) of the private key. Under such circumstances, the certificate authority (CA) that issued the certificate must revoke it. The following task describes how to revoke a certificate for which the firewall is the CA.

STEP 1 | Select **Device** > **Certificate Management** > **Certificates** > **Device Certificates**.

STEP 2 | If the firewall supports multiple virtual systems, the tab displays a **Location** drop-down. Select the virtual system to which the certificate belongs.

STEP 3 | Select the certificate to revoke.

STEP 4 | Click **Revoke**. PAN-OS immediately sets the status of the certificate to revoked and adds the serial number to the Online Certificate Status Protocol (OCSP) responder cache or certificate revocation list (CRL). You need not perform a commit.

Renew a Certificate

If a certificate expires, or soon will, you can reset the validity period. If an external certificate authority (CA) signed the certificate and the firewall uses the Online Certificate Status Protocol (OCSP) to verify certificate revocation status, the firewall uses the OCSP responder information to update the certificate status (see Configure an OCSP Responder). If the firewall is the CA that issued the certificate, the firewall replaces it with a new certificate that has a different serial number but the same attributes as the old certificate.

STEP 1 | Select **Device** > **Certificate Management** > **Certificates** > **Device Certificates**.

STEP 2 | If the firewall has more than one virtual system (vsys), select a **Location** (vsys or **Shared**) for the certificate.

STEP 3 | Select a certificate to renew and click **Renew**.

STEP 4 | Enter a **New Expiration Interval** (in days).

STEP 5 | Click **OK** and **Commit**.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Secure Keys with a Hardware Security Module

A hardware security module (HSM) is a physical device that manages digital keys. An HSM provides secure storage and generation of digital keys. It provides both logical and physical protection of these materials from non-authorized use and potential adversaries.

HSM clients integrated with Palo Alto Networks firewalls and Panorama enable enhanced security for the private keys used in SSL/TLS decryption (both SSL forward proxy and SSL inbound inspection). In addition, you can use the HSM to encrypt master keys.

The following topics describe how to integrate an HSM with your firewall or Panorama:

- Set Up Connectivity with an HSM
- Encrypt a Master Key Using an HSM
- Store Private Keys on an HSM
- Manage the HSM Deployment

Set Up Connectivity with an HSM

HSM clients are integrated with PA-3200 Series, PA-5200 Series, PA-7000 Series, and VM-Series firewalls and with the Panorama management server (both virtual and M-Series appliances) for use with the following HSM vendors:

- **nCipher nShield Connect**—The supported client versions depend on the PAN-OS release:
  - PAN-OS 10.1 supports client version 12.40.2 (backward compatible up to client version 11.50 for older appliances).
  - PAN-OS 9.1, 9.0, and 8.1 support client version 12.30.
  - PAN-OS 8.0 and earlier releases support client version 11.62.
- **SafeNet Network**—The supported client versions depend on the PAN-OS release:
  - PAN-OS 10.1 supports client versions 5.4.2 and 7.2.
  - PAN-OS 9.1 and 9.0 support client versions 5.4.2 and 6.3.
  - PAN-OS 8.1 supports client versions 5.4.2 and 6.2.2.
  - PAN-OS 8.0.2 and later PAN-OS 8.0 releases (also PAN-OS 7.1.10 and later PAN-OS 7.1 releases) support client versions 5.2.1, 5.4.2, and 6.2.2.

The HSM server version must be compatible with these client versions. Refer to the HSM vendor documentation for the client-server version compatibility matrix. On the firewall or Panorama, use the following procedure to select the SafeNet Network client version that is compatible with your SafeNet HSM server.

Downgrading HSM servers might not be an option after you upgrade them.

- Set Up Connectivity with a SafeNet Network HSM
- Set Up Connectivity with an nCipher nShield Connect HSM

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

- Install the SafeNet Client RPM Packet Manager.
  1. Select **Device** > **Setup** > **HSM** and **Select HSM Client Version** (Hardware Security Operations settings).
  2. Select **Version 5.4.2** (default) or **7.2** as appropriate for your HSM server version.
  3. Click **OK**.
  4. (Required only if you change the HSM version on the firewall) If the version change succeeds, the firewall prompts you to reboot to change to the new HSM version. If prompted, click **Yes**.
  5. If the master key isn't on the firewall, the client version upgrade will fail. **Close** the message and make the master key local to the firewall:
     - Edit the Hardware Security Module Provider and disable (clear) the **Master Key Secured by HSM** option.
     - Click **OK**.
     - Select **Device** > **Master Key and Diagnostics** to edit the Master Key.
     - Enter the **Current Master Key**; you can then enter that same key to be the **New Master Key** and then **Confirm New Master Key**.
     - Click **OK**.
     - Repeat the first four steps to **Select HSM Client Version** and reboot again.

Set Up Connectivity with a SafeNet Network HSM

To set up connectivity between the Palo Alto Networks firewall (HSM client) and a SafeNet Network HSM server, you must specify the IP address of the server, enter a password for authenticating the firewall to the server, and then register the firewall with the server. Before you being configuring your HSM client, create a partition for the firewall on the HSM server and then confirm that the SafeNet Network client version on the firewall is compatible with your SafeNet Network HSM server (see Set Up Connectivity with an HSM).

Before the HSM and firewall connect, the HSM authenticates the firewall based on the firewall IP address. Therefore, you must configure the firewall to use a static IP address—not a dynamic address assigned through DHCP. Operations on the HSM stop working if the firewall IP address changes during runtime.

HSM configurations are not synchronized between high availability (HA) firewall peers. Consequently, you must configure the HSM separately on each peer. In active/passive HA configurations, you must manually perform one failover to individually configure and authenticate each HA peer to the HSM. After this initial manual failover, user interaction is not required for failover to function properly.

STEP 1 | Define connection settings for each SafeNet Network HSM.

1. Log in to the firewall web interface and select **Device** > **Setup** > **HSM**.
2. Edit the Hardware Security Module Provider settings and set the **Provider Configured** to **SafeNet Network HSM**.
3. **Add** each HSM server as follows. A high availability (HA) HSM configuration requires at least two servers; you can have a cluster of up to 16 HSM servers. All HSM servers in the cluster must run the same SafeNet version and must authenticate separately. You

https://t.me/learningnets

should use a SafeNet cluster only when you want to replicate the keys across the cluster. Alternatively, you can add up to 16 SafeNet HSM servers to function independently.

1. Enter a **Module Name** (an ASCII string of up to 31 characters) for the HSM server.
2. Enter an IPv4 address for the HSM **Server Address**.

4. (HA only) Select **High Availability**, specify the **Auto Recovery Retry** value (maximum number of times the HSM client tries to recover its connection to an HSM server before failing over to an HSM HA peer server; range is 0 to 500; default is 0), and enter a **High Availability Group Name** (an ASCII string up to 31 characters long).

   *If you configure two or more HSM servers, the best practice is to enable **High Availability**. Otherwise the firewall does not use the additional HSM servers.*

5. Click **OK** and **Commit** your changes.

STEP 2 | (Optional) Configure a service route to connect to the HSM if you don't want the firewall to connect through the Management interface (default).

*If you configure a service route for the HSM, running the **`clear session all`** CLI command clears all existing HSM sessions, which brings all HSM states down and then up again. During the several seconds required for HSM to recover, all SSL/TLS operations will fail.*

1. Select **Device** > **Setup** > **Services** and click **Service Route Configuration**.
2. **Customize** a service route. The **IPv4** tab is active by default.
3. Click **HSM** in the Service column.
4. Select a **Source Interface** for the HSM.
5. Click **OK** and **Commit** your changes.

STEP 3 | Configure the firewall to authenticate to the HSM.

1. Select **Device** > **Setup** and **Setup Hardware Security Module**.
2. Select the HSM **Server Name**.
3. Select **Automatic** or **Manual** for your authentication and trust certificate.
4. Enter the **Administrator Password** to authenticate the firewall to the HSM.
5. Click **OK**.

   The firewall tries to authenticate to the HSM and displays a status message.

6. Click **OK** again.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 4 | Register the firewall as an HSM client with the HSM server and assign the firewall to a partition on the HSM server.

If the HSM has a firewall with the same <cl-name> already registered, you must first remove the duplicate registration by running the ***client delete -client <cl-name>*** *command, where <cl-name> is the name of the registered client (firewall) you want to delete.*

1. Log in to the HSM from a remote system.
2. Register the firewall using the **client register -c <*cl-name*> -ip <*fw-ip-addr*>** CLI command, where *<cl-name>* is a name that you assign to the firewall for use on the HSM and *<fw-ip-addr>* is the IP address for that firewall.
3. Assign a partition to the firewall using the **client assignpartition -c <*cl-name*> -p <*partition-name*>** CLI command, where *<cl-name>* is the name you assigned to the firewall using the **client register** command and *<partition-name>* is the name of a previously configured partition that you want to assign to this firewall.

STEP 5 | Configure the firewall to connect to the HSM partition.

1. Select **Device** > **Setup** > **HSM** and refresh ( ) the display.
2. **Setup HSM Partition** (Hardware Security Operations settings).
3. Enter the **Partition Password** to authenticate the firewall to the partition on the HSM.
4. Click **OK**.

STEP 6 | (HA only) Repeat the previous authentication, registration, and partition connection steps to add another HSM to the existing HA group.

If you remove an HSM from your configuration, repeat the previous partition connection step to remove the deleted HSM from the HA group.

STEP 7 | Verify firewall connectivity and authentication with the HSM.

1. Select **Device** > **Setup** > **HSM** and check the authentication and connection Status:
   - **Green**—The firewall is successfully authenticated and connected to the HSM.
   - **Red**—The firewall failed to authenticate to the HSM or network connectivity to the HSM is down.
2. View the following columns in Hardware Security Module Status to determine the authentication status:
   - **Serial Number**—The serial number of the HSM partition if the firewall successfully authenticated to the HSM.
   - **Partition**—The partition name on the HSM that is assigned to the firewall.
   - **Module State**—The current state of the HSM connection. This value is always `Authenticated` if the Hardware Security Module Status displays the HSM.

Set Up Connectivity with an nCipher nShield Connect HSM

You must set up a remote file system (RFS) as a hub to synchronize key data for all firewalls (HSM clients) in your organization that use the nCipher nShield Connect HSM. To ensure the nShield

https://t.me/learningnets

Connect client version on your firewalls is compatible with your nShield Connect server, see Set Up Connectivity with an HSM.

Before the HSM and firewalls connect, the HSM authenticates the firewalls based on their IP addresses. Therefore, you must configure the firewalls to use static IP addresses—not dynamic addresses assigned through DHCP. (Operations on the HSM stop working if a firewall IP address changes during runtime).

HSM configurations are not synchronized between high availability (HA) firewall peers. Consequently, you must configure the HSM separately on each peer. In active/passive HA configurations, you must manually perform one failover to individually configure and authenticate each HA peer to the HSM. After this initial manual failover, user interaction is not required for failover to function properly.

STEP 1 | Define connection settings for each nCipher nShield Connect HSM.

1. Log in to the firewall web interface and select **Device** > **Setup** > **HSM**.
2. Edit the Hardware Security Module Provider settings and set the **Provider Configured** to **nShield Connect**.
3. **Add** each HSM server as follows. An HA HSM configuration requires two servers.
    1. Enter a **Module Name** for the HSM server. This can be any ASCII string of up to 31 characters.
    2. Enter an IPv4 address for the HSM **Server Address**.
4. Enter an IPv4 address for the **Remote Filesystem Address**.
5. Click **OK** and **Commit** your changes.

STEP 2 | (Optional) Configure a service route to connect to the HSM if you don't want the firewall to connect through the Management interface (default).

If you configure a service route for the HSM, running the ***`clear session all`*** *CLI command clears all existing HSM sessions, which brings all HSM states down and then up again. During the several seconds required for HSM to recover, all SSL/TLS operations will fail.*

1. Select **Device** > **Setup** > **Services** and click **Service Route Configuration**.
2. **Customize** a service route. The **IPv4** tab is active by default.
3. Click **HSM** in the Service column.
4. Select a **Source Interface** for the HSM.
5. Click **OK** and **Commit** your changes.

https://t.me/learningnets

STEP 3 | Register the firewall as an HSM client with the HSM server.

This step briefly describes the procedure for using the front panel interface of the nShield Connect HSM. For more details, refer to nCipher documentation.

1. Log in to the front panel display of the nCipher nShield Connect HSM.
2. Use the right-hand navigation button to select **System** > **System configuration** > **Client config** > **New client**.
3. Enter the firewall IP address.
4. Select **System** > **System configuration** > **Client config** > **Remote file system** and enter the IP address of the client computer where you set up the RFS.

STEP 4 | Configure the RFS to accept connections from the firewall.

1. Log in to the RFS from a Linux client.
2. Obtain the electronic serial number (ESN) and the hash of the K_{NETI} key, which authenticates the HSM to clients, by running the **anonkneti *<ip-address>*** CLI command, where *<ip-address>* is the HSM IP address.

   For example:

   ```
   anonkneti 192.0.2.1
   ```

   ```
   B1E2-2D4C-E6A2 5a2e5107e70d525615a903f6391ad72b1c03352c
   ```

   In this example, `B1E2-2D4C-E6A2` is the ESN and `5a2e5107e70d525615a903f6391ad72b1c03352c` is the hash of the K_{NETI} key.
3. Use the following command from a superuser account to set up the RFS:

   ```
   rfs-setup --force <ip-address> <ESN> <hash-Kneti-key>
   ```

   The *<ip-address>* is the IP address of the HSM, *<ESN>* is the electronic serial number, and *<hash-Kneti-key>* is the hash of the K_{NETI} key.

   The following example uses the values obtained in this procedure:

   ```
   rfs-setup --force 192.0.2.1 B1E2-2D4C-E6A2
    5a2e5107e70d525615a903f6391ad72b1c03352c
   ```
4. Use the following command to permit HSM client submissions on the RFS:

   ```
   rfs-setup --gang-client --write-noauth <FW-IPaddress>
   ```

   where *<FW-IPaddress>* is the firewall IP address.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 5 | Authenticate the firewall to the HSM.

1. In the firewall web interface, select **Device** > **Setup** > **HSM** and **Setup Hardware Security Module**.
2. Click **OK**.

   The firewall tries to authenticate to the HSM and displays a status message.
3. Click **OK**.

STEP 6 | Synchronize the firewall with the RFS by selecting **Device** > **Setup** > **HSM** and **Synchronize with Remote Filesystem**.

STEP 7 | Verify firewall connectivity and authentication with the HSM.

1. Select **Device** > **Setup** > **HSM** and check the authentication and connection Status:
   - **Green**—The firewall is successfully authenticated and connected to the HSM.
   - **Red**—The firewall failed to authenticate to the HSM or network connectivity to the HSM is down.
2. Check the Hardware Security Module Status to determine the authentication status.
   - **Name**—The name of the HSM.
   - **IP address**—The IP address of the HSM.
   - **Module State**—The current state of the HSM connection: `Authenticated` or `NotAuthenticated`.

Encrypt a Master Key Using an HSM

A master key encrypts all private keys and passwords on the firewall and Panorama. If you have security requirements to store your private keys in a secure location, you can encrypt the master key using an encryption key that is stored on an HSM. The firewall or Panorama then requests the HSM to decrypt the master key whenever it is required to decrypt a password or private key on the firewall. Typically, the HSM is in a highly secure location that is separate from the firewall or Panorama for greater security.

The HSM encrypts the master key using a wrapping key. To maintain security, you must occasionally change (refresh) this wrapping key.

Firewalls configured in FIPS/CC mode do not support master key encryption using an HSM.

The following topics describe how to encrypt the master key initially and how to refresh the master key encryption:

- Encrypt the Master Key
- Refresh the Master Key Encryption

Encrypt the Master Key

If you have not previously encrypted the master key on a firewall, use the following procedure to encrypt it. Use this procedure for first time encryption of a key, or if you define a new master key

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

and you want to encrypt it. If you want to refresh the encryption on a previously encrypted key, see Refresh the Master Key Encryption.

STEP 1 | Select **Device** > **Master Key and Diagnostics**.

STEP 2 | Specify the key that is currently used to encrypt all of the private keys and passwords on the firewall in the **Master Key** field.

STEP 3 | If changing the master key, enter the new master key and confirm.

STEP 4 | Select the **HSM** check box.

- **Life Time**—The number of days and hours after which the master key expires (range 1-730 days).
- **Time for Reminder**—The number of days and hours before expiration when the user is notified of the impending expiration (range 1–365 days).

STEP 5 | Click **OK**.

Refresh the Master Key Encryption

As a best practice, periodically refresh the master key encryption by rotating the wrapping key that encrypts it. The frequency of the rotation depends on your application. The wrapping key resides on your HSM. The following command is the same for SafeNet Network and nCipher nShield Connect HSMs.

STEP 1 | Log in to the firewall CLI.

STEP 2 | Use the following CLI command to rotate the wrapping key for the master key on an HSM:

```
> request hsm mkey-wrapping-key-rotation
```

If the master key is encrypted on the HSM, the CLI command will generate a new wrapping key on the HSM and encrypt the master key with the new wrapping key.

If the master key is not encrypted on the HSM, the CLI command will generate new wrapping key on the HSM for future use.

The old wrapping key is not deleted by this command.

Store Private Keys on an HSM

For added security, you can use an HSM to secure the private keys used in SSL/TLS decryption for:

- SSL Forward Proxy—The HSM can store the private key of the Forward Trust certificate that signs certificates in SSL/TLS forward proxy operations. The firewall will then send the certificates that it generates during such operations to the HSM for signing before forwarding the certificates to the client.
- SSL Inbound Inspection—The HSM can store the private keys for the internal servers for which you are performing SSL/TLS inbound inspection.

https://t.me/learningnets

If you use the DHE or ECDHE key exchange algorithms to enable perfect forward secrecy (PFS) support for SSL decryption, you can use an HSM to store the private keys for SSL Inbound Inspection. You can also use an HSM to store ECDSA keys used for SSL Forward Proxy or SSL Inbound Inspection decryption unless you are using TLSv1.3. For TLSv1.3 traffic, PAN-OS supports HSMs only for SSL Forward Proxy. It does not support HSMs for SSL Inbound Inspection.

STEP 1 | On the HSM, import or generate the certificate and private key used in your decryption deployment.

For instructions on importing or generating a certificate and private key on the HSM, refer to your HSM documentation.

STEP 2 | (nCipher nShield Connect only) Synchronize the key data from the nCipher nShield remote file system to the firewall.

Synchronization with the SafeNet Network HSM is automatic.

1. Access the firewall web interface and select **Device** > **Setup** > **HSM**.
2. **Synchronize with Remote Filesystem** (Hardware Security Operations settings).

STEP 3 | Import the certificate that corresponds to the HSM-stored key.

1. Select **Device** > **Certificate Management** > **Certificates** > **Device Certificates** and click **Import**.
2. Enter the **Certificate Name**.
3. **Browse** to the **Certificate File** on the HSM.
4. Select a **File Format**.
5. Select **Private Key resides on Hardware Security Module**.
6. Click **OK** and **Commit** your changes.

STEP 4 | (Forward Trust certificates only) Enable the certificate for use in SSL/TLS Forward Proxy.

1. Open the certificate you imported in Step 3 for editing.
2. Select **Forward Trust Certificate**.
3. Click **OK** and **Commit** your changes.

STEP 5 | Verify that you successfully imported the certificate onto the firewall.

Locate the certificate you imported in Step 3 and check the icon in the Key column:

- **Lock icon**—The private key for the certificate is on the HSM.
- **Error icon**—The private key is not on the HSM or the HSM is not properly authenticated or connected.

Manage the HSM Deployment

You can perform the following tasks to manage your HSM deployment:

- View the HSM configuration settings.

  Select **Device** > **Setup** > **HSM**.

https://t.me/learningnets

- Display detailed HSM information.

  Select **Show Detailed Information** from the Hardware Security Operations section.

  Information regarding the HSM servers, HSM HA status, and HSM hardware is displayed.

- Export Support file.

  Select **Export Support File** from the Hardware Security Operations section.

  A test file is created to help customer support when addressing a problem with an HSM configuration on the firewall.

- Reset HSM configuration.

  Select **Reset HSM Configuration** from the Hardware Security Operations section.

  Selecting this option removes all HSM connections. All authentication procedures must be repeated after using this option.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

High Availability

High availability (HA) is a deployment in which two firewalls are placed in a group or up to 16 firewalls are placed in an HA cluster and their configuration is synchronized to prevent a single point of failure on your network. A heartbeat connection between the firewall peers ensures seamless failover in the event that a peer goes down. Setting up HA provides redundancy and allows you to ensure business continuity.

- HA Overview
- HA Concepts
- Set Up Active/Passive HA
- Set Up Active/Active HA
- HA Clustering Overview
- HA Clustering Best Practices and Provisioning
- Configure HA Clustering
- Refresh HA1 SSH Keys and Configure Key Options
- HA Firewall States
- Reference: HA Synchronization
- CLI Cheat Sheet - HA

https://t.me/learningnets

HA Overview

You can configure two Palo Alto Networks firewalls as an HA pair or configure up to 16 firewalls as peer members of an HA cluster. The peers in the cluster can be HA pairs or standalone firewalls. HA allows you to minimize downtime by making sure that an alternate firewall is available in the event that a peer firewall fails. The firewalls in an HA pair or cluster use dedicated or in-band HA ports on the firewall to synchronize data—network, object, and policy configurations—and to maintain state information. Firewall-specific configuration such as management interface IP address or administrator profiles, HA specific configuration, log data, and the Application Command Center (ACC) information is not shared between peers.

For a consolidated application and log view across an HA pair, you must use Panorama, the Palo Alto Networks centralized management system. See Context Switch—Firewall or Panorama in the Panorama Administrator's Guide. Consult the Prerequisites for Active/Passive HA and Prerequisites for Active/Active HA. It is highly recommended that you use Panorama to provision HA cluster members. Consult the HA Clustering Best Practices and Provisioning.

When a failure occurs on a firewall in an HA pair or HA cluster and a peer firewall takes over the task of securing traffic, the event is called a Failover. The conditions that trigger a failover are:

- One or more of the monitored interfaces fail. (Link Monitoring)
- One or more of the destinations specified on the firewall cannot be reached. (Path Monitoring)
- The firewall does not respond to heartbeat polls. (Heartbeat Polling and Hello messages)
- A critical chip or software component fails, known as packet path health monitoring.

Palo Alto Networks firewalls support stateful active/passive or active/active high availability with session and configuration synchronization with a few exceptions:

- The VM-Series firewall on Azure and VM-Series firewall on AWS support active/passive HA only.

  On AWS, when you deploy the firewall with the Amazon Elastic Load Balancing (ELB) service, it does not support HA (in this case, ELB service provides the failover capabilities).
- The VM-Series firewall on Google Cloud Platform does not support HA.

Begin by understanding the HA Concepts and the HA Clustering Overview if you are going to configure HA clustering.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

HA Concepts

The following topics provide conceptual information about how HA works on a Palo Alto Networks firewall:

- HA Modes
- HA Links and Backup Links
- Device Priority and Preemption
- Failover
- LACP and LLDP Pre-Negotiation for Active/Passive HA
- Floating IP Address and Virtual MAC Address
- ARP Load-Sharing
- Route-Based Redundancy
- HA Timers
- Session Owner
- Session Setup
- NAT in Active/Active HA Mode
- ECMP in Active/Active HA Mode

HA Modes

You can set up the firewalls in an HA pair in one of two modes:

- **Active/Passive**— One firewall actively manages traffic while the other is synchronized and ready to transition to the active state, should a failure occur. In this mode, both firewalls share the same configuration settings, and one actively manages traffic until a path, link, system, or network failure occurs. When the active firewall fails, the passive firewall transitions to the active state and takes over seamlessly and enforces the same policies to maintain network security. Active/passive HA is supported in the virtual wire, Layer 2, and Layer 3 deployments.
- **Active/Active**— Both firewalls in the pair are active and processing traffic and work synchronously to handle session setup and session ownership. Both firewalls individually maintain session tables and routing tables and synchronize to each other. Active/active HA is supported in virtual wire and Layer 3 deployments.

  In active/active HA mode, the firewall does not support DHCP client. Furthermore, only the active-primary firewall can function as a DHCP Relay. If the active-secondary firewall receives DHCP broadcast packets, it drops them.

  An active/active configuration does not load-balance traffic. Although you can load-share by sending traffic to the peer, no load balancing occurs. Ways to load share sessions to both firewalls include using ECMP, multiple ISPs, and load balancers.

When deciding whether to use active/passive or active/active mode, consider the following differences:

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

- Active/passive mode has simplicity of design; it is significantly easier to troubleshoot routing and traffic flow issues in active/passive mode. Active/passive mode supports a Layer 2 deployment; active/active mode does not.
- Active/active mode requires advanced design concepts that can result in more complex networks. Depending on how you implement active/active HA, it might require additional configuration such as activating networking protocols on both firewalls, replicating NAT pools, and deploying floating IP addresses to provide proper failover. Because both firewalls are actively processing traffic, the firewalls use additional concepts of session owner and session setup to perform Layer 7 content inspection. Active/active mode is recommended if each firewall needs its own routing instances and you require full, real-time redundancy out of both firewalls all the time. Active/active mode has faster failover and can handle peak traffic flows better than active/passive mode because both firewalls are actively processing traffic.



In active/active mode, the HA pair can be used to temporarily process more traffic than what one firewall can normally handle. However, this should not be the norm because a failure of one firewall causes all traffic to be redirected to the remaining firewall in the HA pair. Your design must allow the remaining firewall to process the maximum capacity of your traffic loads with content inspection enabled. If the design oversubscribes the capacity of the remaining firewall, high latency and/or application failure can occur.

For information on setting up your firewalls in active/passive mode, see Set Up Active/Passive HA. For information on setting up your firewalls in active/active mode, see Set Up Active/Active HA.

In an HA cluster, all members are considered active; there is no concept of passive firewalls except for HA pairs in the clusters, which can keep their active/passive relationship after you add them to an HA cluster.

HA Links and Backup Links

The firewalls in an HA pair use HA links to synchronize data and maintain state information. Some models of the firewall have dedicated HA ports—Control link (HA1) and Data link (HA2), while others require you to use the in-band ports as HA links.

- For firewalls with dedicated HA ports, use these ports to manage communication and synchronization between the firewalls. For details, see HA Ports on Palo Alto Networks Firewalls.
- For firewalls without dedicated HA ports such as the PA-220 and PA-220R firewalls, as a best practice use the management port for the HA1 port, and use the dataplane port for the HA1 backup.



For firewalls without dedicated HA ports, decide which ports to use for HA1 and HA1 backup based on your environment and understanding which are the least used and least congested. Assign HA1 to the best interface and HA1 backup to the other one.

HA peers in an HA cluster can be a combination of standalone members and HA pairs. HA cluster members use an HA4 link and HA4 backup link to perform session state syncronization. HA1 (control link), HA2 (data link), and HA3 (packet-forwarding link) are not supported between cluster members that aren't HA pairs.

https://t.me/learningnets

HA Links and Backup Links	Description
Control Link	The HA1 link is used to exchange hellos, heartbeats, and HA state information, and management plane sync for routing, and User-ID information. The firewalls also use this link to synchronize configuration changes with its peer. The HA1 link is a Layer 3 link and requires an IP address. ICMP is used to exchange heartbeats between HA peers. Ports used for HA1—TCP port 28769 and 28260 for clear text communication; port 28 for encrypted communication (SSH over TCP). If you enable encryption on the HA1 link, you can also Refresh HA1 SSH Keys and Configure Key Options.
Data Link	The HA2 link is used to synchronize sessions, forwarding tables, IPSec security associations and ARP tables between firewalls in an HA pair. Data flow on the HA2 link is always unidirectional (except for the HA2 keep-alive); it flows from the active or active-primary firewall to the passive or active-secondary firewall. The HA2 link is a Layer 2 link, and it uses ether type 0x7261 by default. Ports used for HA2—The HA data link can be configured to use either IP (protocol number 99) or UDP (port 29281) as the transport, and thereby allow the HA data link to span subnets.
HA1 and HA2 Backup Links	Provide redundancy for the HA1 and the HA2 links. In-band ports can be used for backup links for both HA1 and HA2 connections when dedicated backup links are not available. Consider the following guidelines when configuring backup HA links: • The IP addresses of the primary and backup HA links must not overlap each other. • HA backup links must be on a different subnet from the primary HA links. • HA1-backup and HA2-backup ports must be configured on separate physical ports. The HA1-backup link uses port 28770 and 28260. • PA-3200 Series firewalls don't support an IPv6 address for the HA1-backup link; use an IPv4 address. *Palo Alto Networks recommends enabling heartbeat backup (uses port 28771 on the MGT interface) if you use an in-band port for the HA1 or the HA1 backup links.*
Packet-Forwarding Link	In addition to HA1 and HA2 links, an active/active deployment also requires a dedicated HA3 link. The firewalls use this link for forwarding packets to the peer during session setup and asymmetric traffic flow.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

HA Links and Backup Links	Description
	The HA3 link is a Layer 2 link that uses MAC-in-MAC encapsulation. It does not support Layer 3 addressing or encryption. PA-7000 Series firewalls synchronize sessions across the NPCs one-for-one. On PA-800 Series, PA-3200 Series, and PA-5200 Series firewalls, you can configure aggregate interfaces as an HA3 link. The aggregate interfaces can also provide redundancy for the HA3 link; you cannot configure backup links for the HA3 link. On PA-3200 Series, PA-5200 Series, and PA-7000 Series firewalls, the dedicated HSCI ports support the HA3 link. The firewall adds a proprietary packet header to packets traversing the HA3 link, so the MTU over this link must be greater than the maximum packet length forwarded.
HA4 Link and HA4 Backup Link	The HA4 link and HA4 backup link perform session cache synchronization among all HA cluster members having the same cluster ID. The HA4 link between cluster members detects connectivity failures between cluster members by sending and receiving Layer 2 keepalive messages. View the status of the HA4 and HA4 backup links on the firewall dashboard.

HA Ports on Palo Alto Networks Firewalls

When connecting two Palo Alto Networks® firewalls in a high availability (HA) configuration, we recommend that you use the dedicated HA ports for HA Links and Backup Links. These dedicated ports include: the HA1 ports labeled HA1, HA1-A, and HA1-B used for HA control and synchronization traffic; and HA2 and the High Speed Chassis Interconnect (HSCI) ports used for HA session setup traffic. The PA-5200 Series firewalls have multipurpose auxiliary ports labeled AUX-1 and AUX-2 that you can configure for HA1 traffic.

You can also configure the HSCI port for HA3, which is used for packet forwarding to the peer firewall during session setup and asymmetric traffic flow (active/active HA only). The HSCI port can be used for HA2 traffic, HA3 traffic, or both.

The HA1 and AUX links provide synchronization for functions that reside on the management plane. Using the dedicated HA interfaces on the management plane is more efficient than using the in-band ports as this eliminates the need to pass the synchronization packets over the dataplane.

If your firewall does not have dedicated HA ports, you can configure data ports as HA interfaces. If your firewall does have dedicated HA ports but does not have a dedicated HA backup port, you can also configure data ports as backups to dedicated HA ports.

Whenever possible, connect HA ports directly between the two firewalls in an HA pair (not through a switch or router) to avoid HA link and communications problems that could occur if there is a network issue.

Use the following table to learn about dedicated HA ports and how to connect the HA Links and Backup Links:

https://t.me/learningnets

Model	Front-Panel Dedicated Port(s)
PA-800 Series Firewalls	• **HA1 and HA2**—Ethernet 10Mbps/100Mbps/1000Mbps ports used for HA1 and HA2 in both HA Modes. • **For HA1 traffic**—Connect the HA1 port on the first firewall directly to the HA1 port on the second firewall in the pair or connect these ports together through a switch or router. • **For HA2 traffic**—Connect the HA2 port on the first firewall directly to the HA2 port on the second firewall in the pair or connect these ports together through a switch or router.
PA-3200 Series Firewalls	• **HA1-A and HA1-B**—Ethernet 10Mbps/100Mbps/1000Mbps ports used for HA1 traffic in both HA Modes. • **For HA1 traffic**—Connect the HA1-A port on the first firewall directly to the HA1-A port on the second firewall in the pair or connect them together through a switch or router. • **For a backup to the HA1-A connection**—Connect the HA1-B port on the first firewall directly to the HA1-B port on the second firewall in the pair or connect them together through a switch or router. *If the firewall dataplane restarts due to a failure or manual restart, the HA1-B link will also restart. If this occurs and the HA1-A link is not connected and configured, then a split brain condition occurs. Therefore, we recommend that you connect and configure the HA1-A ports and the HA1-B ports to provide redundancy and to avoid split brain issues.* *You can remap the firewall's SFP ports as HA1-A and HA1-B ports via PAN-OS or Panorama.* • **HSCI**—The HSCI port is a Layer 1 SFP+ interface that connects two PA-3200 Series firewalls in an HA configuration. Use this port for an HA2 connection, HA3 connection, or both. The traffic carried on the HSCI ports is raw Layer 1 traffic, which is not routable or switchable. Therefore, you must connect the HSCI ports directly to each other (from the HSCI port on the first firewall to the HSCI port on the second firewall).
PA-5200 Series Firewalls	• **HA1-A and HA1-B**—Ethernet 10Mbps/100Mbps/1000Mbps ports used for HA1 traffic in both HA Modes. • **For HA1 traffic**—Connect the HA1-A port on the first firewall directly to the HA1-A port on the second firewall in the pair or connect them together through a switch or router. • **For a backup to the HA1-A connection**—Connect the HA1-B port on the first firewall directly to the HA1-B port on the

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Model	Front-Panel Dedicated Port(s)
	second firewall in the pair or connect them together through a switch or router. • **HSCI**—The HSCI port is a Layer 1 interface that connects two PA-5200 Series firewalls in an HA configuration. Use this port for an HA2 connection, HA3 connection, or both. *The HSCI port on the PA-5220 firewall is a QSFP+ port and the HSCI port on the PA-5250, PA-5260, and PA-5280 firewalls is a QSFP28 port.* The traffic carried on the HSCI port is raw Layer 1 traffic, which is not routable or switchable. Therefore, you must connect the HSCI ports directly to each other (from the HSCI port on the first firewall to the HSCI port on the second firewall).
PA-5200 Series Firewalls (continued)	• **AUX-1 and AUX-2**—The auxiliary SFP+ ports are multipurpose ports that you can configure for HA1, management functions, or log forwarding to Panorama. Use these ports when you need a fiber connection for one of these functions. • **For HA1 traffic**—Connect the AUX-1 port on the first firewall directly to the AUX-1 port on the second firewall in the pair or connect them together through a switch or router. • **For a backup to the AUX-1 connection**—Connect the AUX-2 port on the first firewall directly to the AUX-2 port on the second firewall in the pair or connect them together through a switch or router.
PA-5450 Firewall	• **HA1-A and HA1-B**—Ethernet 1Gbps/10Gbps ports used for HA1 traffic in both HA Modes. • **For HA1 traffic**—Connect the HA1-A port on the first firewall directly to the HA1-A port on the second firewall in the pair or connect them together through a switch or router. • **For a backup to the HA1-A connection**—Connect the HA1-B port on the first firewall directly to the HA1-B port on the second firewall in the pair or connect them together through a switch or router. • **HSCI-A and HSCI-B**—The HSCI ports are Layer 1 QSFP + interfaces that connect two PA-5450 firewalls in an HA

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Model	Front-Panel Dedicated Port(s)
	configuration. Use these ports for an HA2 connection, HA3 connection, or both. The traffic carried on the HSCI ports is raw Layer 1 traffic, which is not routable or switchable. Therefore, you must connect these ports as follows: • **For HA2 and HA3 traffic**—Connect the HSCI-A port on the first firewall directly to the HSCI-A port on the second firewall. • **For a backup to the HSCI-A connection**—Connect the HSCI-B port on the first firewall directly to the HSCI-B port on the second firewall.
PA-7000 Series Firewalls	• **HA1-A and HA1-B**—Ethernet 10Mbps/100Mbps/1000Mbps ports used for HA1 traffic in both HA Modes. • **For HA1 traffic**—Connect the HA1-A port on the first firewall directly to the HA1-A port on the second firewall in the pair or connect them together through a switch or router. • **For a backup to the HA1-A connection**—Connect the HA1-B port on the first firewall directly to the HA1-B port on the second firewall in the pair or connect them together through a switch or router. *You cannot configure an HA1 connection on the NPC data ports or the management (MGT) port.* • **HSCI-A and HSCI-B**—The HSCI ports are Layer 1 QSFP+ interfaces that connect two PA-7000 Series firewalls in an HA configuration. Use these ports for an HA2 connection, HA3 connection, or both. The traffic carried on the HSCI ports is raw Layer 1 traffic, which is not routable or switchable. Therefore, you must connect these ports as follows: • **For HA2 and HA3 traffic**—Connect the HSCI-A port on the first firewall directly to the HSCI-A port on the second firewall. *For HA2 or HA2/HA3 traffic, the PA-7000 Series firewalls synchronize sessions across the NPCs one-for-one.* • **For a backup to the HSCI-A connection**—Connect the HSCI-B port on the first firewall directly to the HSCI-B port on the second firewall.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Model	Front-Panel Dedicated Port(s)
	HA2 and HA2-Backup links can be configured to use a dataplane interface instead of the HSCI ports. However, if configured this way, both the HA2 and HA2-Backup links need to use dataplane interfaces. A mix of a dataplane port and an HSCI port for either HA2 or HA2-Backup will result in a commit failure. This applies to the PA-7050-SMC, PA-7080-SMC, PA-7050-SMC-B, and PA-7080-SMC-B.

Device Priority and Preemption

The firewalls in an Active-Passive HA pair can be assigned a *device priority* value to indicate a preference for which firewall should assume the active role. If you need to use a specific firewall in the HA pair for actively securing traffic, you must enable the preemptive behavior on both the firewalls and assign a device priority value for each firewall. The firewall with the lower numerical value, and therefore *higher priority*, is designated as active. The other firewall is the passive firewall.

The same is true for an Active-Active HA pair; however, the *device ID* is used to assign a device priority value. Similarly, the lower numerical value in device ID corresponds to a higher priority. The firewall with the higher priority becomes active-primary and the paired firewall becomes active-secondary.

By default, preemption is disabled on the firewalls and must be enabled on both firewalls. When enabled, the preemptive behavior allows the firewall with the *higher priority* (lower numerical value) to resume as active or active-primary after it recovers from a failure. When preemption occurs, the event is logged in the system logs.

Failover

When a failure occurs on one firewall and the peer in the HA pair (or a peer in the HA cluster) takes over the task of securing traffic, the event is called a *failover*. A failover is triggered, for example, when a monitored metric on a firewall in the HA pair fails. The metrics that the firewall monitors for detecting a firewall failure are:

- **Heartbeat Polling and Hello messages**

  The firewalls use hello message and heartbeats to verify that the peer firewall is responsive and operational. Hello messages are sent from one peer to the other at the configured *Hello Interval* to verify the state of the firewall. The heartbeat is an ICMP ping to the HA peer over the control link, and the peer responds to the ping to establish that the firewalls are connected and responsive. By default, the interval for the heartbeat is 1000 milliseconds. A ping is sent every 1000 milliseconds and if there are three consecutive heartbeat losses, a failovers occurs. For details on the HA timers that trigger a failover, see HA Timers.

- **Link Monitoring**

  You can specify a group of physical interfaces that the firewall will monitor (a link group) and the firewall monitors the state of each link in the group (link up or link down). You determine

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

the failure condition for the link group: **Any** link down or **All** links down in the group constitutes a link group failure (but not necessarily a failover).

You can create multiple link groups. Therefore, you also determine the failure condition of the set of link groups: **Any** link group fails or **All** link groups fail, which determines when a failover is triggered. The default behavior is that failure of **Any** one link in **Any** link group causes the firewall to change the HA state to non-functional (or to tentative state in active/active mode) to indicate a failure of a monitored object.

- **Path Monitoring**

  You can specify a destination IP group of IP address that the firewall will monitor. The firewall monitors the full path through the network to mission-critical IP addresses using ICMP pings to verify reachability of the IP address. The default interval for pings is 200ms. An IP address is considered unreachable when 10 consecutive pings (the default value) fail. You specify the failure condition for the IP addresses in a destination IP group: **Any** IP address unreachable or **All** IP addresses unreachable in the group. You can specify multiple destination IP groups for a path group for a virtual wire, VLAN, or virtual router; you specify the failure condition of destination IP groups in a path group: **Any** or **All**, which constitutes a path group failure. You can configure multiple virtual wire path groups, VLAN path groups, and virtual router path groups.

  You also determine the global failure condition: **Any** path group fails or **All** path groups fail, which determines when a failover is triggered. The default behavior is that **Any** one of the IP addresses becoming unreachable in **Any** destination IP group in **Any** virtual wire, VLAN, or virtual router path group causes the firewall to change the HA state to non-functional (or to tentative state in active/active mode) to indicate a failure of a monitored object.

In addition to the failover triggers listed above, a failover also occurs when the administrator suspends the firewall or when preemption occurs.

On PA-3200 Series, PA-5200 Series, and PA-7000 Series firewalls, a failover can occur when an internal health check fails. This health check is not configurable and is enabled to monitor the critical components, such as the FPGA and CPUs. Additionally, general health checks occur on any platform, causing failover.

The following describes what occurs in the event of a failure of a Network Processing Card (NPC) on a PA-7000 Series firewall that is a member of an HA cluster:

- If the NPC that is being used to hold the HA clustering session cache (a copy of the other members' sessions) goes down, the firewall goes non-functional. When this occurs, the session distribution device (such as a load balancer) must detect that the firewall is down and distribute session load to the other members of the cluster.
- If the NPC of a cluster member goes down and no link monitoring or path monitoring was enabled on that NPC, the PA-7000 Series firewall member will stay up, but with a lower capacity because one NPC is down.
- If the NPC of a cluster member goes down and link monitoring or path monitoring was enabled on that NPC, the PA-7000 Series firewall will go non-functional and the session distribution device (such as a load balancer) must detect that the firewall is down and distribute session load to the other members of the cluster.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

LACP and LLDP Pre-Negotiation for Active/Passive HA

If a firewall uses LACP or LLDP, negotiation of those protocols upon failover prevents sub-second failover. However, you can enable an interface on a passive firewall to negotiate LACP and LLDP prior to failover. Thus, a firewall in Passive or Non-functional HA state can communicate with neighboring devices using LACP or LLDP. Such pre-negotiation speeds up failover.

All firewall models except VM-Series firewalls support a pre-negotiation configuration, which depends on whether the Ethernet or AE interface is in a Layer 2, Layer 3, or virtual wire deployment. An HA passive firewall handles LACP and LLDP packets in one of two ways:

- **Active**—The firewall has LACP or LLDP configured on the interface and actively participates in LACP or LLDP pre-negotiation, respectively.
- **Passive**—LACP or LLDP is not configured on the interface and the firewall does not participate in the protocol, but allows the peers on either side of the firewall to pre-negotiate LACP or LLDP, respectively.

The following table displays which deployments are supported on Aggregate Ethernet (AE) and Ethernet interfaces.

Interface Deployment	AE Interface	Ethernet Interface
LACP in Layer 2	Active	Not supported
LACP in Layer 3	Active	Not supported
LACP in Virtual Wire	Not supported	Passive
LLDP in Layer 2	Active	Active
LLDP in Layer 3	Active	Active
LLDP in Virtual Wire	Active	• Active if LLDP itself is configured. • Passive if LLDP itself is not configured.

Pre-negotiation is not supported on subinterfaces or tunnel interfaces.

To configure LACP or LLDP pre-negotiation, see the step (Optional) Enable LACP and LLDP Pre-Negotiation for Active/Passive HA for faster failover if your network uses LACP or LLDP.

Floating IP Address and Virtual MAC Address

In a Layer 3 deployment of HA active/active mode, you can assign floating IP addresses, which move from one HA firewall to the other if a link or firewall fails. The interface on the firewall that owns the floating IP address responds to ARP requests with a virtual MAC address.

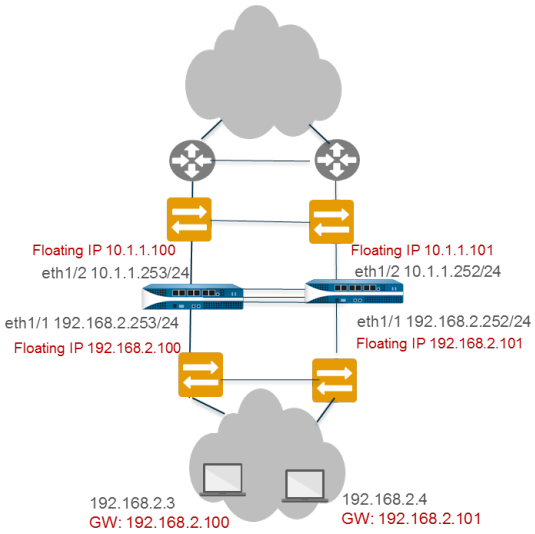
Floating IP addresses are recommended when you need functionality such as Virtual Router Redundancy Protocol (VRRP). Floating IP addresses can also be used to implement VPNs and source NAT, allowing for persistent connections when a firewall offering those services fails.

https://t.me/learningnets

As shown in the figure below, each HA firewall interface has its own IP address and floating IP address. The interface IP address remains local to the firewall, but the floating IP address moves between the firewalls upon firewall failure. You configure the end hosts to use a floating IP address as its default gateway, allowing you to load balance traffic to the two HA peers. You can also use external load balancers to load balance traffic.

If a link or firewall fails or a path monitoring event causes a failover, the floating IP address and virtual MAC address move over to the functional firewall. (In the figure below, each firewall has two floating IP addresses and virtual MAC addresses; they all move over if the firewall fails.) The functioning firewall sends a gratuitous ARP to update the MAC tables of the connected switches to inform them of the change in floating IP address and MAC address ownership to redirect traffic to itself.

After the failed firewall recovers, by default the floating IP address and virtual MAC address move back to firewall with the Device ID [0 or 1] to which the floating IP address is bound. More specifically, after the failed firewall recovers, it comes on line. The currently active firewall determines that the firewall is back online and checks whether the floating IP address it is handling belongs natively to itself or the other firewall. If the floating IP address was originally bound to the other Device ID, the firewall automatically gives it back. (For an alternative to this default behavior, see Use Case: Configure Active/Active HA with Floating IP Address Bound to Active-Primary Firewall.)



Each firewall in the HA pair creates a virtual MAC address for each of its interfaces that has a floating IP address or ARP Load-Sharing IP address.

The format of the virtual MAC address (on firewalls other than PA-7000, PA-5200, and PA-3200 Series firewalls) is 00-1B-17-00-xx-yy, where 00-1B-17 is the vendor ID (of Palo Alto Networks in this case), 00 is fixed, xx indicates the Device ID and Group ID as shown in the following figure, and yy is the Interface ID:

7	6	5 4 3 2 1 0	7 6 5 4 3 2 1 0
Device-ID	0	Group-ID	Interface-ID

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

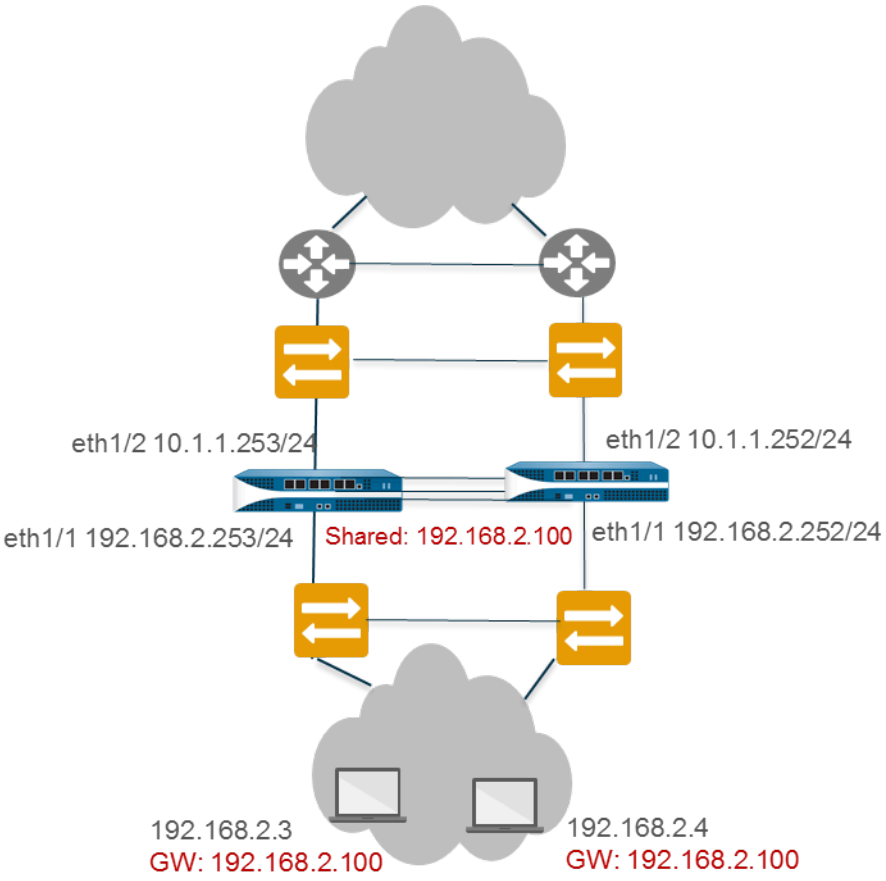
The format of the virtual MAC address on PA-7000, PA-5200, and PA-3200 Series firewalls is B4-0C-25-xx-xx-xx, where B4-0C-25 is the vendor ID (of Palo Alto Networks in this case), and the next 24 bits indicate the Device ID, Group ID and Interface ID as follows:

765	4	321076	5432	1076543210
111	Device-ID	Group-ID	0000	Interface-ID

When a new active firewall takes over, it sends gratuitous ARPs from each of its connected interfaces to inform the connected Layer 2 switches of the new location of the virtual MAC address. To configure floating IP addresses, see Use Case: Configure Active/Active HA with Floating IP Addresses.

ARP Load-Sharing

In a Layer 3 interface deployment and active/active HA configuration, ARP load-sharing allows the firewalls to share an IP address and provide gateway services. Use ARP load-sharing only when no Layer 3 device exists between the firewall and end hosts, that is, when end hosts use the firewall as their default gateway.



In such a scenario, all hosts are configured with a single gateway IP address. One of the firewalls responds to ARP requests for the gateway IP address with its virtual MAC address. Each firewall has a unique virtual MAC address generated for the shared IP address. The load-sharing algorithm

 ©2022 Palo Alto Networks, Inc.

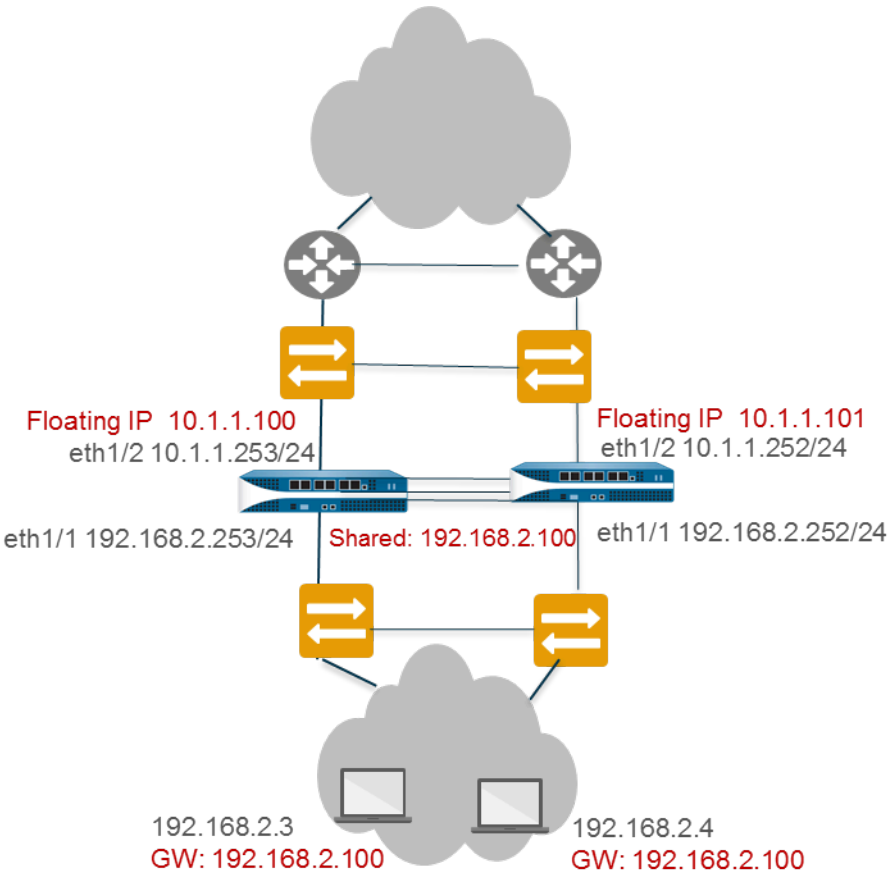
https://t.me/learningnets

that controls which firewall will respond to the ARP request is configurable; it is determined by computing the hash or modulo of the source IP address of the ARP request.

After the end host receives the ARP response from the gateway, it caches the MAC address and all traffic from the host is routed via the firewall that responded with the virtual MAC address for the lifetime of the ARP cache. The lifetime of the ARP cache depends on the end host operating system.

If a link or firewall fails, the floating IP address and virtual MAC address move over to the functional firewall. The functional firewall sends gratuitous ARPs to update the MAC table of the connected switches to redirect traffic from the failed firewall to itself. See Use Case: Configure Active/Active HA with ARP Load-Sharing.

You can configure interfaces on the WAN side of the HA firewalls with floating IP addresses, and configure interfaces on the LAN side of the HA firewalls with a shared IP address for ARP load-sharing. For example, the figure below illustrates floating IP addresses for the upstream WAN edge routers and an ARP load-sharing address for the hosts on the LAN segment.

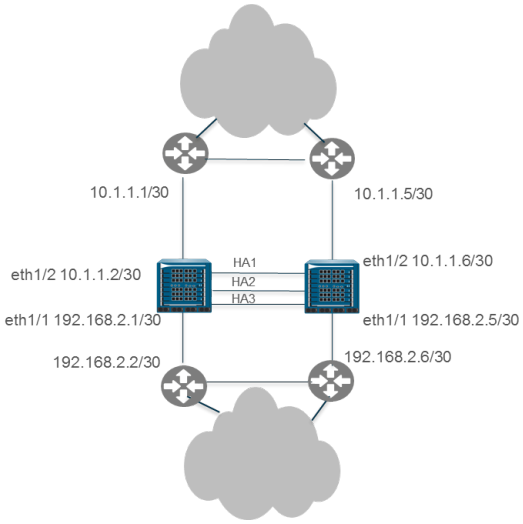


Route-Based Redundancy

In a Layer 3 interface deployment and active/active HA configuration, the firewalls are connected to routers, not switches. The firewalls use dynamic routing protocols to determine the best path (asymmetric route) and to load share between the HA pair. In such a scenario, no floating IP addresses are necessary. If a link, monitored path, or firewall fails, or if Bidirectional Forwarding Detection (BFD) detects a link failure, the routing protocol (RIP, OSPF, or BGP) handles the

https://t.me/learningnets

rerouting of traffic to the functioning firewall. You configure each firewall interface with a unique IP address. The IP addresses remain local to the firewall where they are configured; they do not move between devices when a firewall fails. See Use Case: Configure Active/Active HA with Route-Based Redundancy.



HA Timers

High availability (HA) timers facilitate a firewall to detect a firewall failure and trigger a failover. To reduce the complexity in configuring timers for an HA pair, you can select from three profiles: **Recommended**, **Aggressive** and **Advanced**. These profiles auto-populate the optimum HA timer values for the specific firewall platform to enable a speedier HA deployment.

Use the **Recommended** profile for typical failover timer settings and the **Aggressive** profile for faster failover timer settings. The **Advanced** profile allows you to customize the timer values to suit your network requirements.

The following table describes each timer included in the profiles and the current preset values (Recommended/Aggressive) across the different hardware models; these values are for current reference only and can change in a subsequent release.

Timers that affect members of an HA cluster are described in Configure HA Clustering.

Timers	Description	PA-7000 Series PA-5200 Series PA-3200 Series	PA-800 Series PA-220 VM-Series	Panorama Virtual Appliance Panorama M-Series
Monitor Fail Hold Up Time (ms)	Interval during which the firewall will remain active following a path monitor or link monitor failure. This setting is	0/0	0/0	0/0

©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Timers	Description	PA-7000 Series PA-5200 Series PA-3200 Series	PA-800 Series PA-220 VM-Series	Panorama Virtual Appliance Panorama M-Series
	recommended to avoid an HA failover due to the occasional flapping of neighboring devices.			
Preemption Hold Time (min)	Time that a passive or active-secondary firewall will wait before taking over as the active or active-primary firewall.	1/1	1/1	1/1
Heartbeat Interval (ms)	Frequency at which the HA peers exchange heartbeat messages in the form of an ICMP (ping).	1000/1000	2000/1000	2000/1000
Promotion Hold Time (ms)	Time that the passive firewall (in active/passive mode) or the active-secondary firewall (in active/active mode) will wait before taking over as the active or active-primary firewall after communications with the HA peer have been lost. This hold time will begin only after the peer failure declaration has been made.	2000/500	2000/500	2000/500
Additional Master Hold Up Time (ms)	Time interval in milliseconds that is applied to the same event as Monitor Fail Hold Up Time (range is 0 to 60,000; default is 500). The additional time interval is applied only to the active	500/500	500/500	7000/5000

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Timers	Description	PA-7000 Series PA-5200 Series PA-3200 Series	PA-800 Series PA-220 VM-Series	Panorama Virtual Appliance Panorama M-Series
	firewall in active/passive mode and to the active-primary firewall in active/active mode. This timer is recommended to avoid a failover when both firewalls experience the same link/path monitor failure simultaneously.			
Hello Interval (ms)	Interval in milliseconds between hello packets that are sent to verify that the HA functionality on the other firewall is operational (range is 8,000 to 60,000; default is 8,000).	8000/8000	8000/8000	8000/8000
Flap Max	A flap is counted when one of the following occurs: • A preemption-enabled firewall leaves the active state within 20 minutes after becoming active. • A link or path fails to stay up for 10 minutes after becoming functional. In the case of a failed preemption or non-functional loop, this value indicates the maximum number of flaps that are permitted before the firewall is	3/3	3/3	Not Applicable

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Timers	Description	PA-7000 Series PA-5200 Series PA-3200 Series	PA-800 Series PA-220 VM-Series	Panorama Virtual Appliance Panorama M-Series
	suspended (range 0 to 16; default is 3).			

Session Owner

In an HA active/active configuration, both firewalls are active simultaneously, which means packets can be distributed between them. Such distribution requires the firewalls to fulfill two functions: session ownership and session setup. Typically, each firewall of the pair performs one of these functions, thereby avoiding race conditions that can occur in asymmetrically routed environments.

You configure the session owner of sessions to be either the firewall that receives the First Packet of a new session from the end host or the firewall that is in active-primary state (the Primary device). If Primary device is configured, but the firewall that receives the first packet is not in active-primary state, the firewall forwards the packet to the peer firewall (the session owner) over the HA3 link.

The session owner performs all Layer 7 processing, such as App-ID, Content-ID, and threat scanning for the session. The session owner also generates all traffic logs for the session.

If the session owner fails, the peer firewall becomes the session owner. The existing sessions fail over to the functioning firewall and no Layer 7 processing is available for those sessions. When a firewall recovers from a failure, by default, all sessions it owned before the failure revert back to that original firewall; Layer 7 processing does not resume.

If you configure session ownership to be Primary device, the session setup defaults to Primary device also.

Palo Alto Networks recommends setting the Session Owner to First Packet and the Session Setup to IP Modulo unless otherwise indicated in a specific use case. Setting the Session Owner to First Packet reduces traffic across the HA3 link and helps distribute the dataplane load across peers.

Setting Session Owner and Session Setup to Primary Device causes the active-primary firewall to perform all traffic processing. You might want to configure this for one of these reasons:

- *You are troubleshooting and capturing logs and pcaps, so that packet processing is not split between the firewalls.*
- *You want to force the active/active HA pair to function like an active/passive HA pair. See* Use Case: Configure Active/Active HA with Floating IP Address Bound to Active-Primary Firewall.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Session Setup

The session setup firewall performs the Layer 2 through Layer 4 processing necessary to set up a new session. The session setup firewall also performs NAT using the NAT pool of the session owner. You determine the session setup firewall in an active/active configuration by selecting one of the following session setup load sharing options.

Session Setup Option	Description
IP Modulo	The firewall distributes the session setup load based on parity of the source IP address. This is a deterministic method of sharing the session setup.
IP Hash	The firewall uses a hash of the source and destination IP addresses to distribute session setup responsibilities.
Primary Device	The active-primary firewall always sets up the session; only one firewall performs all session setup responsibilities.
First Packet	The firewall that receives the first packet of a session performs session setup.

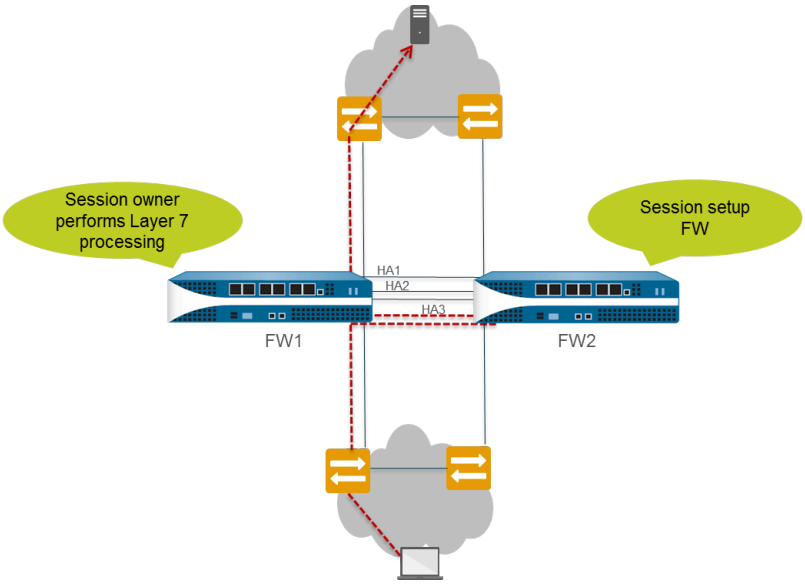


- *If you want to load-share the session owner and session setup responsibilities, set session owner to First Packet and session setup to IP modulo. These are the recommended settings.*
- *If you want to do troubleshooting or capture logs or pcaps, or if you want an active/active HA pair to function like an active/passive HA pair, set both the session owner and session setup to Primary device so that the active-primary device performs all traffic processing. See* Use Case: Configure Active/Active HA with Floating IP Address Bound to Active-Primary Firewall.

The firewall uses the HA3 link to send packets to its peer for session setup if necessary. The following figure and text describe the path of a packet that firewall FW1 receives for a new session. The red dotted lines indicate FW1 forwarding the packet to FW2 and FW2 forwarding the packet back to FW1 over the HA3 link.

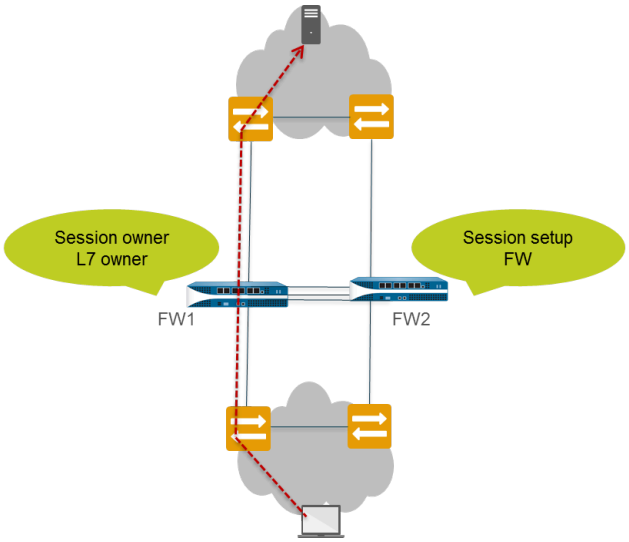
 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets



- The end host sends a packet to FW1.
- FW1 examines the contents of the packet to match it to an existing session. If there is no session match, FW1 determines that it has received the first packet for a new session and therefore becomes the session owner (assuming **Session Owner Selection** is set to **First Packet**).
- FW1 uses the configured session setup load-sharing option to identify the session setup firewall. In this example, FW2 is configured to perform session setup.
- FW1 uses the HA3 link to send the first packet to FW2.
- FW2 sets up the session and returns the packet to FW1 for Layer 7 processing, if any.
- FW1 then forwards the packet out the egress interface to the destination.

The following figure and text describe the path of a packet that matches an existing session:



- The end host sends a packet to FW1.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

- FW1 examines the contents of the packet to match it to an existing session. If the session matches an existing session, FW1 processes the packet and sends the packet out the egress interface to the destination.

NAT in Active/Active HA Mode

In an active/active HA configuration:

- You must bind each Dynamic IP (DIP) NAT rule and Dynamic IP and Port (DIPP) NAT rule to either Device ID 0 or Device ID 1.
- You must bind each static NAT rule to either Device ID 0, Device ID 1, both Device IDs, or the firewall in active-primary state.

Thus, when one of the firewalls creates a new session, the Device ID **0** or Device ID **1** binding determines which NAT rules match the firewall. The device binding must include the session owner firewall to produce a match.

The session setup firewall performs the NAT policy match, but the NAT rules are evaluated based on the session owner. That is, the session is translated according to NAT rules that are bound to the session owner firewall. While performing NAT policy matching, a firewall skips all NAT rules that are not bound to the session owner firewall.

For example, suppose the firewall with Device ID 1 is the session owner and session setup firewall. When the firewall with Device ID 1 tries to match a session to a NAT rule, it skips all rules bound to Device ID 0. The firewall performs the NAT translation only if the session owner and the Device ID in the NAT rule match.

You will typically create device-specific NAT rules when the peer firewalls use different IP addresses for translation.

If one of the peer firewalls fails, the active firewall continues to process traffic for synchronized sessions from the failed firewall, including NAT traffic. In a source NAT configuration, when one firewall fails:

- The floating IP address that is used as the Translated IP address of the NAT rule transfers to the surviving firewall. Hence, the existing sessions that fail over will still use this IP address.
- All new sessions will use the device-specific NAT rules that the surviving firewall naturally owns. That is, the surviving firewall translates new sessions using only the NAT rules that match its Device ID; it ignores any NAT rules bound to the failed Device ID.

For examples of active/active HA with NAT, see:

- Use Case: Configure Active/Active HA with Source DIPP NAT Using Floating IP Addresses
- Use Case: Configure Separate Source NAT IP Address Pools for Active/Active HA Firewalls
- Use Case: Configure Active/Active HA for ARP Load-Sharing with Destination NAT
- Use Case: Configure Active/Active HA for ARP Load-Sharing with Destination NAT in Layer 3

ECMP in Active/Active HA Mode

When an active/active HA peer fails, its sessions transfer to the new active-primary firewall, which tries to use the same egress interface that the failed firewall was using. If the firewall finds that interface among the ECMP paths, the transferred sessions will take the same egress interface

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

and path. This behavior occurs regardless of the ECMP algorithm in use; using the same interface is desirable.

Only if no ECMP path matches the original egress interface will the active-primary firewall select a new ECMP path.

If you did not configure the same interfaces on the active/active peers, upon failover the active-primary firewall selects the next best path from the FIB table. Consequently, the existing sessions might not be distributed according to the ECMP algorithm.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Set Up Active/Passive HA

- Prerequisites for Active/Passive HA
- Configuration Guidelines for Active/Passive HA
- Configure Active/Passive HA
- Define HA Failover Conditions
- Verify Failover

Prerequisites for Active/Passive HA

To set up high availability on your Palo Alto Networks firewalls, you need a pair of firewalls that meet the following requirements:

- **The same model**—Both the firewalls in the pair must be of the same hardware model or virtual machine model.
- **The same PAN-OS version**—Both the firewalls should be running the same PAN-OS version and must each be up-to-date on the application, URL, and threat databases.
- The same multi virtual system capability—Both firewalls must have **Multi Virtual System Capability** either enabled or not enabled. When enabled, each firewall requires its own multiple virtual systems licenses.
- **The same type of interfaces**—Dedicated HA links, or a combination of the management port and in-band ports that are set to *interface type* HA.
  - Determine the IP address for the HA1 (control) connection between the HA peers. The HA1 IP address for both peers must be on the same subnet if they are directly connected or are connected to the same switch.

    For firewalls without dedicated HA ports, you can use the management port for the control connection. Using the management port provides a direct communication link between the management planes on both firewalls. However, because the management ports will not be directly cabled between the peers, make sure that you have a route that connects these two interfaces across your network.
  - If you use Layer 3 as the transport method for the HA2 (data) connection, determine the IP address for the HA2 link. Use Layer 3 only if the HA2 connection must communicate over a routed network. The IP subnet for the HA2 links must not overlap with that of the HA1 links or with any other subnet assigned to the data ports on the firewall.
- **The same set of licenses**—Licenses are unique to each firewall and cannot be shared between the firewalls. Therefore, you must license both firewalls identically. If both firewalls do not have an identical set of licenses, they cannot synchronize configuration information and maintain parity for a seamless failover.

> *As a best practice, if you have an existing firewall and you want to add a new firewall for HA purposes and the new firewall has an existing configuration* Reset the Firewall to Factory Default Settings *on the new firewall. This ensures that the new firewall has a clean configuration. After HA is configured, you will then sync the configuration on the primary firewall to the newly introduced firewall with the clean configuration.*

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Configuration Guidelines for Active/Passive HA

To set up an active (PeerA) passive (PeerB) pair in HA, you must configure some options identically on both firewalls and some independently (non-matching) on each firewall. These HA settings are not synchronized between the firewalls. For details on what is/is not synchronized, see Reference: HA Synchronization.

The following checklist details the settings that you must configure identically on both firewalls:

- You must enable HA on both firewalls.
- You must configure the same Group ID value on both firewalls. The firewall uses the Group ID value to create a virtual MAC address for all the configured interfaces. See Floating IP Address and Virtual MAC Address for information about virtual MAC addresses. When a new active firewall takes over, it sends Gratuitous ARP messages from each of its connected interfaces to inform the connected Layer 2 switches of the virtual MAC address' new location.
- If you are using in-band ports as HA links, you must set the interfaces for the HA1 and HA2 links to type HA.
- Set the HA Mode to Active Passive on both firewalls.
- If required, enable preemption on both firewalls. The device priority value, however, must not be identical.
- If required, configure encryption on the HA1 link (for communication between the HA peers) on both firewalls.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

- Based on the combination of HA1 and HA1 Backup ports you are using, use the following recommendations to decide whether you should enable heartbeat backup:

  *HA functionality (HA1 and HA1 backup) is not supported on the management interface if it's configured for DHCP addressing (**IP Type** set to **DHCP Client**). The exceptions are AWS and Azure, where the management interface is configured as DHCP Client and it supports HA1 and HA1 Backup links.*

  - HA1: Dedicated HA1 port

    HA1 Backup: Dedicated HA1 port

    Recommendation: Enable Heartbeat Backup
  - HA1: Dedicated HA1 port

    HA1 Backup: In-band port

    Recommendation: Enable Heartbeat Backup
  - HA1: Dedicated HA1 port

    HA1 Backup: Management port

    Recommendation: Do not enable Heartbeat Backup
  - HA1: In-band port

    HA1 Backup: In-band port

    Recommendation: Enable Heartbeat Backup
  - HA1: Management port

    HA1 Backup: In-band port

    Recommendation: Do not enable Heartbeat Backup

The following table lists the HA settings that you must configure independently on each firewall. See Reference: HA Synchronization for more information about other configuration settings are not automatically synchronized between peers.

Independent Configuration Settings	PeerA	PeerB
Control Link	IP address of the HA1 link configured on this firewall (PeerA).	IP address of the HA1 link configured on this firewall (PeerB).
	For firewalls without dedicated HA ports, use the management port IP address for the control link.	
Data Link The data link information is synchronized	By default, the HA2 link uses Ethernet/Layer 2.	By default, the HA2 link uses Ethernet/Layer 2.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

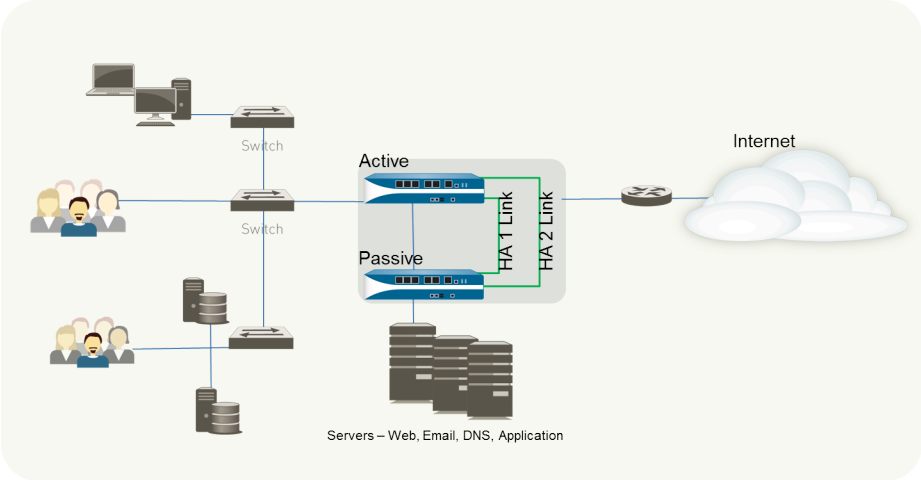
Independent Configuration Settings	PeerA	PeerB
between the firewalls after HA is enabled and the control link is established between the firewalls.	If using a Layer 3 connection, configure the IP address for the data link on this firewall (PeerA).	If using a Layer 3 connection, configure the IP address for the data link on this firewall (PeerB).
Device Priority (required, if preemption is enabled)	The firewall you plan to make active must have a lower numerical value than its peer. So, if Peer A is to function as the active firewall, keep the default value of 100 and increment the value on PeerB. If the firewalls have the same device priority value, they use the MAC address of their HA1 as the tie-breaker.	If PeerB is passive, set the device priority value to a number larger than the setting on PeerA. For example, set the value to 110.
Link Monitoring—Monitor one or more physical interfaces that handle vital traffic on this firewall and define the failure condition.	Select the physical interfaces on the firewall that you would like to monitor and define the failure condition (all or any) to trigger a failover.	Pick a similar set of physical interfaces that you would like to monitor on this firewall and define the failure condition (all or any) to trigger a failover.
Path Monitoring—Monitor one or more destination IP addresses that the firewall can use ICMP pings to ascertain responsiveness.	Define the failure condition (all or any), ping interval and the ping count. This is particularly useful for monitoring the availability of other interconnected networking devices. For example, monitor the availability of a router that connects to a server, connectivity to the server itself, or some other vital device that is in the flow of traffic. Make sure that the node/device that you are monitoring is not likely to be unresponsive, especially when it comes under load, as this could cause a a path monitoring failure and trigger a failover.	Pick a similar set of devices or destination IP addresses that can be monitored for determining the failover trigger for PeerB. Define the failure condition (all or any), ping interval and the ping count.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Configure Active/Passive HA

The following procedure shows how to configure a pair of firewalls in an active/passive deployment as depicted in the following example topology.



To configure an active/passive HA pair, first complete the following workflow on the first firewall and then repeat the steps on the second firewall.

STEP 1 | Connect the HA ports to set up a physical connection between the firewalls.

- For firewalls with dedicated HA ports, use an Ethernet cable to connect the dedicated HA1 ports and the HA2 ports on peers. Use a crossover cable if the peers are directly connected to each other.
- For firewalls without dedicated HA ports, select two data interfaces for the HA2 link and the backup HA1 link. Then, use an Ethernet cable to connect these in-band HA interfaces across both firewalls.

Use the management port for the HA1 link and ensure that the management ports can connect to each other across your network.

STEP 2 | Enable ping on the management port.

Enabling ping allows the management port to exchange heartbeat backup information.

1. Select **Device** > **Setup** > **Management** and edit the Management Interface Settings.
2. Select **Ping** as a service that is permitted on the interface.

STEP 3 | If the firewall does not have dedicated HA ports, set up the data ports to function as HA ports.

For firewalls with dedicated HA ports continue to the next step.

1. Select **Network** > **Interfaces**.
2. Confirm that the link is up on the ports that you want to use.
3. Select the interface and set **Interface Type** to **HA**.
4. Set the **Link Speed** and **Link Duplex** settings, as appropriate.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 4 | Set the HA mode and group ID.

1. Select **Device** > **High Availability** > **General** and edit the Setup section.
2. Set a **Group ID** and optionally a **Description** for the pair. The Group ID uniquely identifies each HA pair on your network. If you have multiple HA pairs that share the same broadcast domain you must set a unique Group ID for each pair.
3. Set the mode to **Active Passive**.

STEP 5 | Set up the control link connection.

This example shows an in-band port that is set to interface type HA.

For firewalls that use the management port as the control link, the IP address information is automatically pre-populated.

1. In **Device** > **High Availability** > **General**, edit the Control Link (HA1) section.
2. Select the **Port** that you have cabled for use as the HA1 link.
3. Set the **IPv4/IPv6 Address** and **Netmask**.

   If the HA1 interfaces are on separate subnets, enter the IP address of the **Gateway**. Do not add a gateway address if the firewalls are directly connected or are on the same VLAN.

STEP 6 | (Optional) Enable encryption for the control link connection.

This is typically used to secure the link if the two firewalls are not directly connected, that is if the ports are connected to a switch or a router.

1. Export the HA key from one firewall and import it into the peer firewall.
   1. Select **Device** > **Certificate Management** > **Certificates**.
   2. Select **Export HA key**. Save the HA key to a network location that the peer can access.
   3. On the peer firewall, select **Device** > **Certificate Management** > **Certificates**, and select **Import HA key** to browse to the location that you saved the key and import it in to the peer.
   4. Repeat this process on the second firewall to exchange HA keys on both devices.
2. Select **Device** > **High Availability** > **General**, edit the Control Link (HA1) section.
3. Select **Encryption Enabled**.

   If you enable encryption, after you finish configuring the HA firewalls, you can Refresh HA1 SSH Keys and Configure Key Options.

STEP 7 | Set up the backup control link connection.

1. In **Device** > **High Availability** > **General**, edit the Control Link (HA1 Backup) section.
2. Select the HA1 backup interface and set the **IPv4/IPv6 Address** and **Netmask**.

   PA-3200 Series firewalls don't support an IPv6 address for the HA1 backup control link; use an IPv4 address.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 8 | Set up the data link connection (HA2) and the backup HA2 connection between the firewalls.

1. In **Device** > **High Availability** > **General**, edit the Data Link (HA2) section.
2. Select the **Port** to use for the data link connection.
3. Select the **Transport** method. The default is **ethernet**, and will work when the HA pair is connected directly or through a switch. If you need to route the data link traffic through the network, select **IP** or **UDP** as the transport mode.
4. If you use IP or UDP as the transport method, enter the **IPv4/IPv6 Address** and **Netmask**.
5. Verify that **Enable Session Synchronization** is selected.
6. Select **HA2 Keep-alive** to enable monitoring on the HA2 data link between the HA peers. If a failure occurs based on the threshold that is set (default is 10000 ms), the defined action will occur. For active/passive configuration, a critical system log message is generated when an HA2 keep-alive failure occurs.

   You can configure the HA2 keep-alive option on both firewalls, or just one firewall in the HA pair. If the option is only enabled on one firewall, only that firewall will send the keep-alive messages. The other firewall will be notified if a failure occurs.
7. Edit the **Data Link (HA2 Backup)** section, select the interface, and add the **IPv4/IPv6 Address** and **Netmask**.

STEP 9 | Enable heartbeat backup if your control link uses a dedicated HA port or an in-band port.

You do not need to enable heartbeat backup if you are using the management port for the control link.

1. In **Device** > **High Availability** > **General**, edit the Election Settings.
2. Select **Heartbeat Backup**.

   To allow the heartbeats to be transmitted between the firewalls, you must verify that the management port across both peers can route to each other.

   Enabling heartbeat backup also allows you to prevent a split-brain situation. Split brain occurs when the HA1 link goes down causing the firewall to miss heartbeats, although the firewall is still functioning. In such a situation, each peer believes that the other is down and attempts to start services that are running, thereby causing a split brain. When the heartbeat backup link is enabled, split brain is prevented because redundant heartbeats and hello messages are transmitted over the management port.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 10 | Set the device priority and enable preemption.

This setting is only required if you wish to make sure that a specific firewall is the preferred active firewall. For information, see Device Priority and Preemption.

1. In **Device** > **High Availability** > **General**, edit the Election Settings.
2. Set the numerical value in **Device Priority**. Make sure to set a lower numerical value on the firewall that you want to assign a higher priority to.

   If both firewalls have the same device priority value, the firewall with the lowest MAC address on the HA1 control link will become the active firewall.
3. Select **Preemptive**.

   You must enable preemptive on both the active firewall and the passive firewall.

STEP 11 | (Optional) Modify the HA Timers.

By default, the HA timer profile is set to the **Recommended** profile and is suited for most HA deployments.

1. In **Device** > **High Availability** > **General**, edit the Election Settings.
2. Select the **Aggressive** profile for triggering failover faster; select **Advanced** to define custom values for triggering failover in your set up.

   *To view the preset value for an individual timer included in a profile, select **Advanced** and click **Load Recommended** or **Load Aggressive**. The preset values for your hardware model will be displayed on screen.*

STEP 12 | (Optional) Modify the link status of the HA ports on the passive firewall.

*The passive link state is **shutdown**, by default. After you enable HA, the link state for the HA ports on the active firewall will be green and those on the passive firewall will be down and display as red.*

Setting the link state to **Auto** allows for reducing the amount of time it takes for the passive firewall to take over when a failover occurs and it allows you to monitor the link state.

To enable the link status on the passive firewall to stay up and reflect the cabling status on the physical interface:

1. In **Device** > **High Availability** > **General**, edit the Active Passive Settings.
2. Set the **Passive Link State** to **Auto**.

   The auto option decreases the amount of time it takes for the passive firewall to take over when a failover occurs.

   Although the interface displays green (as cabled and up) it continues to discard all traffic until a failover is triggered.

   When you modify the passive link state, make sure that the adjacent devices do not forward traffic to the passive firewall based only on the link status of the firewall.

https://t.me/learningnets

STEP 13 | Enable HA.

1. Select **Device** > **High Availability** > **General** and edit the Setup section.
2. Select **Enable HA**.
3. Select **Enable Config Sync**. This setting enables the synchronization of the configuration settings between the active and the passive firewall.
4. Enter the IP address assigned to the control link of the peer in **Peer HA1 IP Address**.

   For firewalls without dedicated HA ports, if the peer uses the management port for the HA1 link, enter the management port IP address of the peer.
5. Enter the **Backup HA1 IP Address**.

STEP 14 | (Optional) Enable LACP and LLDP Pre-Negotiation for Active/Passive HA for faster failover if your network uses LACP or LLDP.



Enable LACP and LLDP before configuring HA pre-negotiation for the protocol if you want pre-negotiation to function in active mode.

1. Ensure that in Step 12 you set the link state to **Auto**.
2. Select **Network** > **Interfaces** > **Ethernet**.
3. To enable LACP active pre-negotiation:
   1. Select an AE interface in a Layer 2 or Layer 3 deployment.
   2. Select the **LACP** tab.
   3. Select **Enable in HA Passive State**.
   4. Click **OK**.

   *You cannot also select **Same System MAC Address for Active-Passive HA** because pre-negotiation requires unique interface MAC addresses on the active and passive firewalls.*
4. To enable LACP passive pre-negotiation:
   1. Select an Ethernet interface in a virtual wire deployment.
   2. Select the **Advanced** tab.
   3. Select the **LACP** tab.
   4. Select **Enable in HA Passive State**.
   5. Click **OK**.
5. To enable LLDP active pre-negotiation:
   1. Select an Ethernet interface in a Layer 2, Layer 3, or virtual wire deployment.
   2. Select the **Advanced** tab.
   3. Select the **LLDP** tab.
   4. Select **Enable in HA Passive State**.
   5. Click **OK**.

   If you want to allow LLDP passive pre-negotiation for a virtual wire deployment, perform Step 14.e but do not enable LLDP itself.

https://t.me/learningnets

STEP 15 | Save your configuration changes.

Click **Commit**.

STEP 16 | After you finish configuring both firewalls, verify that the firewalls are paired in active/passive HA.

1. Access the **Dashboard** on both firewalls, and view the High Availability widget.
2. On the active firewall, click the **Sync to peer** link.
3. Confirm that the firewalls are paired and synced, as shown as follows:
   - On the passive firewall: the state of the local firewall should display **passive** and the Running Config should show as **synchronized**.
   - On the active firewall: The state of the local firewall should display **active** and the Running Config should show as **synchronized**.

Define HA Failover Conditions

Perform the following task to use link monitoring or path monitoring to define Failover conditions and thus establish what will cause a firewall in an HA pair to fail over, an event where the task of securing traffic passes from the previously active firewall to its HA peer. The HA Overview describes conditions that cause a failover.

You can monitor multiple IP path groups per virtual router, VLAN, or virtual wire. You can enable each path group with one or more IP addresses and give each its own peer failure conditions. Additionally, you can set these failure conditions at both the path-group level and the broader virtual router or VLAN or virtual wire group level using “any” or “all” fail checks to determine the status of the active firewall.

When you upgrade to PAN-OS 10.0, the firewall automatically transfers your currently monitored destination IP addresses to a newly created destination group and gives that group a default path-monitoring name. The new destination group retains your previous failover condition at the path-group level.

Ensure that you delete all VLAN path monitoring configurations in active/active HA before you upgrade to PAN-OS 10.1 because VLAN path monitoring is not compatible with active/active HA pairing in PAN-OS 10.0; retaining an earlier active/active HA configuration results in an autocommit failure.

Before you enable path monitoring, you must set up your virtual routers, VLAN, or virtual wires or a combination of these logical networking components. Path monitoring in virtual routers and virtual wires is compatible with both active/active and active/passive HA deployments; however, path monitoring in VLANs is supported only on active/passive pairs.

Before you enable path monitoring, you must also:

- Check reachability for destination IP groups in your virtual routers.
- Ensure that the VLANs (for which you intend to enable path monitoring) include configured interfaces.
- Obtain the source IP address that you will use to receive pings from the appropriate destination IP address.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

If you are using SNMPv3 to monitor the firewalls, note that the SNMPv3 Engine ID is unique to each firewall; the EngineID is not synchronized between the HA pair and, therefore, allows you to independently monitor each firewall in the HA pair. For information on setting up SNMP, see Forward Traps to an SNMP Manager. Because the EngineID is generated using the firewall serial number, on the VM-Series firewall you must apply a valid license in order to obtain a unique EngineID for each firewall.

STEP 1 | To configure HA link monitoring, specify a group of physical interfaces for the firewall to monitor (link up or link down).

1. Select **Device** > **High Availability** > **Link and Path Monitoring**.
2. In the Link Monitoring section, **Add** a link group by **Name**.
3. Select **Enabled** to enable the link group.
4. Select the **Failure Condition** for the interfaces in the link group: **Any** (default) or **All**.
5. **Add** the **Interface**(s) to monitor.
6. Click **OK**.

STEP 2 | (Optional) Modify the failure condition for the set of Link Groups configured on the firewall.

By default, the firewall triggers a failover when any monitored Link Group fails.

1. Edit the **Link Monitoring** section.
2. Set the **Failure Condition** to **Any** (default) or **All**.
3. Click **OK**.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 3 | To configure HA path monitoring for a virtual wire, VLAN, or virtual router, specify the destination IP addresses that the firewall will ping to verify network connectivity.

1. In the Path Monitoring section, select **Add Virtual Wire Path**, **Add VLAN Path**, or **Add Virtual Router Path**.
2. Enter a **Name** for the virtual wire, VLAN, or virtual router path group.
3. (Virtual Wire Path or VLAN Path only) Enter the **Source IP** address to use to ping the destination IP address through the virtual wire or VLAN.
4. Select **Enabled** to enable the path group.
5. Select the **Failure Condition** that results in a failure for this path group: **Any** (default) to issue a failure when one or more Destination IP groups in this path group fail or **All** to issue a failure when all Destination IP groups in this path group fail.
6. Enter the **Ping Interval** in milliseconds; the interval between ICMP messages sent to the Destination IP address (range is 200 to 60,000; default is 200).
7. Enter the **Ping Count** of pings that must fail before declaring a failure (range is 3 to 10; default is 10).
8. **Add** and enter a **Destination IP Group** name.
9. **Add** one or more **Destination IP** addresses to ping.
10. Select **Enabled** to enable path monitoring for the Destination IP group.
11. Select the **Failure Condition** that results in a failure for this Destination IP group: **Any** (default) to issue a failure when one or more listed IP addresses is unreachable or **All** to issue a failure when all listed IP addresses are unreachable.
12. Click **OK** twice.
13. (Panorama only) Select the appropriate Panorama template to push the path monitoring configuration to your appliance.



You can push HA path monitoring for a virtual wire, VLAN, or virtual router only to firewalls running PAN-OS 10.0 or a later releases. If you try to push the configuration to firewalls running a release earlier than PAN-OS 10.0 (such as 9.1.x or 9.0.x), the commit may fail or the commit may remove destination IP addresses from the path group.

Only HA Path Groups containing one Destination IP Group are supported for managed firewalls running PAN-OS 9.1 and earlier releases.



To manage the destination IP addresses from Panorama for managed firewalls running different PAN-OS releases, create a separate template for managed firewalls running PAN-OS 10.0 and later releases and a separate template for managed firewalls running PAN-OS 9.1 and earlier releases. This allows you to more accurately control the destination IP address configuration if you created multiple destination IP groups and ensures your managed firewall successfully fails over.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 4 | (Optional) Modify the failure condition for the set of Path Groups configured on the firewall.

By default, the firewall triggers a failover when any monitored Path Group fails.

1. Edit the **Path Monitoring** section.
2. Select **Enabled** to enable path monitoring on the appliance.
3. Set the **Failure Condition** to **Any** (default) to issue a failure for this firewall when one or more monitored virtual routers, VLANs, or virtual wires is down. Select **All** to issue a failure for this firewall when all monitored virtual routers, VLANs, or virtual wires are down.
4. Click **OK**.

STEP 5 | **Commit**.

Verify Failover

To test that your HA configuration works properly, trigger a manual failover and verify that the firewalls transition states successfully.

STEP 1 | Suspend the active firewall.

Select **Device** > **High Availability** > **Operational Commands** and click the **Suspend local device** link.

STEP 2 | Verify that the passive firewall has taken over as active.

On the **Dashboard**, verify that the state of the passive firewall changes to **active** in the High Availability widget.

STEP 3 | Restore the suspended firewall to a functional state. Wait for a couple of minutes, and then verify that preemption has occurred, if **Preemptive** is enabled.

1. On the firewall you previously suspended, select **Device** > **High Availability** > **Operational Commands** and click the **Make local device functional** link.
2. In the High Availability widget on the **Dashboard**, confirm that the firewall has taken over as the active firewall and that the peer is now in a passive state.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Set Up Active/Active HA

- Prerequisites for Active/Active HA
- Configure Active/Active HA
- Determine Your Active/Active Use Case

Prerequisites for Active/Active HA

To set up active/active HA on your firewalls, you need a pair of firewalls that meet the following requirements:

- **The same model**—The firewalls in the pair must be of the same hardware model.
- **The same PAN-OS version**—The firewalls must be running the same PAN-OS version and must each be up-to-date on the application, URL, and threat databases.
- **The same multi virtual system capability**—Both firewalls must have **Multi Virtual System Capability** either enabled or not enabled. When enabled, each firewall requires its own multiple virtual systems licenses.
- **The same type of interfaces**—Dedicated HA links, or a combination of the management port and in-band ports that are set to *interface type* HA.
  - The HA interfaces must be configured with static IP addresses only, not IP addresses obtained from DHCP (except AWS can use DHCP addresses). Determine the IP address for the HA1 (control) connection between the HA peers. The HA1 IP address for the peers must be on the same subnet if they are directly connected or are connected to the same switch.

    For firewalls without dedicated HA ports, you can use the management port for the control connection. Using the management port provides a direct communication link between the management planes on both firewalls. However, because the management ports will not be directly cabled between the peers, make sure that you have a route that connects these two interfaces across your network.
  - If you use Layer 3 as the transport method for the HA2 (data) connection, determine the IP address for the HA2 link. Use Layer 3 only if the HA2 connection must communicate over a routed network. The IP subnet for the HA2 links must not overlap with that of the HA1 links or with any other subnet assigned to the data ports on the firewall.
  - Each firewall needs a dedicated interface for the HA3 link. The PA-7000 Series firewalls use the HSCI port for HA3. The PA-5200 Series firewalls can use the HSCI port for HA3 or you can configure aggregate interfaces on the dataplane ports for HA3 for redundancy. On the remaining platforms, you can configure aggregate interfaces on dataplane ports as the HA3 link for redundancy.
- **The same set of licenses**—Licenses are unique to each firewall and cannot be shared between the firewalls. Therefore, you must license both firewalls identically. If both firewalls do not have

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

an identical set of licenses, they cannot synchronize configuration information and maintain parity for a seamless failover.

If you have an existing firewall and you want to add a new firewall for HA purposes and the new firewall has an existing configuration, it is recommended that you Reset the Firewall to Factory Default Settings on the new firewall. This will ensure that the new firewall has a clean configuration. After HA is configured, you will then sync the configuration on the primary firewall to the newly introduced firewall with the clean config. You will also have to configure local IP addresses.

Configure Active/Active HA

The following procedure describes the basic workflow for configuring your firewalls in an active/active configuration. However, before you begin, Determine Your Active/Active Use Case for configuration examples more tailored to your specific network environment.

If you have a switch located between your HA firewalls, the switch ports that connect the HA3 link must support jumbo frames to handle the overhead associated with the MAC-in-MAC encapsulation on the HA3 link.

To configure active/active, first complete the following steps on one peer and then complete them on the second peer, ensuring that you set the Device ID to different values (0 or 1) on each peer.

STEP 1 | Connect the HA ports to set up a physical connection between the firewalls.

For each use case, the firewalls could be any hardware model; choose the HA3 step that corresponds with your model.

- For firewalls with dedicated HA ports, use an Ethernet cable to connect the dedicated HA1 ports and the HA2 ports on peers. Use a crossover cable if the peers are directly connected to each other.
- For firewalls without dedicated HA ports, select two data interfaces for the HA2 link and the backup HA1 link. Then, use an Ethernet cable to connect these in-band HA interfaces across both firewalls. Use the management port for the HA1 link and ensure that the management ports can connect to each other across your network.
- For HA3:
  - On PA-7000 Series firewalls, connect the High Speed Chassis Interconnect (HSCI-A) on the first chassis to the HSCI-A on the second chassis, and the HSCI-B on the first chassis to the HSCI-B on the second chassis.
  - On PA-5200 Series firewalls (which have one HSCI port), connect the HSCI port on the first chassis to the HSCI port on the second chassis. You can also use data ports for HA3 on PA-5200 Series firewalls.
  - On PA-3200 Series firewalls (which have one HSCI port), connect the HSCI port on the first chassis to the HSCI port on the second chassis.
  - On any other hardware model, use dataplane interfaces for HA3.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 2 | Enable ping on the management port.

Enabling ping allows the management port to exchange heartbeat backup information.

1. In **Device** > **Setup** > **Management**, edit Management Interface Settings.
2. Select **Ping** as a service that is permitted on the interface.

STEP 3 | If the firewall does not have dedicated HA ports, set up the data ports to function as HA ports.

For firewalls with dedicated HA ports continue to the next step.

1. Select **Network** > **Interfaces**.
2. Confirm that the link is up on the ports that you want to use.
3. Select the interface and set **Interface Type** to **HA**.
4. Set the **Link Speed** and **Link Duplex** settings, as appropriate.

STEP 4 | Enable active/active HA and set the group ID.

1. In **Device** > **High Availability** > **General**, edit Setup.
2. Select **Enable HA**.
3. Enter a **Group ID**, which must be the same for both firewalls. The firewall uses the Group ID to calculate the virtual MAC address (range is 1-63).
4. (Optional) Enter a **Description**.
5. For **Mode**, select **Active Active**.

STEP 5 | Set the Device ID, enable synchronization, and identify the control link on the peer firewall

1. In **Device** > **High Availability** > **General**, edit Setup.
2. Select **Device ID** as follows:
   - When configuring the first peer, set the **Device ID** to **0**.
   - When configuring the second peer, set the **Device ID** to **1**.
3. Select **Enable Config Sync**. This setting is required to synchronize the two firewall configurations (enabled by default).
4. Enter the **Peer HA1 IP Address**, which is the IP address of the HA1 control link on the peer firewall.
5. (Optional) Enter a **Backup Peer HA1 IP Address**, which is the IP address of the backup control link on the peer firewall.
6. Click **OK**.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 6 | Determine whether or not the firewall with the lower Device ID preempts the active-primary firewall upon recovery from a failure.

1. In **Device** > **High Availability** > **General**, edit Election Settings.
2. Select **Preemptive** to cause the firewall with the lower Device ID to automatically resume active-primary operation after either firewall recovers from a failure. Both firewalls must have **Preemptive** selected for preemption to occur.

   Leave **Preemptive** unselected if you want the active-primary role to remain with the current firewall until you manually make the recovered firewall the active-primary firewall.

STEP 7 | Enable heartbeat backup if your control link uses a dedicated HA port or an in-band port.

You need not enable heartbeat backup if you are using the management port for the control link.

1. In **Device** > **High Availability** > **General**, edit Election Settings.
2. Select **Heartbeat Backup**.

   To allow the heartbeats to be transmitted between the firewalls, you must verify that the management port across both peers can route to each other.

   Enabling heartbeat backup allows you to prevent a split-brain situation. Split brain occurs when the HA1 link goes down, causing the firewall to miss heartbeats, although the firewall is still functioning. In such a situation, each peer believes the other is down and attempts to start services that are running, thereby causing a split brain. Enabling heartbeat backup prevents split brain because redundant heartbeats and hello messages are transmitted over the management port.

STEP 8 | (Optional) Modify the HA Timers.

By default, the HA timer profile is set to the **Recommended** profile and is suited for most HA deployments.

1. In **Device** > **High Availability** > **General**, edit Election Settings.
2. Select **Aggressive** to trigger faster failover. Select **Advanced** to define custom values for triggering failover in your setup.

   *To view the preset value for an individual timer included in a profile, select **Advanced** and click **Load Recommended** or **Load Aggressive**. The preset values for your hardware model will be displayed on screen.*

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 9 | Set up the control link connection.

This example uses an in-band port that is set to interface type HA.

For firewalls that use the management port as the control link, the IP address information is automatically pre-populated.

1. In **Device** > **High Availability** > **General**, edit Control Link (HA1).
2. Select the **Port** that you have cabled for use as the HA1 link.
3. Set the **IPv4/IPv6 Address** and **Netmask**.

   If the HA1 interfaces are on separate subnets, enter the IP address of the **Gateway**. Do not add a gateway address if the firewalls are directly connected.

STEP 10 | (Optional) Enable encryption for the control link connection.

This is typically used to secure the link if the two firewalls are not directly connected, that is if the ports are connected to a switch or a router.

1. Export the HA key from one firewall and import it into the peer firewall.
   1. Select **Device** > **Certificate Management** > **Certificates**.
   2. Select **Export HA key**. Save the HA key to a network location that the peer can access.
   3. On the peer firewall, select **Device** > **Certificate Management** > **Certificates**, and select **Import HA key** to browse to the location that you saved the key and import it in to the peer.
2. In **Device** > **High Availability** > **General**, edit the Control Link (HA1).
3. Select **Encryption Enabled**.

   If you enable encryption, after you finish configuring the HA firewalls, you can Refresh HA1 SSH Keys and Configure Key Options.

STEP 11 | Set up the backup control link connection.

1. In **Device** > **High Availability** > **General**, edit Control Link (HA1 Backup).
2. Select the HA1 backup interface and set the **IPv4/IPv6 Address** and **Netmask**.

   PA-3200 Series firewalls don't support an IPv6 address for the HA1 backup control link; use an IPv4 address.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 12 | Set up the data link connection (HA2) and the backup HA2 connection between the firewalls.

1. In **Device** > **High Availability** > **General**, edit Data Link (HA2).
2. Select the **Port** to use for the data link connection.
3. Select the **Transport** method. The default is **ethernet**, and will work when the HA pair is connected directly or through a switch. If you need to route the data link traffic through the network, select **IP** or **UDP** as the transport mode.
4. If you use IP or UDP as the transport method, enter the **IPv4/IPv6 Address** and **Netmask**.
5. Verify that **Enable Session Synchronization** is selected.
6. Select **HA2 Keep-alive** to enable monitoring on the HA2 data link between the HA peers. If a failure occurs based on the threshold that is set (default is 10000 ms), the defined action will occur. When an HA2 Keep-alive failure occurs, the system either generates a critical system log message or causes a split dataplane depending on your configuration.

   You can configure the HA2 Keep-alive option on both firewalls, or just one firewall in the HA pair. If the option is only enabled on one firewall, only that firewall sends the Keep-alive messages. The other firewall is notified if a failure occurs.

   A split dataplane causes the dataplanes of both peers to operate independently while leaving the high-available state as Active-Primary and Active-Secondary. If only one firewall is configured to split dataplane, then split dataplane applies to the other device as well.
7. Edit the **Data Link (HA2 Backup)** section, select the interface, and add the **IPv4/IPv6 Address** and **Netmask**.
8. Click **OK**.

STEP 13 | Configure the HA3 link for packet forwarding.

1. In **Device** > **High Availability** > **Active/Active Config**, edit Packet Forwarding.
2. For **HA3 Interface**, select the interface you want to use to forward packets between active/active HA peers. It must be a dedicated interface capable of Layer 2 transport and set to **Interface Type HA**.
3. Select **VR Sync** to force synchronization of all virtual routers configured on the HA peers. Select when the virtual router is not configured for dynamic routing protocols. Both peers must be connected to the same next-hop router through a switched network and must use static routing only.
4. Select **QoS Sync** to synchronize the QoS profile selection on all physical interfaces. Select when both peers have similar link speeds and require the same QoS profiles on all physical interfaces. This setting affects the synchronization of QoS settings on the **Network** tab. QoS policy is synchronized regardless of this setting.

STEP 14 | (Optional) Modify the Tentative Hold time.

1. In **Device** > **High Availability** > **Active/Active Config**, edit Packet Forwarding.
2. For **Tentative Hold Time (sec)**, enter the number of seconds that a firewall stays in Tentative state after it recovers post-failure (range is 10-600, default is 60).

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 15 | Configure Session Owner and Session Setup.

1. In **Device** > **High Availability** > **Active/Active Config**, edit Packet Forwarding.
2. For **Session Owner Selection**, select one of the following:
   - **First Packet**—The firewall that receives the first packet of a new session is the session owner (recommended setting). This setting minimizes traffic across HA3 and load shares traffic across peers.
   - **Primary Device**—The firewall that is in active-primary state is the session owner.
3. For **Session Setup**, select one of the following:
   - **IP Modulo**—The firewall performs an XOR operation on the source and destination IP addresses from the packet and based on the result, the firewall chooses which HA peer will set up the session.
   - **Primary Device**—The active-primary firewall sets up all sessions.
   - **First Packet**—The firewall that receives the first packet of a new session performs session setup (recommended setting).

     Start with First Packet for Session Owner and Session Setup, and then based on load distribution, you can change to one of the other options.
   - **IP Hash**—The firewall uses a hash of either the source IP address or a combination of the source and destination IP addresses to distribute session setup responsibilities.
4. Click **OK**.

STEP 16 | Configure an HA virtual address.

You need a virtual address to use a Floating IP Address and Virtual MAC Address or ARP Load-Sharing.

1. In **Device** > **High Availability** > **Active/Active Config**, **Add** a Virtual Address.
2. Enter or select an **Interface**.
3. Select the **IPv4** or **IPv6** tab and click **Add**.
4. Enter an **IPv4 Address** or **IPv6 Address**.
5. For **Type**:
   - Select **Floating** to configure the virtual IP address to be a floating IP address.
   - Select **ARP Load Sharing** to configure the virtual IP address to be a shared IP address and skip to Configure ARP Load-Sharing.

STEP 17 | Configure the floating IP address.

1. Do not select **Floating IP bound to the Active-Primary device** unless you want the active/active HA pair to behave like an active/passive HA pair.
2. For **Device 0 Priority** and **Device 1 Priority**, enter a priority for the firewall configured with Device ID 0 and Device ID 1, respectively. The relative priorities determine which

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

peer owns the floating IP address you just configured (range is 0-255). The firewall with the lowest priority value (highest priority) owns the floating IP address.

3. Select **Failover address if link state is down** to cause the firewall to use the failover address when the link state on the interface is down.
4. Click **OK**.

STEP 18 | Configure ARP Load-Sharing.

The device selection algorithm determines which HA firewall responds to the ARP requests to provide load sharing.

1. For **Device Selection Algorithm**, select one of the following:
   - **IP Modulo**—The firewall that will respond to ARP requests is based on the parity of the ARP requester's IP address.
   - **IP Hash**—The firewall that will respond to ARP requests is based on a hash of the ARP requester's IP address.
2. Click **OK**.

STEP 19 | Define HA Failover Conditions.

STEP 20 | **Commit** the configuration.

Determine Your Active/Active Use Case

Determine which type of use case you have and then select the corresponding procedure to configure active/active HA.

If you are using Route-Based Redundancy, Floating IP Address and Virtual MAC Address, or ARP Load-Sharing, select the corresponding procedure:

- Use Case: Configure Active/Active HA with Route-Based Redundancy
- Use Case: Configure Active/Active HA with Floating IP Addresses
- Use Case: Configure Active/Active HA with ARP Load-Sharing

If you want a Layer 3 active/active HA deployment that behaves like an active/passive deployment, select the following procedure:

- Use Case: Configure Active/Active HA with Floating IP Address Bound to Active-Primary Firewall

If you are configuring NAT in Active/Active HA Mode, see the following procedures:

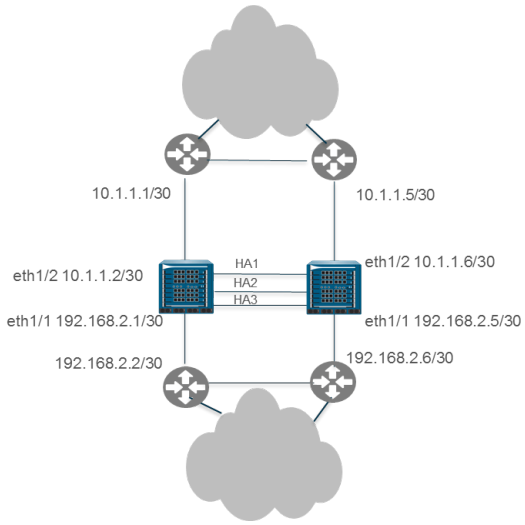
- Use Case: Configure Active/Active HA with Source DIPP NAT Using Floating IP Addresses
- Use Case: Configure Separate Source NAT IP Address Pools for Active/Active HA Firewalls
- Use Case: Configure Active/Active HA for ARP Load-Sharing with Destination NAT
- Use Case: Configure Active/Active HA for ARP Load-Sharing with Destination NAT in Layer 3

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Use Case: Configure Active/Active HA with Route-Based Redundancy

The following Layer 3 topology illustrates two PA-7050 firewalls in an active/active HA environment that use Route-Based Redundancy. The firewalls belong to an OSPF area. When a link or firewall fails, OSPF handles the redundancy by redirecting traffic to the functioning firewall.



STEP 1 | Configure Active/Active HA.

Perform Step 1 through Step 15.

STEP 2 | Configure OSPF.

See OSPF.

STEP 3 | Define HA failover conditions.

Define HA Failover Conditions.

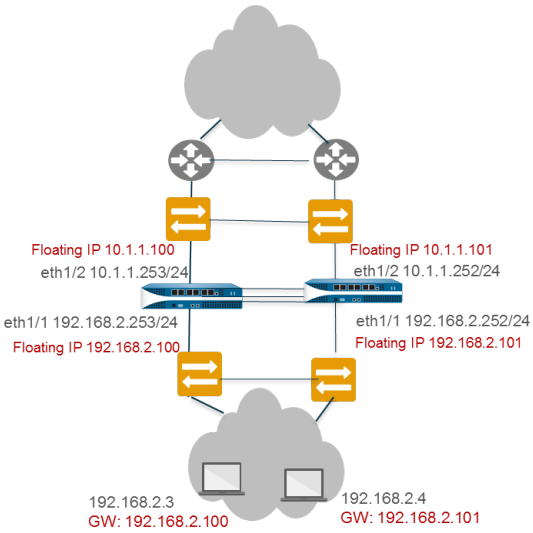
STEP 4 | **Commit** the configuration.

STEP 5 | Configure the peer firewall in the same way, except in Step 5, if you selected Device ID **0** for the first firewall, select Device ID **1** for the peer firewall.

Use Case: Configure Active/Active HA with Floating IP Addresses

In this Layer 3 interface example, the HA firewalls connect to switches and use floating IP addresses to handle link or firewall failures. The end hosts are each configured with a gateway, which is the floating IP address of one of the HA firewalls. See Floating IP Address and Virtual MAC Address.

https://t.me/learningnets



STEP 1 | Configure Active/Active HA.

Perform Step 1 through Step 15.

STEP 2 | Configure an HA virtual address.

You need a virtual address to use a Floating IP Address and Virtual MAC Address.

1. In **Device** > **High Availability** > **Active/Active Config**, **Add** a Virtual Address.
2. Enter or select an **Interface**.
3. Select the **IPv4** or **IPv6** tab and click **Add**.
4. Enter an **IPv4 Address** or **IPv6 Address**.
5. For **Type**, select **Floating** to configure the virtual IP address to be a floating IP address.

STEP 3 | Configure the floating IP address.

1. Do not select **Floating IP bound to the Active-Primary device**.
2. For **Device 0 Priority** and **Device 1 Priority**, enter a priority for the firewall configured with Device ID 0 and Device ID 1, respectively. The relative priorities determine which peer owns the floating IP address you just configured (range is 0 to 255). The firewall with the lowest priority value (highest priority) owns the floating IP address.
3. Select **Failover address if link state is down** to cause the firewall to use the failover address when the link state on the interface is down.
4. Click **OK**.

STEP 4 | Enable jumbo frames on firewalls other than PA-7000 Series firewalls.

Perform Step 19 of Configure Active/Active HA.

STEP 5 | Define HA Failover Conditions

STEP 6 | **Commit** the configuration.

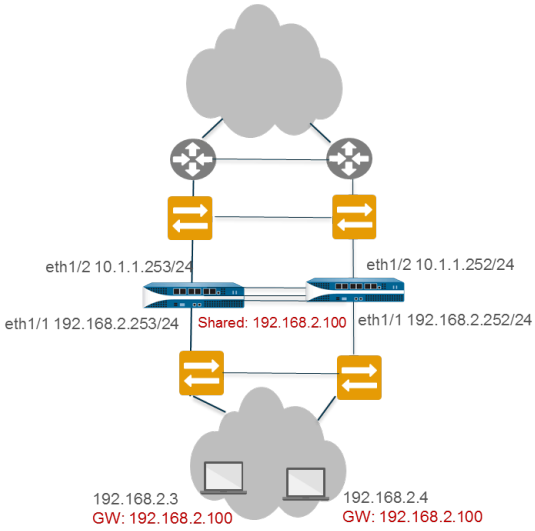
https://t.me/learningnets

STEP 7 | Configure the peer firewall in the same way, except selecting a different Device ID.

For example, if you selected Device ID **0** for the first firewall, select Device ID **1** for the peer firewall.

Use Case: Configure Active/Active HA with ARP Load-Sharing

In this example, hosts in a Layer 3 deployment need gateway services from the HA firewalls. The firewalls are configured with a single shared IP address, which allows ARP Load-Sharing. The end hosts are configured with the same gateway, which is the shared IP address of the HA firewalls.



STEP 1 | Perform Step 1 through Step 15 of Configure Active/Active HA.

STEP 2 | Configure an HA virtual address.

The virtual address is the shared IP address that allows ARP Load-Sharing.

1. Select **Device** > **High Availability** > **Active/Active Config** > **Virtual Address** and click **Add**.
2. Enter or select an **Interface**.
3. Select the **IPv4** or **IPv6** tab and click **Add**.
4. Enter an **IPv4 Address** or **IPv6 Address**.
5. For **Type**, select **ARP Load Sharing**, which allows both peers to use the virtual IP address for ARP Load-Sharing.

https://t.me/learningnets

STEP 3 | Configure ARP Load-Sharing.

The device selection algorithm determines which HA firewall responds to the ARP requests to provide load sharing.

1. For **Device Selection Algorithm**, select one of the following:
   - **IP Modulo**—The firewall that will respond to ARP requests is based on the parity of the ARP requester's IP address.
   - **IP Hash**—The firewall that will respond to ARP requests is based on a hash of the ARP requester's IP address.
2. Click **OK**.

STEP 4 | Enable jumbo frames on firewalls other than PA-7000 Series firewalls.

STEP 5 | Define HA Failover Conditions

STEP 6 | **Commit** the configuration.

STEP 7 | Configure the peer firewall in the same way, except selecting a different Device ID.

For example, if you selected Device ID **0** for the first firewall, select Device ID **1** for the peer firewall.

Use Case: Configure Active/Active HA with Floating IP Address Bound to Active-Primary Firewall

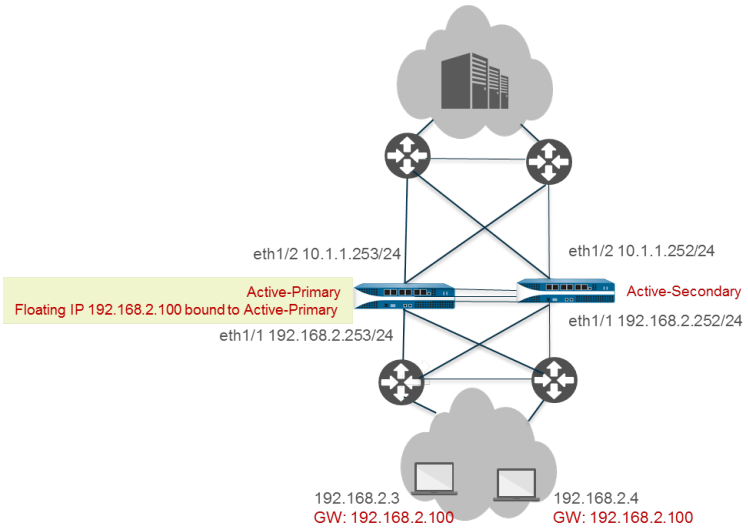
In mission-critical data centers, you may want both Layer 3 HA firewalls to participate in path monitoring so that they can detect path failures upstream from both firewalls. Additionally, you prefer to control if and when the floating IP address returns to the recovered firewall after it comes back up, rather than the floating IP address returning to the device ID to which it is bound. (That default behavior is described in Floating IP Address and Virtual MAC Address.)

In this use case, you control when the floating IP address and therefore the active-primary role move back to a recovered HA peer. The active/active HA firewalls share a single floating IP address that you bind to whichever firewall is in the active-primary state. With only one floating IP address, network traffic flows predominantly to a single firewall, so this active/active deployment functions like an active/passive deployment.

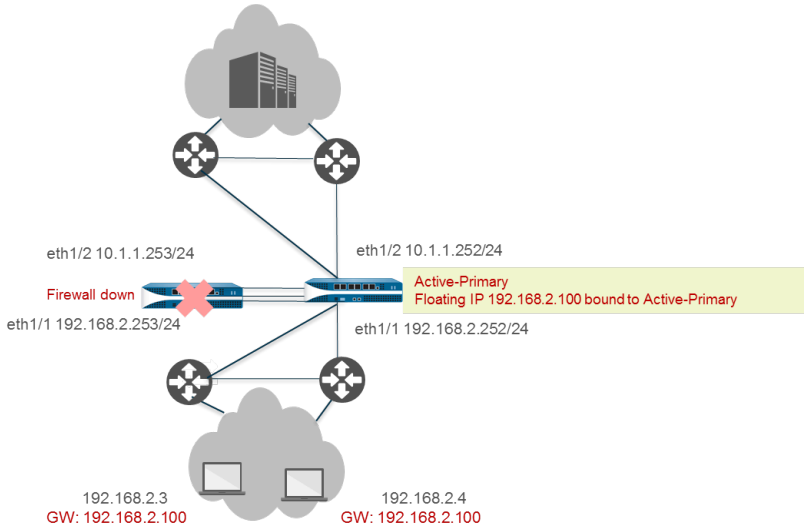
In this use case, Cisco Nexus 7010 switches with virtual PortChannels (vPCs) operating in Layer 3 connect to the firewalls. You must configure the Layer 3 switches (router peers) north and south of the firewalls with a route preference to the floating IP address. That is, you must design your network so the route tables of the router peers have the best path to the floating IP address. This example uses static routes with the proper metrics so that the route to the floating IP address uses a lower metric (the route to the floating IP address is preferred) and receives the traffic. An alternative to using static routes would be to design the network to redistribute the floating IP address into the OSPF routing protocol (if you are using OSPF).

The following topology illustrates the floating IP address bound to the active-primary firewall, which is initially Peer A, the firewall on the left.

https://t.me/learningnets



Upon a failover, when the active-primary firewall (Peer A) goes down and the active-secondary firewall (Peer B) takes over as the active-primary peer, the floating IP address moves to Peer B (shown in the following figure). Peer B remains the active-primary firewall and traffic continues to go to Peer B, even when Peer A recovers and becomes the active-secondary firewall. You decide if and when to make Peer A the active-primary firewall again.



Binding the floating IP address to the active-primary firewall provides you with more control over how the firewalls determine floating IP address ownership as they move between various HA Firewall States. The following advantages result:

- You can have an active/active HA configuration for path monitoring out of both firewalls, but have the firewalls function like an active/passive HA configuration because traffic directed to the floating IP address always goes to the active-primary firewall.

When you disable preemption on both firewalls, you have the following additional benefits:

- The floating IP address does not move back and forth between HA firewalls if the active-secondary firewall flaps up and down.
- You can review the functionality of the recovered firewall and the adjacent components before manually directing traffic to it again, which you can do at a convenient down time.

https://t.me/learningnets

- You have control over which firewall owns the floating IP address so that you keep all flows of new and existing sessions on the active-primary firewall, thereby minimizing traffic on the HA3 link.



- *We strongly recommended you configure HA link monitoring on the interface(s) that support the floating IP address(es) to allow each HA peer to quickly detect a link failure and fail over to its peer. Both HA peers must have link monitoring for it to function.*
- *We strongly recommend you configure HA path monitoring to notify each HA peer when a path has failed so a firewall can fail over to its peer. Because the floating IP address is always bound to the active-primary firewall, the firewall cannot automatically fail over to the peer when a path goes down and path monitoring is not enabled.*



You cannot configure NAT for a floating IP address that is bound to an active-primary firewall.

STEP 1 | Perform Step 1 through Step 5 of Configure Active/Active HA.

STEP 2 | (Optional) Disable preemption.



Disabling preemption allows you full control over when the recovered firewall becomes the active-primary firewall.

1. In **Device** > **High Availability** > **General**, edit the Election Settings.
2. Clear **Preemptive** if it is enabled.
3. Click **OK**.

STEP 3 | Perform Step 7 through Step 14 of Configure Active/Active HA.

STEP 4 | Configure Session Owner and Session Setup.

1. In **Device** > **High Availability** > **Active/Active Config**, edit Packet Forwarding.
2. For **Session Owner Selection**, we recommend you select **Primary Device**. The firewall that is in active-primary state is the session owner.

   Alternatively, for **Session Owner Selection** you can select **First Packet** and then for **Session Setup**, select **Primary Device** or **First Packet**.
3. For **Session Setup**, select **Primary Device**—The active-primary firewall sets up all sessions. This is the recommended setting if you want your active/active configuration to behave like an active/passive configuration because it keeps all activity on the active-primary firewall.

   *You must also engineer your network to eliminate the possibility of asymmetric traffic going to the HA pair. If you don't do so and traffic goes to the active-secondary firewall, setting **Session Owner Selection** and **Session Setup** to **Primary Device** causes the traffic to traverse HA3 to get to the active-primary firewall for session ownership and session setup.*
4. Click **OK**.

https://t.me/learningnets

STEP 5 | Configure an HA virtual address.

1. Select **Device** > **High Availability** > **Active/Active Config** > **Virtual Address** and click **Add**.
2. Enter or select an **Interface**.
3. Select the **IPv4** or **IPv6** tab and **Add** an **IPv4 Address** or **IPv6 Address**.
4. For **Type**, select **Floating**, which configures the virtual IP address to be a floating IP address.
5. Click **OK**.

STEP 6 | Bind the floating IP address to the active-primary firewall.

1. Select **Floating IP bound to the Active-Primary device**.
2. Select **Failover address if link state is down** to cause the firewall to use the failover address when the link state on the interface is down.
3. Click **OK**.

STEP 7 | Enable jumbo frames on firewalls other than PA-7000 Series firewalls.

STEP 8 | **Commit** the configuration.

STEP 9 | Configure the peer firewall in the same way, except selecting a different Device ID.

For example, if you selected Device ID **0** for the first firewall, select Device ID **1** for the peer firewall.

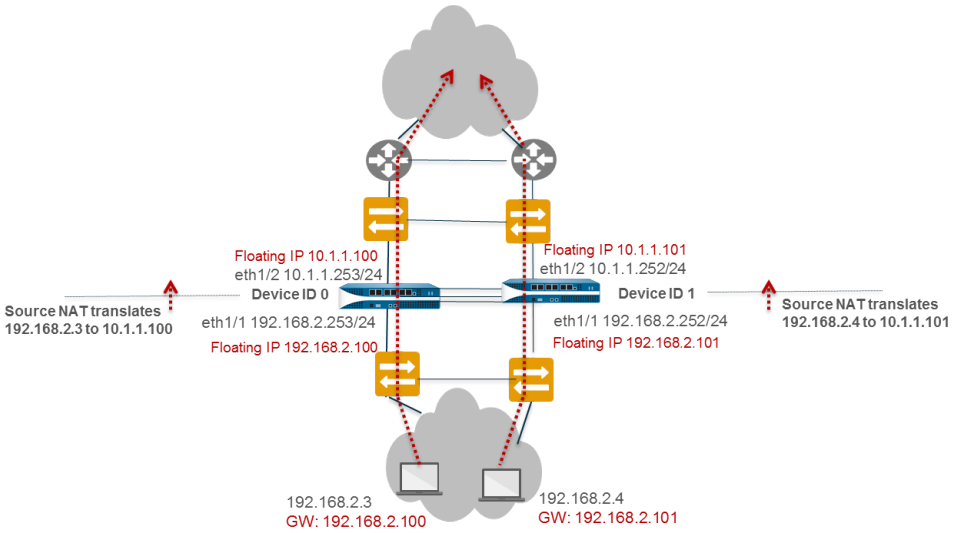
Use Case: Configure Active/Active HA with Source DIPP NAT Using Floating IP Addresses

This Layer 3 interface example uses source NAT in Active/Active HA Mode. The Layer 2 switches create broadcast domains to ensure users can reach everything north and south of the firewalls.

PA-3050-1 has Device ID 0 and its HA peer, PA-3050-2, has Device ID 1. In this use case, NAT translates the source IP address and port number to the floating IP address configured on the egress interface. Each host is configured with a default gateway address, which is the floating IP address on Ethernet1/1 of each firewall. The configuration requires two source NAT rules, one bound to each Device ID, although you configure both NAT rules on a single firewall and they are synchronized to the peer firewall.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets



STEP 1 | On PA-3050-2 (Device ID 1), perform Step 1 through Step 3 of Configure Active/Active HA.

STEP 2 | Enable active/active HA.

1. In **Device** > **High Availability** > **General**, edit Setup.
2. Select **Enable HA**.
3. Enter a **Group ID**, which must be the same for both firewalls. The firewall uses the Group ID to calculate the virtual MAC address (range is 1-63).
4. For **Mode**, select **Active Active**.
5. Set the **Device ID** to **1**.
6. Select **Enable Config Sync**. This setting is required to synchronize the two firewall configurations (enabled by default).
7. Enter the **Peer HA1 IP Address**, which is the IP address of the HA1 control link on the peer firewall.
8. (Optional) Enter a **Backup Peer HA1 IP Address**, which is the IP address of the backup control link on the peer firewall.
9. Click **OK**.

STEP 3 | Configure Active/Active HA.

Complete Step 6 through Step 14.

STEP 4 | Configure Session Owner and Session Setup.

1. In **Device** > **High Availability** > **Active/Active Config**, edit Packet Forwarding.
2. For **Session Owner Selection**, select **First Packet**—The firewall that receives the first packet of a new session is the session owner.
3. For **Session Setup**, select **IP Modulo**—Distributes session setup load based on parity of the source IP address.
4. Click **OK**.

https://t.me/learningnets

STEP 5 | Configure an HA virtual address.

1. Select **Device** > **High Availability** > **Active/Active Config** > **Virtual Address** and click **Add**.
2. Select **Interface** eth1/1.
3. Select **IPv4** and **Add** an **IPv4 Address** of 10.1.1.101.
4. For **Type**, select **Floating**, which configures the virtual IP address to be a floating IP address.

STEP 6 | Configure the floating IP address.

1. Do not select **Floating IP bound to the Active-Primary device**.
2. Select **Failover address if link state is down** to cause the firewall to use the failover address when the link state on the interface is down.
3. Click **OK**.

STEP 7 | Enable jumbo frames on firewalls other than the PA-7000 Series.

STEP 8 | Define HA Failover Conditions.

STEP 9 | **Commit** the configuration.

STEP 10 | Configure the peer firewall, PA-3050-1 with the same settings, except for the following changes:

- Select **Device ID 0**.
- Configure an HA virtual address of 10.1.1.100.
- For **Device 1 Priority**, enter 255. For **Device 0 Priority**, enter 0.

In this example, Device ID 0 has a lower priority value so a higher priority; therefore, the firewall with Device ID 0 (PA-3050-1) owns the floating IP address 10.1.1.100.

STEP 11 | Still on PA-3050-1, create the source NAT rule for Device ID 0.

1. Select **Policies** > **NAT** and click **Add**.
2. Enter a **Name** for the rule that in this example identifies it as a source NAT rule for Device ID 0.
3. For **NAT Type**, select **ipv4** (default).
4. On the **Original Packet**, for **Source Zone**, select **Any**.
5. For **Destination Zone**, select the zone you created for the external network.
6. Allow **Destination Interface**, **Service**, **Source Address**, and **Destination Address** to remain set to **Any**.
7. For the **Translated Packet**, select **Dynamic IP And Port** for **Translation Type**.
8. For **Address Type**, select **Interface Address**, in which case the translated address will be the IP address of the interface. Select an **Interface** (eth1/1 in this example) and an **IP Address** of the floating IP address 10.1.1.100.
9. On the **Active/Active HA Binding** tab, for **Active/Active HA Binding**, select **0** to bind the NAT rule to Device ID 0.
10. Click **OK**.

https://t.me/learningnets

STEP 12 | Create the source NAT rule for Device ID 1.

1. Select **Policies** > **NAT** and click **Add**.
2. Enter a **Name** for the policy rule that in this example helps identify it as a source NAT rule for Device ID 1.
3. For **NAT Type**, select **ipv4** (default).
4. On the **Original Packet**, for **Source Zone**, select **Any**. For **Destination Zone**, select the zone you created for the external network.
5. Allow **Destination Interface**, **Service**, **Source Address**, and **Destination Address** to remain set to **Any**.
6. For the **Translated Packet**, select **Dynamic IP And Port** for **Translation Type**.
7. For **Address Type**, select **Interface Address**, in which case the translated address will be the IP address of the interface. Select an **Interface** (eth1/1 in this example) and an **IP Address** of the floating IP address 10.1.1.101.
8. On the **Active/Active HA Binding** tab, for the **Active/Active HA Binding**, select **1** to bind the NAT rule to Device ID 1.
9. Click **OK**.

STEP 13 | **Commit** the configuration.

Use Case: Configure Separate Source NAT IP Address Pools for Active/Active HA Firewalls

If you want to use IP address pools for source NAT in Active/Active HA Mode, each firewall must have its own pool, which you then bind to a Device ID in a NAT rule.

Address objects and NAT rules are synchronized (in both active/passive and active/active mode), so they need to be configured on only one of the firewalls in the HA pair.

This example configures an address object named Dyn-IP-Pool-dev0 containing the IP address pool 10.1.1.140-10.1.1.150. It also configures an address object named Dyn-IP-Pool-dev1 containing the IP address pool 10.1.1.160-10.1.1.170. The first address object is bound to Device ID 0; the second address object is bound to Device ID 1.

STEP 1 | On one HA firewall, create address objects.

1. Select **Objects** > **Addresses** and **Add** an address object **Name**, in this example, Dyn-IP-Pool-dev0.
2. For **Type**, select **IP Range** and enter the range 10.1.1.140-10.1.1.150.
3. Click **OK**.
4. Repeat this step to configure another address object named Dyn-IP-Pool-dev1 with the **IP Range** of 10.1.1.160-10.1.1.170.

https://t.me/learningnets

STEP 2 | Create the source NAT rule for Device ID 0.

1. Select **Policies** > **NAT** and **Add** a NAT policy rule with a **Name**, for example, Src-NAT-dev0.
2. For **Original Packet**, for **Source Zone**, select **Any**.
3. For **Destination Zone**, select the destination zone for which you want to translate the source address, such as Untrust.
4. For **Translated Packet**, for **Translation Type**, select **Dynamic IP and Port**.
5. For **Translated Address**, **Add** the address object you created for the pool of addresses belonging to Device ID 0: Dyn-IP-Pool-dev0.
6. For **Active/Active HA Binding**, select **0** to bind the NAT rule to Device ID 0.
7. Click **OK**.

STEP 3 | Create the source NAT rule for Device ID 1.

1. Select **Policies** > **NAT** and **Add** a NAT policy rule with a **Name**, for example, Src-NAT-dev1.
2. For **Original Packet**, for **Source Zone**, select **Any**.
3. For **Destination Zone**, select the destination zone for which you want to translate the source address, such as Untrust.
4. For **Translated Packet**, for **Translation Type**, select **Dynamic IP and Port**.
5. For **Translated Address**, **Add** the address object you created for the pool of addresses belonging to Device ID 1: Dyn-IP-Pool-dev1.
6. For **Active/Active HA Binding**, select **1** to bind the NAT rule to Device ID 1.
7. Click **OK**.

STEP 4 | **Commit** the configuration.

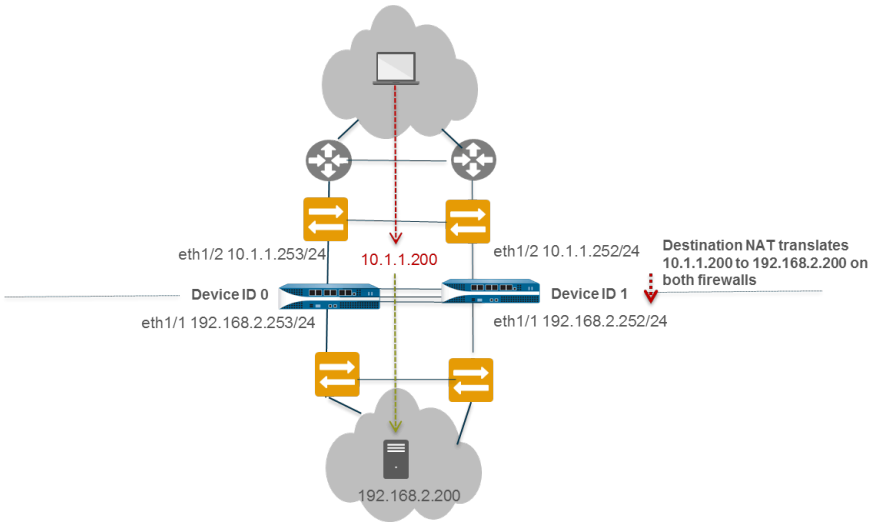
Use Case: Configure Active/Active HA for ARP Load-Sharing with Destination NAT

This Layer 3 interface example uses NAT in Active/Active HA Mode and ARP Load-Sharing with destination NAT. Both HA firewalls respond to an ARP request for the destination NAT address with the ingress interface MAC address. Destination NAT translates the public, shared IP address (in this example, 10.1.1.200) to the private IP address of the server (in this example, 192.168.2.200).

When the HA firewalls receive traffic for the destination 10.1.1.200, both firewalls could possibly respond to the ARP request, which could cause network instability. To avoid the potential issue, configure the firewall that is in active-primary state to respond to the ARP request by binding the destination NAT rule to the active-primary firewall.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets



STEP 1 | On PA-3050-2 (Device ID 1), perform Step 1 through Step 3 of Configure Active/Active HA.

STEP 2 | Enable active/active HA.

1. In **Device** > **High Availability** > **General**, edit Setup.
2. Select **Enable HA**.
3. Enter a **Group ID**, which must be the same for both firewalls. The firewall uses the Group ID to calculate the virtual MAC address (range is 1 to 63).
4. (Optional) Enter a **Description**.
5. For **Mode**, select **Active Active**.
6. Select **Device ID** to be **1**.
7. Select **Enable Config Sync**. This setting is required to synchronize the two firewall configurations (enabled by default).
8. Enter the **Peer HA1 IP Address**, which is the IP address of the HA1 control link on the peer firewall.
9. (Optional) Enter a **Backup Peer HA1 IP Address**, which is the IP address of the backup control link on the peer firewall.
10. Click **OK**.

STEP 3 | Perform Step 6 through Step 15 in Configure Active/Active HA.

STEP 4 | Configure an HA virtual address.

1. Select **Device** > **High Availability** > **Active/Active Config** > **Virtual Address** and click **Add**.
2. Select **Interface** eth1/1.
3. Select **IPv4** and **Add** an **IPv4 Address** of 10.1.1.200.
4. For **Type**, select **ARP Load Sharing**, which configures the virtual IP address to be for both peers to use for ARP Load-Sharing.

https://t.me/learningnets

STEP 5 | Configure ARP Load-Sharing.

The device selection algorithm determines which HA firewall responds to the ARP requests to provide load sharing.

1. For **Device Selection Algorithm**, select **IP Modulo**. The firewall that will respond to ARP requests is based on the parity of the ARP requester's IP address.
2. Click **OK**.

STEP 6 | Enable jumbo frames on firewalls other than the PA-7000 Series.

STEP 7 | Define HA Failover Conditions.

STEP 8 | **Commit** the configuration.

STEP 9 | Configure the peer firewall, PA-3050-1 (Device ID 0), with the same settings, except in Step 2 select **Device ID 0**.

STEP 10 | Still on PA-3050-1 (Device ID 0), create the destination NAT rule so that the active-primary firewall responds to ARP requests.

1. Select **Policies** > **NAT** and click **Add**.
2. Enter a **Name** for the rule that, in this example, identifies it as a destination NAT rule for Layer 2 ARP.
3. For **NAT Type**, select **ipv4** (default).
4. On the **Original Packet**, for **Source Zone**, select **Any**.
5. For **Destination Zone**, select the Untrust zone you created for the external network.
6. Allow **Destination Interface**, **Service**, and **Source Address** to remain set to **Any**.
7. For **Destination Address**, specify 10.1.1.200.
8. For the **Translated Packet**, Source Address Translation remains **None**.
9. For **Destination Address Translation**, enter the private IP address of the destination server, in this example, 192.168.1.200.
10. On the **Active/Active HA Binding** tab, for **Active/Active HA Binding**, select **primary** to bind the NAT rule to the firewall in active-primary state.
11. Click **OK**.

STEP 11 | **Commit** the configuration.

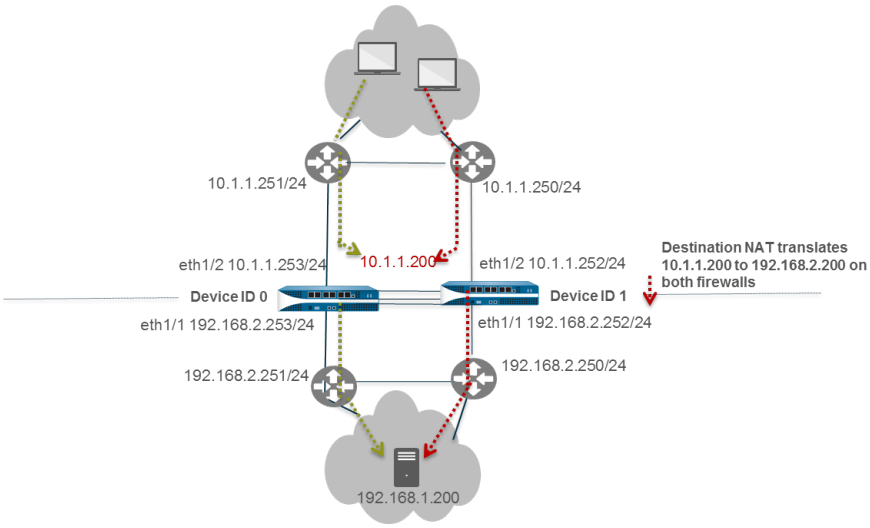
Use Case: Configure Active/Active HA for ARP Load-Sharing with Destination NAT in Layer 3

This Layer 3 interface example uses NAT in Active/Active HA Mode and ARP Load-Sharing. PA-3050-1 has Device ID 0 and its HA peer, PA-3050-2, has Device ID 1.

In this use case, both of the HA firewalls must respond to an ARP request for the destination NAT address. Traffic can arrive at either firewall from either WAN router in the untrust zone. Destination NAT translates the public-facing, shared IP address to the private IP address of the server. The configuration requires one destination NAT rule bound to both Device IDs so that both firewalls can respond to ARP requests.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets



STEP 1 | On PA-3050-2 (Device ID 1), perform Step 1 through Step 3 of Configure Active/Active HA.

STEP 2 | Enable active/active HA.

1. Select **Device** > **High Availability** > **General** > **Setup** and edit.
2. Select **Enable HA**.
3. Enter a **Group ID**, which must be the same for both firewalls. The firewall uses the Group ID to calculate the virtual MAC address (range is 1-63).
4. (Optional) Enter a **Description**.
5. For **Mode**, select **Active Active**.
6. Select **Device ID** to be **1**.
7. Select **Enable Config Sync**. This setting is required to synchronize the two firewall configurations (enabled by default).
8. Enter the **Peer HA1 IP Address**, which is the IP address of the HA1 control link on the peer firewall.
9. (Optional) Enter a **Backup Peer HA1 IP Address**, which is the IP address of the backup control link on the peer firewall.
10. Click **OK**.

STEP 3 | Configure Active/Active HA.

Perform Step 6 through Step 15.

STEP 4 | Configure an HA virtual address.

1. Select **Device** > **High Availability** > **Active/Active Config** > **Virtual Address** and click **Add**.
2. Select **Interface** eth1/2.
3. Select **IPv4** and **Add** an **IPv4 Address** of 10.1.1.200.
4. For **Type**, select **ARP Load Sharing**, which configures the virtual IP address to be for both peers to use for ARP Load-Sharing.

https://t.me/learningnets

STEP 5 | Configure ARP Load-Sharing.

The device selection algorithm determines which HA firewall responds to the ARP requests to provide load sharing.

1. For **Device Selection Algorithm**, select one of the following
   - **IP Modulo**—The firewall that will respond to ARP requests is based on the parity of the ARP requester's IP address.
   - **IP Hash**—The firewall that will respond to ARP requests is based on a hash of the ARP requester's source IP address and destination IP address.
2. Click **OK**.

STEP 6 | Enable jumbo frames on firewalls other than PA-7000 Series firewalls.

STEP 7 | Define HA Failover Conditions.

STEP 8 | **Commit** the configuration.

STEP 9 | Configure the peer firewall, PA-3050-1 (Device ID 0), with the same settings, except set the **Device ID** to **0** instead of **1**.

STEP 10 | Still on PA-3050-1 (Device ID 0), create the destination NAT rule for both Device ID 0 and Device ID 1.

1. Select **Policies** > **NAT** and click **Add**.
2. Enter a **Name** for the rule that in this example identifies it as a destination NAT rule for Layer 3 ARP.
3. For **NAT Type**, select **ipv4** (default).
4. On the **Original Packet**, for **Source Zone**, select **Any**.
5. For **Destination Zone**, select the Untrust zone you created for the external network.
6. Allow **Destination Interface**, **Service**, and **Source Address** to remain set to **Any**.
7. For **Destination Address**, specify 10.1.1.200.
8. For the **Translated Packet**, Source Address Translation remains None.
9. For **Destination Address Translation**, enter the private IP address of the destination server, in this example 192.168.1.200.
10. On the **Active/Active HA Binding** tab, for **Active/Active HA Binding**, select **both** to bind the NAT rule to both Device ID 0 and Device ID 1.
11. Click **OK**.

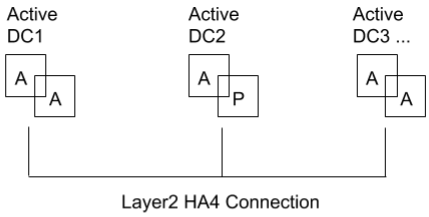
STEP 11 | **Commit** the configuration.

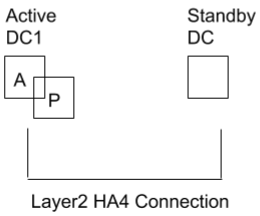
https://t.me/learningnets

HA Clustering Overview

A number of Palo Alto Networks® firewall models now support session state synchronization among firewalls in a high availability (HA) cluster of up to 16 firewalls. The HA cluster peers synchronize sessions to protect against failure of the data center or a large security inspection point with horizontally scaled firewalls. In the case of a network outage or a firewall going down, the sessions fail over to a different firewall in the cluster. Such synchronization is especially helpful in the following use cases.

One use case is when HA peers are spread across multiple data centers so that there is no single point of failure within or between data centers. A second multi-data center use case is when one data center is active and the other is standby.

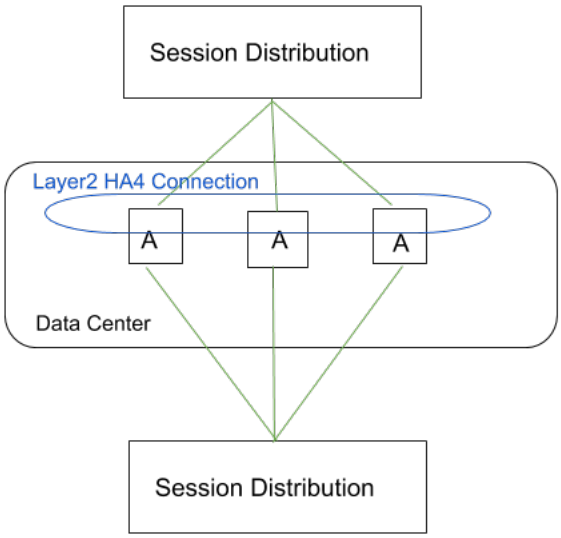




A third HA clustering use case is horizontal scaling, in which you add HA cluster members to a single data center to scale security and ensure session survivability.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets



HA clusters support a Layer 3 or virtual wire deployment. HA peers in the cluster can be a combination of HA pairs and standalone cluster members. In an HA cluster, all members are considered active; there is no concept of passive firewalls except for HA pairs, which can keep their active/passive relationship after you add them to an HA cluster.

All cluster members share session state. When a new firewall joins an HA cluster, that triggers all firewalls in the cluster to synchronize all existing sessions. HA4 and HA4 backup connections are the dedicated cluster links that synchronize session state among all cluster members having the same cluster ID. The HA4 link between cluster members detects connectivity failures between cluster members. HA1 (control link), HA2 (data link), and HA3 (packet-forwarding link) are not supported between cluster members that aren't HA pairs.

For a normal session that has not failed over, only the firewall that is the session owner creates a traffic log. For a session that failed over, the new session owner (the firewall that receives the failed over traffic) creates the traffic log.

The firewall models that support HA clustering and the maximum number of members supported per cluster are as follows:

Firewall Model	Number of Members Supported Per Cluster
PA-3200 Series	6
PA-5200 Series	16
PA-5450	8
PA-7000 Series firewalls that have at least one of the following cards: PA-7000-100G-NPC, PA-7000-20GQXM-NPC, PA-7000-20GXM-NPC	PA-7080: 4 PA-7050: 6

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Firewall Model	Number of Members Supported Per Cluster
VM-300	6
VM-500	6
VM-700	16

HA clustering is not supported in public cloud deployments. Consider the HA Clustering Best Practices and Provisioning before you start to Configure HA Clustering.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

HA Clustering Best Practices and Provisioning

These are the provisioning requirements and best practices for HA clustering.

- Provisioning Requirements and Best Practices
  - HA cluster members must be the same firewall model and run the same PAN-OS® version.

    When upgrading, firewall members will continue to synchronize sessions with one member at a different version.
  - It is highly recommended and a best practice to use Panorama to provision HA cluster members to keep all configuration and policies synchronized among all cluster members.
  - HA cluster members must be licensed for the same components to ensure consistent policy enforcement and content inspection capabilities.
  - The licenses must expire at the same time to prevent mismatched licenses and loss of functionality.
  - All cluster members should be running with the same version of dynamic Content Updates for consistent security enforcement.
  - HA cluster members must share the same zone names in order for sessions to successfully fail over to another cluster member. For example, suppose sessions going to an ingress zone named **`internal`** are dropped because the link is down. For those sessions to fail over to an HA firewall peer in the cluster, that peer must also have a zone named **`internal`**.
  - Client-to-server and server-to-client flows must go back to the same firewall under normal (non-failure) conditions in order for security content scanning to occur. Asymmetric traffic won't be dropped, but it cannot be scanned for security purposes.
- Session Synchronization Best Practices
  - Dedicated HA communication interfaces should be used over dataplane interfaces. HSCI interfaces aren't used for HA4. This allows separation of HA pair and cluster session synchronization to ensure maximum bandwidth and reliability for session syncing.
  - HA4 should be adequately sized if you use dataplane interfaces. This ensures best effort session state synchronizing between cluster members.
  - Best practice is to have a dedicated cluster network for the HA4 communications link to ensure adequate bandwidth and non-congested, low-latency connections between cluster members.
  - Architect your networks and perform traffic engineering to avoid possible race conditions, in which a network steers traffic from the session owner to a cluster member before the session is successfully synced between the firewalls. Layer2 HA4 connections must have sufficient bandwidth and low latency to allow timely synchronization between HA members. The HA4 latency must be lower than the latency incurred when the peering devices switch traffic between cluster members.
  - Architect your networks to minimize asymmetric flows. Session setup requires one cluster member to see the complete TCP three-way handshake.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

- Health Check Best Practices
  - On HA pairs in a cluster, configure an Active/Passive pair with HA backup communication links for HA1, HA2, and HA4. Configure an Active/Active pair with HA backup communications links for HA1, HA2, HA3, and HA4.
  - Configure HA4 backup links on all cluster members.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Configure HA Clustering

Learn about HA clustering and follow the HA Clustering Best Practices and Provisioning before you configure HA firewalls as members of a cluster.

STEP 1 | Establish an interface as an HA interface (to later assign as the HA4 link).

1. Select **Network** > **Interfaces** > **Ethernet** and select an interface; for example, ethernet1/1.
2. Select the **Interface Type** to be **HA**.
3. Click **OK**.
4. Repeat this step to configure another interface to use as the HA4 backup link.

STEP 2 | Enable HA clustering.

1. Select **Device** > **High Availability** > **General** and edit the Clustering Settings.
2. **Enable Cluster Participation**.
3. Enter the **Cluster ID**, a unique numeric ID for an HA cluster in which all members can share session state; range is 1 to 99.
4. Enter a short, helpful **Cluster Description**.
5. (Optional) Change **Cluster Synchronization Timeout (min)**, which is the maximum number of minutes that the local firewall waits before going to Active state when another cluster member (for example, in unknown state) is preventing the cluster from fully synchronizing; range is 0 to 30; default is 0.
6. (Optional) Change **Monitor Fail Hold Down Time (min)**, which is the number of minutes after which a down link is retested to see if it is back up; range is 1 to 60; default is 1.
7. Click **OK**.

STEP 3 | Configure the HA4 link.

1. Select **HA Communications** and in the Clustering Links section, edit the HA4 section.
2. Select the interface you configured in the first step as an **HA** interface to be the **Port** for the HA4 link; for example, ethernet1/1.
3. Enter the **IPv4/IPv6 Address** of the local HA4 interface.
4. Enter the **Netmask**.
5. (Optional) Change the **HA4 Keep-alive Threshold (ms)** to specify the timeframe within which the firewall must receive keepalives from a cluster member to know that the cluster member is functional; range is 5,000 to 60,000; default is 10,000.
6. Click **OK**.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 4 | Configure the HA4 Backup link.

1. Edit the HA4 Backup section.
2. Select the other interface you configured in the first step as an **HA** interface to be the **Port** for the HA4 backup link.
3. Enter the **IPv4/IPv6 Address** of the local HA4 backup interface.
4. Enter the **Netmask**.
5. Click **OK**.

STEP 5 | Specify all members of the HA cluster, including the local member and both HA peers in any HA pair.

1. Select **Cluster Config**.
2. (On a supported firewall) **Add** a peer member's **Device Serial Number**.
3. (On Panorama) **Add** and select a **Device** from the dropdown and enter a **Device Name**.
4. Enter the **HA4 IP Address** of the HA peer in the cluster.
5. Enter the **HA4 Backup IP Address** of the HA peer in the cluster.
6. Enable **Session Synchronization** with the peer you identified.
7. (Optional) Enter a helpful **Description**.
8. Click **OK**.
9. Select the device and **Enable** it.

STEP 6 | Define HA failover conditions with link and path monitoring.

STEP 7 | **Commit**.

STEP 8 | (Panorama only) Refresh the list of HA firewalls in the HA cluster.

1. Under Templates, select **Device** > **High Availability** > **Cluster Config**.
2. Click **Refresh** at the bottom of the screen.

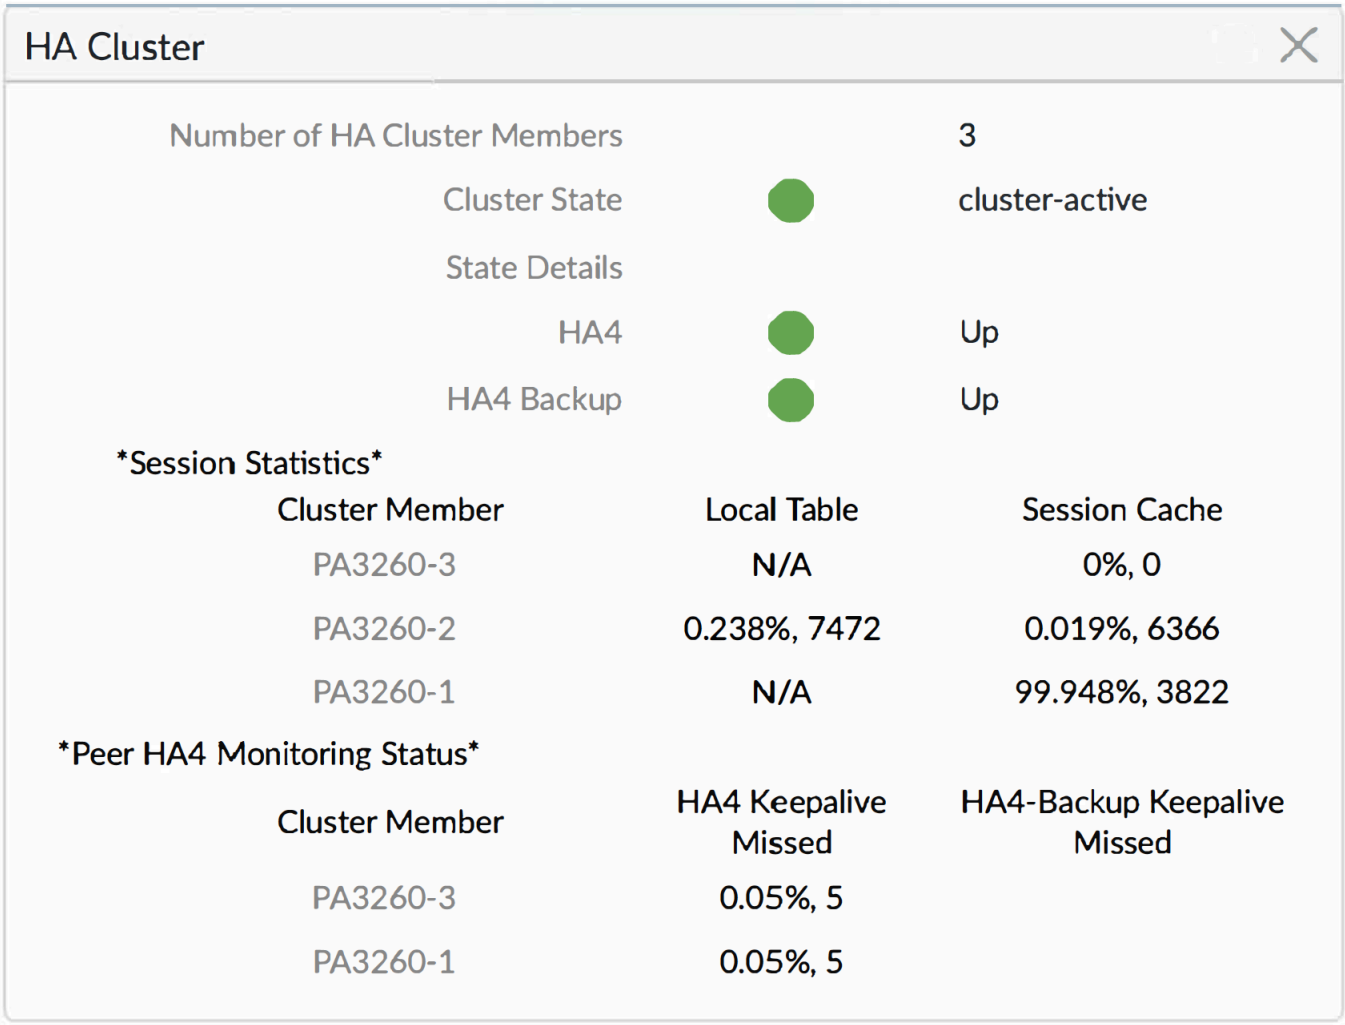
STEP 9 | View HA cluster information in the UI.

1. Select **Dashboard**.
2. View the HA cluster fields. The top section displays cluster state and HA4 connections to provide cluster health at a glance. The HA4 and HA4 Backup indicators will be one of the following: Green indicates the link status of the cluster members is Up. Red indicates the link status of all the cluster members is Down. Yellow indicates the link status of some cluster members is Up while the status of other cluster members is Down. Grey indicates not configured. The center section displays the capacity of the local session table and session cache table so you can monitor how full the tables are and plan for firewall upgrades. The lower section displays communication errors on the HA4 and

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

HA4 backup links, signifying possible problems with synchronizing information between members.



STEP 10 | Access the CLI to view HA cluster and HA4 link information and perform other HA clustering tasks.

You can view HA cluster flap statistics. The cluster flap count is reset when the HA device moves from suspended to functional and vice versa. The cluster flap count also resets when the non-functional hold time expires.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Refresh HA1 SSH Keys and Configure Key Options

All Palo Alto Networks firewalls come with Secure Shell (SSH) pre-configured, and the high availability (HA) firewalls can act as SSH server and SSH client simultaneously. When you configure active/passive or active/active HA, you can enable encryption for the HA1 (control link) connection between the HA firewalls. We recommend you secure the HA1 traffic between the HA peers with encryption, particularly if the firewalls aren't located in the same site. After you enable encryption on the HA1 control link, you can use the CLI to create an SSH service profile and secure the connection between the HA firewalls.

SSH service profiles enable you to change the default host key type, generate a new pair of public and private SSH host keys for the HA1 control link, and configure other SSH HA1 settings. You can apply the new host keys and configured settings to the firewalls without restarting the HA peers. The firewall will reestablish HA1 sessions with its peer to synchronize the configuration changes. It also generates system logs (subtype is ha) for reestablishing HA1 and HA1-backup sessions.

The following examples show how to configure various SSH settings for your HA1 after you enable encryption and access the CLI. (See Refresh SSH Keys and Configure Key Options for Management Interface Connection for SSH management server profile examples.)

You must enable encryption and it must be functioning properly on an HA pair before you can perform the following tasks.

If you are configuring the HA1 control link in FIPS-CC mode, you must set automatic rekeying parameters for session keys.

*To use the same SSH connection settings for each Dedicated Log Collector (M-series or Panorama virtual appliance in Log Collector mode) in a Collector Group, configure an SSH service profile from the Panorama management server, **Commit** the changes to Panorama, and then **Push** the configuration to the Log Collectors. You can use the **`set log-collector-group <name> general-setting management ssh`** commands.*

- Create an SSH service profile to exercise greater control over SSH connections between your HA firewalls.

  This example creates an HA profile without configuring any settings.

  1. admin@PA-3250> **`configure`**
  2. admin@PA-3250# **`set deviceconfig system ssh profiles ha-profiles <name>`**
  3. admin@PA-3250# **`commit`**
  4. admin@PA-3250# **`exit`**
  5. To verify that the new profile has been created and view the settings for any existing profiles:

     admin@PA-3250> **`configure`**

     admin@PA-3250# **`show deviceconfig system ssh profiles`**

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

- (Optional) Set the SSH server to use only the specified encryption ciphers for the HA1 sessions.

  By default, HA1 SSH allows all supported ciphers for encryption of CLI HA sessions. When you set one or more ciphers, the SSH server advertises only those ciphers while connecting, and if the SSH client (HA peer) tries to connect using a different cipher, the server terminates the connection.

  1. admin@PA-3250> **configure**
  2. admin@PA-3250# **set deviceconfig system ssh profiles ciphers ha-profiles <*name*> ciphers <*cipher*>**

     aes128-cbc—AES 128-bit cipher with Cipher Block Chaining

     aes128-ctr—AES 128-bit cipher with Counter Mode

     aes128-gcm—AES 128-bit cipher with GCM (Galois/Counter Mode)

     aes192-cbc—AES 192-bit cipher with Cipher Block Chaining

     aes192-ctr—AES 192-bit cipher with Counter Mode

     aes256-cbc—AES 256-bit cipher with Cipher Block Chaining

     aes256-ctr—AES 256-bit cipher with Counter Mode

     aes256-gcm—AES 256-bit cipher with GCM
  3. admin@PA-3250# **commit**
  4. admin@PA-3250# **exit**
  5. (HA1 Backup is configured) admin@PA-3250> **request high-availability session-reestablish**
  6. (No HA1 Backup is configured or HA1 Backup link is down) admin@PA-3250> **request high-availability session-reestablish force**

     *You can force the firewall to reestablish HA1 sessions if there is no HA1 backup, which causes a brief split-brain condition between the HA peers. (Using the **force** option when an HA1 backup is configured has no effect.)*
  7. To verify the ciphers have been updated:

     admin@PA-3250> **configure**

     admin@PA-3250# **show deviceconfig system ssh profiles ha-profiles ciphers**

- (Optional) Set the default host key type.

  If you enable encryption on the HA1 control link, the firewall uses a default host key type of RSA 2048 unless you change it. The HA1 SSH connection uses only the `default host key type` to authenticate the HA peers (before an encrypted session is established between them). You can change the default host key type; the choices are ECDSA 256, 384, or 521, or RSA 2048, 3072, or 4096. Change the default host key type if you prefer a longer RSA key length or if you prefer ECDSA rather than RSA. This example sets the default host key type to an ECDSA

https://t.me/learningnets

key of 256 bits. It also re-establishes the HA1 connection using the new host key without restarting the HA peers.

1. admin@PA-3250> **configure**
2. admin@PA-3250# **set deviceconfig system ssh profiles ha-profiles <*name*> default-hostkey key-type ECDSA key-length 256**
3. admin@PA-3250# **commit**
4. admin@PA-3250# **exit**
5. admin@PA-3250> **request high-availability sync-to-remote ssh-key**

   An HA connection must already be established between the HA firewalls. If the firewalls have not yet established an HA connection, you must enable encryption on the control link connection, export the HA key to a network location and import the HA key on the peer. See Configure Active/Passive HA *or* Configure Active/Active HA.
6. (HA1 Backup is configured) admin@PA-3250> **request high-availability session-reestablish**
7. (No HA1 Backup is configured or HA1 Backup link is down) admin@PA-3250> **request high-availability session-reestablish force**

   You can force the firewall to reestablish HA1 sessions if there is no HA1 backup, which causes a brief split-brain condition between the two HA peers. (Using the ***force*** *option when an HA1 backup is configured has no effect.)*
8. To verify the host key has been updated:

   admin@PA-3250> **configure**

   admin@PA-3250# **show deviceconfig system ssh profiles ha-profiles <*name*> default-hostkey**

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

- (Optional) Delete a cipher from the set of ciphers you selected for SSH over the HA1 control link.

  This example deletes the AES CBC cipher with 128-bit key.

  1. admin@PA-3250> **configure**
  2. admin@PA-3250# **delete deviceconfig system ssh profiles ha-profiles <*name*> ciphers aes128-cbc**
  3. admin@PA-3250# **commit**
  4. admin@PA-3250# **exit**
  5. (HA1 Backup is configured) admin@PA-3250> **request high-availability session-reestablish**
  6. (No HA1 Backup is configured or HA1 Backup link is down) admin@PA-3250> **request high-availability session-reestablish force**

     You can force the firewall to reestablish HA1 sessions if there is no HA1 backup, which causes a brief split-brain condition between the two HA peers. (Using the ***force*** *option when an HA1 backup is configured has no effect.*
  7. To verify the cipher has been deleted:

     admin@PA-3250> **configure**

     admin@PA-3250# **show deviceconfig system ssh profiles ha-profiles <*name*> ciphers**

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

- (Optional) Set the session key exchange algorithms the HA1 SSH server will support.

  By default, the SSH server (HA firewall) advertises all the key exchange algorithms to the SSH client (HA peer firewall).

  If you are using an ECDSA default key type, the best practice is to use an ECDH key algorithm.

  1. admin@PA-3250> **configure**
  2. admin@PA-3250# **set deviceconfig system ssh profiles ha-profiles <*name*> kex <*value*>**

     diffie-hellman-group14-sha1—Diffie-Hellman group 14 with SHA1 hash

     ecdh-sha2-nistp256—Elliptic-Curve Diffie-Hellman over National Institute of Standards and Technology (NIST) P-256 with SHA2-256 hash

     ecdh-sha2-nistp384—Elliptic-Curve Diffie-Hellman over NIST P-384 with SHA2-384 hash

     ecdh-sha2-nistp521—Elliptic-Curve Diffie-Hellman over NIST P-521 with SHA2-521 hash
  3. admin@PA-3250# **commit**
  4. admin@PA-3250# **exit**
  5. (HA1 Backup is configured) admin@PA-3250> **request high-availability session-reestablish**
  6. (No HA1 Backup is configured or HA1 Backup link is down) admin@PA-3250> **request high-availability session-reestablish force**

     *You can force the firewall to reestablish HA1 sessions if there is no HA1 backup, which causes a brief split-brain condition between the two HA peers. (Using the **force** option when an HA1 backup is configured has no effect.*
  7. To verify the key exchange algorithms have been updated:

     admin@PA-3250> **configure**

     admin@PA-3250# **show deviceconfig system ssh profiles ha-profiles**

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

- (Optional) Set the message authentication codes (MAC) the HA1 SSH server will support.

  By default, the server advertises all of the MAC algorithms to the client.

  1. admin@PA-3250> **configure**
  2. admin@PA-3250# **set deviceconfig system ssh profiles ha-profiles <*name*> mac <*value*>**

     hmac-sha1—MAC with SHA1 cryptographic hash

     hmac-sha2-256—MAC with SHA2-256 cryptographic hash

     hmac-sha2-512—MAC with SHA2-512 cryptographic hash
  3. admin@PA-3250# **commit**
  4. admin@PA-3250# **exit**
  5. (HA1 Backup is configured) admin@PA-3250> **request high-availability session-reestablish**
  6. (No HA1 Backup is configured or HA1 Backup link is down) admin@PA-3250> **request high-availability session-reestablish force**

     You can force the firewall to reestablish HA1 sessions if there is no HA1 backup, which causes a brief split-brain condition between the two HA peers. (Using the ***force*** *option has no effect when an HA1 backup is configured.*
  7. To verify the MAC algorithms have been updated:

     admin@PA-3250> **configure**

     admin@PA-3250# **show deviceconfig system ssh profiles ha-profiles**

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

- (Optional) Regenerate ECDSA or RSA host keys for HA1 SSH to replace the existing keys, and re-establish HA1 sessions between HA peers using the new keys without restarting the HA peers.

  The HA peers use the host keys to authenticate each other. This example regenerates the ECDSA 256 default host key.

  Regenerating a host key does not change your default host key type. To regenerate the default host key you are using, you must specify your default host key type and length when you regenerate. Regenerating a host key that isn't your default host key type simply regenerates a key that you aren't using and therefore has no effect.

  1. admin@PA-3250> **configure**
  2. admin@PA-3250# **set deviceconfig system ssh regenerate-hostkeys ha key-type ECDSA key-length 256**
  3. admin@PA-3250# **commit**
  4. admin@PA-3250# **exit**
  5. admin@PA-3250> **request high-availability sync-to-remote ssh-key**

     An HA connection must already be established between the HA firewalls. If the firewalls have not yet established an HA connection, you must enable encryption on the control link connection, export the HA key to a network location, and import the HA key on the peer. See Configure Active/Passive HA *or* Configure Active/Active HA.
  6. (HA1 Backup is configured) admin@PA-3250> **request high-availability session-reestablish**
  7. (No HA1 Backup is configured or HA1 Backup link is down) admin@PA-3250> **request high-availability session-reestablish force**

     You can force the firewall to reestablish HA1 sessions if there is no HA1 backup, which causes a brief split-brain condition between the two HA peers. (Using the ***force*** *option when an HA1 backup is configured has no effect.)*

- (Optional) Set rekey parameters to establish when automatic rekeying of the session keys occurs for SSH over the HA1 control link.

  The session keys are used to encrypt the traffic between the HA peers. The parameters you can set are data volume (in megabytes), time interval (seconds), and packet count. After any one rekey parameter reaches its configured value, SSH initiates a key exchange.

  You can set a second or third parameter if you aren't sure the parameter you configured will reach its value as soon as you want rekeying to occur. The first parameter to reach its configured value will prompt a rekey, then the firewall will reset all rekey parameters.

  1. admin@PA-3250> **configure**
  2. admin@PA-3250# **set deviceconfig system ssh profiles ha-profiles <*name*> session-rekey data 32**

     Rekeying occurs after the volume of data (in megabytes) is transmitted following the previous rekey. The default is based on the cipher you use and ranges from 1GB to 4GB;

https://t.me/learningnets

the range is 10MB to 4,000MB. Alternatively, you can enter **set deviceconfig system ssh profiles ha-profiles <*name*> session-rekey data default** command, which sets the data parameter to the default value of the individual cipher you are using.

3. admin@PA-3250# **set deviceconfig system ssh profiles ha-profiles <*name*> session-rekey interval 3600**

   Rekeying occurs after the specified time interval (in seconds) passes following the previous rekeying. By default, time-based rekeying is disabled (set to none). The range is 10 to 3,600.

4. admin@PA-3250# **set deviceconfig system ssh profiles ha-profiles <*name*> session-rekey packets 27**

   Rekeying occurs after the defined number of packets (2^n) are transmitted following the previous rekey. For example, 14 configures that a maximum of 2^{14} packets are transmitted before a rekey occurs. The default is 2^{28}. The range is 12 to 27 (2^{12} to 2^{27}). Alternatively, you can enter **set deviceconfig system ssh profiles ha-profiles <*name*> session-rekey packets default**, which sets the packets parameter to 2^{28}.

   Choose rekeying parameters based on your type of traffic and network speeds (in addition to FIPS-CC requirements if they apply to you). Don't set the parameters so low that they affect SSH performance.

5. admin@PA-3250# **commit**
6. admin@PA-3250# **exit**
7. (HA1 Backup is configured) admin@PA-3250> **request high-availability session-reestablish**
8. (No HA1 Backup is configured or HA1 Backup link is down) admin@PA-3250> **request high-availability session-reestablish force**

   *You can force the firewall to reestablish HA1 sessions if there is no HA1 backup, which causes a brief split-brain condition between the two HA peers. (Using the **force** option when an HA1 backup is configured has no effect.)*

9. To verify the changes:

   admin@PA-3250> **configure**

   admin@PA-3250# **show deviceconfig system ssh profiles ha-profiles <*name*> session-rekey**

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

- Activate the profile by selecting the profile and restarting HA1 SSH service.
  1. admin@PA-3250> **configure**
  2. admin@PA-3250# **set deviceconfig system ssh ha ha-profile *<name>***
  3. admin@PA-3250# **commit**
  4. admin@PA-3250# **exit**
  5. admin@PA-3250> **set ssh service-restart ha**
  6. To verify the correct profile is in use:

     admin@PA-3250> **configure**

     admin@PA-3250# **show deviceconfig system ssh ha**

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

HA Firewall States

An HA firewall can be in one of the following states:

HA Firewall State	Occurs In	Description
Initial	A/P or A/A	Transient state of a firewall when it joins the HA pair. The firewall remains in this state after boot-up until it discovers a peer and negotiations begins. After a timeout, the firewall becomes active if HA negotiation has not started.
Active	A/P	State of the active firewall in an active/passive configuration.
Passive	A/P	State of the passive firewall in an active/passive configuration. The passive firewall is ready to become the active firewall with no disruption to the network. Although the passive firewall is not processing other traffic: • If passive link state auto is configured, the passive firewall is running routing protocols, monitoring link and path state, and the passive firewall will pre-negotiate LACP and LLDP if LACP and LLDP pre-negotiation are configured, respectively. • The passive firewall is synchronizing flow state, runtime objects, and configuration. • The passive firewall is monitoring the status of the active firewall using the hello protocol.
Active-Primary	A/A	In an active/active configuration, state of the firewall that connects to User-ID agents, runs DHCP server and DHCP relay, and matches NAT and PBF rules with the Device ID of the active-primary firewall. A firewall in this state can own sessions and set up sessions.
Active-Secondary	A/A	In an active/active configuration, state of the firewall that connects to User-ID agents, runs DHCP server, and matches NAT and PBF rules with the Device ID of the active-secondary firewall. A firewall in active-secondary state does not support DHCP relay. A firewall in this state can own sessions and set up sessions.
Tentative	A/A	State of a firewall (in an active/active configuration) caused by one of the following: • Failure of a firewall. • Failure of a monitored object (a link or path). • The firewall leaves suspended or non-functional state.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

HA Firewall State	Occurs In	Description
		A firewall in tentative state synchronizes sessions and configurations from the peer. • In a virtual wire deployment, when a firewall enters tentative state due to a path failure and receives a packet to forward, it sends the packet to the peer firewall over the HA3 link for processing. The peer firewall processes the packet and sends it back over the HA3 link to the firewall to be sent out the egress interface. This behavior preserves the forwarding path in a virtual wire deployment. • In a Layer 3 deployment, when a firewall in tentative state receives a packet, it sends that packet over the HA3 link for the peer firewall to own or set up the session. Depending on the network topology, this firewall either sends the packet out to the destination or sends it back to the peer in tentative state for forwarding. After the failed path or link clears or as a failed firewall transitions from tentative state to active-secondary state, the **Tentative Hold Time** is triggered and routing convergence occurs. The firewall attempts to build routing adjacencies and populate its route table before processing any packets. Without this timer, the recovering firewall would enter active-secondary state immediately and would silently discard packets because it would not have the necessary routes. When a firewall leaves suspended state, it goes into tentative state for the **Tentative Hold Time** after links are up and able to process incoming packets. **Tentative Hold Time range (sec)** can be disabled (which is 0 seconds) or in the range 10-600; default is 60.
Non-functional	A/P or A/A	Error state due to a dataplane failure or a configuration mismatch, such as only one firewall configured for packet forwarding, VR sync or QoS sync. In active/passive mode, all of the causes listed for Tentative state cause non-functional state.
Suspended	A/P or A/A	The device is disabled so won't pass data traffic and although HA communications still occur, the device doesn't participate in the HA election process. It can't move to an HA functional state without user intervention.

https://t.me/learningnets

Reference: HA Synchronization

If you have enabled configuration synchronization on both peers in an HA pair, most of the configuration settings you configure on one peer will automatically sync to the other peer upon commit. To avoid configuration conflicts, always make configuration changes on the active (active/passive) or active-primary (active/active) peer and wait for the changes to sync to the peer before making any additional configuration changes.

Only committed configurations synchronize between HA peers. Any configuration in the commit queue at the time of an HA sync will not be synchronized.

The following topics identify which configuration settings you must configure on each firewall independently (these settings are not synchronized from the HA peer).

- What Settings Don't Sync in Active/Passive HA?
- What Settings Don't Sync in Active/Active HA?
- Synchronization of System Runtime Information

What Settings Don't Sync in Active/Passive HA?

You must configure the following settings on each firewall in an HA pair in an active/passive deployment. These settings do not sync from one peer to another.

Configuration Item	What Doesn't Sync in Active/Passive?
Management Interface Settings	All management configuration settings must be configured individually on each firewall, including: • **Device** > **Setup** > **Management** > **General Settings**—Hostname, Domain, Login Banner, SSL/TLS Service Profile (and associated certificates), Time Zone, Locale, Date, Time, Latitude, Longitude. • **Device** > **Setup** > **Management** > **Management Interface Settings**—IP Type, IP Address, Netmask, Default Gateway, IPv6 Address/Prefix Length, Default IPv6 Gateway, Speed, MTU, and Services (HTTP, HTTP OCSP, HTTPS, Telnet, SSH, Ping, SNMP, User-ID, User-ID Syslog Listener-SSL, User-ID Syslog Listener-UDP)
Multi-vsys Capability	You must activate the Virtual Systems license on each firewall in the pair to increase the number of virtual systems beyond the base number provided by default on PA-3200 Series, PA-5200 Series, and PA-7000 Series firewalls. You must also enable **Multi Virtual System Capability** on each firewall (**Device** > **Setup** > **Management** > **General Settings**).
Panorama Settings	Set the following Panorama settings on each firewall (**Device** > **Setup** > **Management** > **Panorama Settings**). • **Panorama Servers**

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Configuration Item	What Doesn't Sync in Active/Passive?
	• **Disable Panorama Policy and Objects** and **Disable Device and Network Template**
SNMP	**Device** > **Setup** > **Operations** > **SNMP Setup**
Services	**Device** > **Setup** > **Services**
Global Service Routes	**Device** > **Setup** > **Services** > **Service Route Configuration**
Telemetry and Threat Intelligence Settings	**Device** > **Setup** > **Telemetry and Threat Intelligence**
Data Protection	**Device** > **Setup** > **Content-ID** > **Manage Data Protection**
Jumbo Frames	**Device** > **Setup** > **Session** > **Session Settings** > **Enable Jumbo Frame**
Packet Buffer Protection	**Device** > **Setup** > **Session** > **Session Settings** > **Packet Buffer Protection** **Network** > **Zones** > **Enable Packet Buffer Protection**
Forward Proxy Server Certificate Settings	**Device** > **Setup** > **Session** > **Decryption Settings** > **SSL Forward Proxy Settings**
Master Key Secured by HSM	**Device** > **Setup** > **HSM** > **Hardware Security Module Provider** > **Master Key Secured by HSM**
Log Export Settings	**Device** > **Scheduled Log Export**
Software Updates	With software updates, you can either download and install them separately on each firewall, or download them on one peer and sync the update to the other peer. You must install the update on each peer (**Device** > **Software**).
GlobalProtect Agent Package	With GlobalProtect app updates, you can either download and install them separately on each firewall, or download them to one peer and sync the update to the other peer. You must activate separately on each peer (**Device** > **GlobalProtect Client**).
Content Updates	With content updates, you can either download and install them separately on each firewall, or download them on one peer and sync the update to the other peer. You must install the update on each peer (**Device** > **Dynamic Updates**).
Licenses/ Subscriptions	**Device** > **Licenses**

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Configuration Item	What Doesn't Sync in Active/Passive?
Support Subscription	**Device** > **Support**
Master Key	The master key must be identical on each firewall in the HA pair, but you must manually enter it on each firewall (**Device** > **Master Key and Diagnostics**). Before changing the master key, you must disable config sync on both peers (**Device** > **High Availability** > **General** > **Setup** and clear the **Enable Config Sync** check box) and then re-enable it after you change the keys.
Reports, logs, and Dashboard Settings	Log data, reports, and Dashboard data and settings (column display, widgets) are not synced between peers. Report configuration settings, however, are synced.
HA settings	**Device** > **High Availability**
Rule Usage Data	Rule usage data, such as hit count, Created, and Modified Dates, are not synced between peers. You need to log in to the each firewall to view the policy rule hit count data for each firewall or use Panorama to view information on the HA firewall peers.
Certificates for Device Management and Syslog Communication over SSL only	**Device** > **Certificate Management** > **Certificates** Certificates used for device management or for syslog communication over SSL don't synchronize with an HA peer. *Although the certificates used for the management interface are not synchronized (and can be different), the name of the certificate entry should be the same for the active and passive devices.*
Certificates in a Certificate Profile	**Device** > **Certificate Management** > **Certificate Profile**
SSL/TLS Service Profile for Device Management only	**Device** > **Certificate Management** > **SSL/TLS Service Profile** SSL/TLS Service Profile for Device Management doesn't synchronize with an HA peer.
Device-ID and IoT Security	IP address-to-device mappings and policy rule recommendations don't synchronize with an HA peer.

What Settings Don't Sync in Active/Active HA?

You must configure the following settings on each firewall in an HA pair in an active/active deployment. These settings do not sync from one peer to another.

https://t.me/learningnets

Configuration Item	What Doesn't Sync in Active/Active?
Management Interface Settings	You must configure all management settings individually on each firewall, including: • **Device** > **Setup** > **Management** > **General Settings**—Hostname, Domain, Login Banner, SSL/TLS Service Profile (and associated certificates), Time Zone, Locale, Date, Time, Latitude, Longitude. • **Device** > **Setup** > **Management** > **Management Interface Settings**—IP Address, Netmask, Default Gateway, IPv6 Address/Prefix Length, Default IPv6 Gateway, Speed, MTU, and Services (HTTP, HTTP OCSP, HTTPS, Telnet, SSH, Ping, SNMP, User-ID, User-ID Syslog Listener-SSL, User-ID Syslog Listener-UDP)
Multi-vsys Capability	You must activate the Virtual Systems license on each firewall in the pair to increase the number of virtual systems beyond the base number provided by default on PA-3200 Series, PA-5200 Series, and PA-7000 Series firewalls. You must also enable **Multi Virtual System Capability** on each firewall (**Device** > **Setup** > **Management** > **General Settings**).
Panorama Settings	Set the following Panorama settings on each firewall (**Device** > **Setup** > **Management** > **Panorama Settings**). • **Panorama Servers** • **Disable Panorama Policy and Objects** and **Disable Device and Network Template**
SNMP	**Device** > **Setup** > **Operations** > **SNMP Setup**
Services	**Device** > **Setup** > **Services**
Global Service Routes	**Device** > **Setup** > **Services** > **Service Route Configuration**
Telemetry and Threat Intelligence Settings	**Device** > **Setup** > **Telemetry and Threat Intelligence**
Data Protection	**Device** > **Setup** > **Content-ID** > **Manage Data Protection**
Jumbo Frames	**Device** > **Setup** > **Session** > **Session Settings** > **Enable Jumbo Frame**
Packet Buffer Protection	**Device** > **Setup** > **Session** > **Session Settings** > **Packet Buffer Protection** **Network** > **Zones** > **Enable Packet Buffer Protection**
Forward Proxy Server Certificate Settings	**Device** > **Setup** > **Session** > **Decryption Settings** > **SSL Forward Proxy Settings**

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Configuration Item	What Doesn't Sync in Active/Active?
HSM Configuration	**Device** > **Setup** > **HSM**
Log Export Settings	**Device** > **Scheduled Log Export**
Software Updates	With software updates, you can either download and install them separately on each firewall, or download them on one peer and sync the update to the other peer. You must install the update on each peer (**Device** > **Software**).
GlobalProtect Agent Package	With GlobalProtect app updates, you can either download and install them separately on each firewall, or download them to one peer and sync the update to the other peer. You must activate separately on each peer (**Device** > **GlobalProtect Client**).
Content Updates	With content updates, you can either download and install them separately on each firewall, or download them on one peer and sync the update to the other peer. You must install the update on each peer (**Device** > **Dynamic Updates**).
Licenses/ Subscriptions	**Device** > **Licenses**
Support Subscription	**Device** > **Support**
Ethernet Interface IP Addresses	All Ethernet interface configuration settings sync except for the IP address (**Network** > **Interface** > **Ethernet**).
Loopback Interface IP Addresses	All Loopback interface configuration settings sync except for the IP address (**Network** > **Interface** > **Loopback**).
Tunnel Interface IP Addresses	All Tunnel interface configuration settings sync except for the IP address (**Network** > **Interface** > **Tunnel**).
LACP System Priority	Each peer must have a unique LACP System ID in an active/active deployment (**Network** > **Interface** > **Ethernet** > **Add Aggregate Group** > **System Priority**).
VLAN Interface IP Address	All VLAN interface configuration settings sync except for the IP address (**Network** > **Interface** > **VLAN**).
Virtual Routers	Virtual router configuration synchronizes only if you have enabled VR Sync (**Device** > **High Availability** > **Active/Active Config** > **Packet Forwarding**). Whether or not to do this depends on your network design, including whether you have asymmetric routing.
IPSec Tunnels	IPSec tunnel configuration synchronization is dependent on whether you have configured the Virtual Addresses to use Floating IP

https://t.me/learningnets

Configuration Item	What Doesn't Sync in Active/Active?
	addresses (**Device** > **High Availability** > **Active/Active Config** > **Virtual Address**). If you have configured a floating IP address, these settings sync automatically. Otherwise, you must configure these settings independently on each peer.
GlobalProtect Portal Configuration	GlobalProtect portal configuration synchronization is dependent on whether you have configured the Virtual Addresses to use Floating IP addresses (**Network** > **GlobalProtect** > **Portals**). If you have configured a floating IP address, the GlobalProtect portal configuration settings sync automatically. Otherwise, you must configure the portal settings independently on each peer.
GlobalProtect Gateway Configuration	GlobalProtect gateway configuration synchronization is dependent on whether you have configured the Virtual Addresses to use Floating IP addresses (**Network** > **GlobalProtect** > **Gateways**). If you have configured a floating IP address, the GlobalProtect gateway configuration settings sync automatically. Otherwise, you must configure the gateway settings independently on each peer.
QoS	QoS configuration synchronizes only if you have enabled **QoS Sync** (**Device** > **High Availability** > **Active/Active Config** > **Packet Forwarding**). You might choose not to sync QoS setting if, for example, you have different bandwidth on each link or different latency through your service providers.
LLDP	No LLDP state or individual firewall data is synchronized in an active/active configuration (**Network** > **Network Profiles** > **LLDP**).
BFD	No BFD configuration or BFD session data is synchronized in an active/active configuration (**Network** > **Network Profiles** > **BFD Profile**).
IKE Gateways	IKE gateway configuration synchronization is dependent on whether you have configured the Virtual Addresses to use floating IP addresses (**Network** > **IKE Gateways**). If you have configured a floating IP address, the IKE gateway configuration settings sync automatically. Otherwise, you must configure the IKE gateway settings independently on each peer.
Master Key	The master key must be identical on each firewall in the HA pair, but you must manually enter it on each firewall (**Device** > **Master Key and Diagnostics**). Before changing the master key, you must disable config sync on both peers (**Device** > **High Availability** > **General** > **Setup** and clear the **Enable Config Sync** check box) and then re-enable it after you change the keys.

https://t.me/learningnets

Configuration Item	What Doesn't Sync in Active/Active?
Reports, logs, and Dashboard Settings	Log data, reports, and dashboard data and settings (column display, widgets) are not synced between peers. Report configuration settings, however, are synced.
HA settings	• **Device** > **High Availability** • (The exception is **Device** > **High Availability** > **Active/Active Configuration** > **Virtual Addresses**, which do sync.)
Rule Usage Data	Rule usage data, such as hit count, Created, and Modified Dates, are not synced between peers. You need to log in to the each firewall to view the policy rule hit count data for each firewall or use Panorama to view information on the HA firewall peers.
Certificates for Device Management and Syslog Communication over SSL only	**Device** > **Certificate Management** > **Certificates** Certificates used for device management or for syslog communication over SSL don't synchronize with an HA peer.
Certificates in a Certificate Profile	**Device** > **Certificate Management** > **Certificate Profile**
SSL/TLS Service Profile for Device Management only	**Device** > **Certificate Management** > **SSL/TLS Service Profile** SSL/TLS Service Profile for Device Management doesn't synchronize with an HA peer.
Device-ID and IoT Security	IP address-to-device mappings and policy rule recommendations don't synchronize with an HA peer.

Synchronization of System Runtime Information

The following table summarizes what system runtime information is synchronized between HA peers.

Runtime Information	Config Synced?		HA Link	Details
	A/P	A/A		
Management Plane				
User to Group Mappings	Yes	Yes	HA1	
User Mappings across Virtual Systems	Yes	Yes	HA1	

https://t.me/learningnets

Runtime Information	Config Synced?		HA Link	Details
	A/P	A/A		
User to IP Address Mappings	Yes	Yes	HA1	
DHCP Lease (as server)	Yes	Yes	HA1	If the PAN-OS versions on the HA peers don't match, the DHCP Lease (as server) config information won't sync.
DNS Cache	No	No	N/A	
FQDN Refresh	No	No	N/A	
IKE Keys (phase 2)	Yes	Yes	HA1	
Forward Information Base (FIB)	Yes	Yes	HA1	
PAN-DB URL Cache	Yes	No	HA1	This is synchronized upon database backup to disk (every eight hours, when URL database version updates), or when the firewall reboots.
Content (manual sync)	Yes	Yes	HA1	
PPPoE, PPPoE Lease	Yes	Yes	HA1	
DHCP Client Settings and Lease	Yes	Yes	HA1	If the PAN-OS versions on the HA peers don't match, the DHCP Client Settings and Lease config information won't sync.
SSL VPN Logged in User List	Yes	Yes	HA1	
Dataplane				
Session Table	Yes	Yes	HA2	• Active/passive peers do not sync ICMP or host session information.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Runtime Information	Config Synced?		HA Link	Details
	A/P	A/A		
				• Active/active peers do not sync host session, multicast session, or BFD session information. *A host session is a session terminated on one of the firewall interfaces, such as an ICMP session pinging one of the firewall interfaces or a GP tunnel.*
ARP Table	Yes	No	HA2	
Neighbor Discovery (ND) Table	Yes	No	HA2	
MAC Table	Yes	No	HA2	
IPSec Sequence Number (anti-replay)	Yes	Yes	HA2	
DoS Block List Entries	No	No	N/A	
Virtual MAC	Yes	Yes	HA2	
SCTP Associations	Yes	No	HA2	

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Monitoring

To forestall potential issues and to accelerate incidence response when needed, the firewall provides intelligence about traffic and user patterns using customizable and informative reports. The dashboard, Application Command Center (ACC), reports, and logs on the firewall allow you to monitor activity on your network. You can monitor the logs and filter the information to generate reports with predefined or customized views. For example, you can use the predefined templates to generate reports on user activities or analyze the reports and logs to interpret unusual behavior on your network and generate a custom report on the traffic pattern. For a visually engaging presentation of network activity, the dashboard and the ACC include widgets, charts, and tables with which you can interact to find the information you care about. In addition, you can configure the firewall to forward monitored information as email notifications, syslog messages, SNMP traps, and NetFlow records to external services.

- Use the Dashboard
- Use the Application Command Center
- Use the App Scope Reports
- Use the Automated Correlation Engine
- Take Packet Captures
- Monitor Applications and Threats
- View and Manage Logs
- Monitor Block List
- View and Manage Reports
- View Policy Rule Usage
- Use External Services for Monitoring
- Configure Log Forwarding
- Configure Email Alerts
- Use Syslog for Monitoring
- SNMP Monitoring and Traps
- Forward Logs to an HTTP(S) Destination
- NetFlow Monitoring

https://t.me/learningnets

Use the Dashboard

The **Dashboard** tab widgets show general firewall information, such as the software version, the operational status of each interface, resource utilization, and up to 10 of the most recent entries in the threat, configuration, and system logs. All of the available widgets are displayed by default, but each administrator can remove and add individual widgets, as needed. Click the refresh icon to update the dashboard or an individual widget. To change the automatic refresh interval, select an interval from the drop-down (**1 min**, **2 mins**, **5 mins**, or **Manual**). To add a widget to the dashboard, click the widget drop-down, select a category and then the widget name. To delete a widget, click in the title bar. The following table describes the dashboard widgets.

Dashboard Charts	Descriptions
Top Applications	Displays the applications with the most sessions. The block size indicates the relative number of sessions (mouse-over the block to view the number), and the color indicates the security risk—from green (lowest) to red (highest). Click an application to view its application profile.
Top High Risk Applications	Similar to Top Applications, except that it displays the highest-risk applications with the most sessions.
General Information	Displays the firewall name, model, PAN-OS software version, the application, threat, and URL filtering definition versions, the current date and time, and the length of time since the last restart.
Interface Status	Indicates whether each interface is up (green), down (red), or in an unknown state (gray).
Threat Logs	Displays the threat ID, application, and date and time for the last 10 entries in the Threat log. The threat ID is a malware description or URL that violates the URL filtering profile.
Config Logs	Displays the administrator username, client (Web or CLI), and date and time for the last 10 entries in the Configuration log.
Data Filtering Logs	Displays the description and date and time for the last 60 minutes in the Data Filtering log.
URL Filtering Logs	Displays the description and date and time for the last 60 minutes in the URL Filtering log.
System Logs	Displays the description and date and time for the last 10 entries in the System log. *A **`Config installed`** entry indicates configuration changes were committed successfully.*

https://t.me/learningnets

Dashboard Charts	Descriptions
System Resources	Displays the Management CPU usage, Data Plane usage, and the Session Count, which displays the number of sessions established through the firewall.
Logged In Admins	Displays the source IP address, session type (Web or CLI), and session start time for each administrator who is currently logged in.
ACC Risk Factor	Displays the average risk factor (1 to 5) for the network traffic processed over the past week. Higher values indicate higher risk.
High Availability	If high availability (HA) is enabled, indicates the HA status of the local and peer firewall—green (active), yellow (passive), or black (other). For more information about HA, see High Availability.
Locks	Shows configuration locks taken by administrators.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Use the Application Command Center

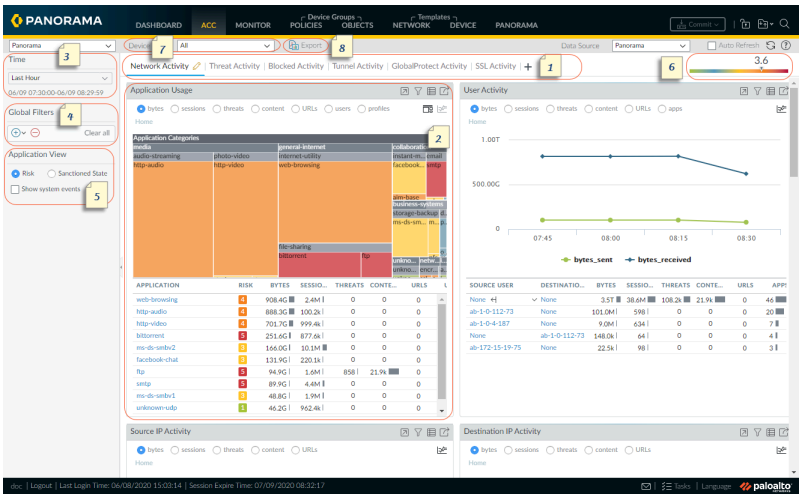
The Application Command Center (ACC) is an interactive, graphical summary of the applications, users, URLs, threats, and content traversing your network.The ACC uses the firewall logs to provide visibility into traffic patterns and actionable information on threats. The ACC layout includes a tabbed view of network activity, threat activity, and blocked activity and each tab includes pertinent widgets for better visualization of network traffic. The graphical representation allows you to interact with the data and visualize the relationships between events on the network, so that you can uncover anomalies or find ways to enhance your network security rules. For a personalized view of your network, you can also add a custom tab and include widgets that allow you to drill down into the information that is most important to you.

*ACC data, including ACC widgets and exported ACC reports, use Security policy rule data that you enabled to **Log at Session End**. If some data you expect to view in the ACC is not displayed, view your Traffic and Threat logs to determine the correct Security policy rule to modify as needed so all new logs generated that match the Security policy rule are viewable in the ACC.*

- ACC—First Look
- ACC Tabs
- ACC Widgets (Widget Descriptions)
- ACC Filters
- Interact with the ACC
- Use Case: ACC—Path of Information Discovery

ACC—First Look

Take a quick tour of the ACC.



https://t.me/learningnets

ACC—First Look		
	**Tabs**	The ACC includes three predefined tabs that provide visibility into network traffic, threat activity, and blocked activity. For information on each tab, see ACC Tabs.
	**Widgets**	Each tab includes a default set of widgets that best represent the events/trends associated with the tab. The widgets allow you to survey the data using the following filters: • bytes (in and out) • sessions • content (files and data) • URL categories • threats (and count) For information on each widget, see ACC Widgets.
	**Time**	The charts or graphs in each widget provide a summary and historic view. You can choose a custom range or use the predefined time periods that range from the last 15 minutes up to the last 90 days or last 30 calendar days. The selected time period applies across all tabs in the ACC. The time period used to render data, by default, is the **Last Hour** updated in 15 minute intervals. The date and time interval are displayed onscreen, for example at 11:40, the time range is 01/12 10:30:00-01/12 11:29:59.
	**Global Filters**	The Global Filters allow you to set the filter across all widgets and all tabs. The charts/graphs apply the selected filters before rendering the data. For information on using the filters, see ACC Filters.
	Application View	The application view allows you filter the ACC view by either the sanctioned and unsanctioned applications in use on your network, or by the risk level of the applications in use on your network. Green indicates sanctioned applications, blue unsanctioned applications, and yellow indicates applications that are partially sanctioned. Partially sanctioned applications are those that have a mixed sanctioned state; it indicates that the application is inconsistently tagged as sanctioned, for example it might be sanctioned on one or more

https://t.me/learningnets

ACC—First Look		
		virtual systems on a firewall enabled for multiple virtual systems or across one or more firewalls within a device group on Panorama.
6	**Risk Factor**	The risk factor (1=lowest to 5=highest) indicates the relative risk based on the applications used on your network. The risk factor uses a variety of factors to assess the associated risk levels, such as whether the application can share files, is it prone to misuse or does it try to evade firewalls, it also factors in the threat activity and malware as seen through the number of blocked threats, compromised hosts or traffic to malware hosts/ domains.
7	**Source**	The data used for the ACC display. The options vary on the firewall and on Panorama. On the firewall, if enabled for multiple virtual systems, you can use the **Virtual System** drop-down to change the ACC display to include data from all virtual systems or just a selected virtual system. On Panorama, you can select the **Device Group** drop-down to change the ACC display to include data from all device groups or just a selected device group. Additionally, on Panorama, you can change the **Data Source** as **Panorama** data or **Remote Device Data**. **Remote Device Data** is only available when all the managed firewalls are on PAN-OS 7.0.0 or later. When you filter the display for a specific device group, **Panorama** data is used as the data source.
8	**Export**	You can export the widgets displayed in the currently selected tab as a PDF. The PDF is downloaded and saved to the downloads folder associated with your web browser, on your computer.

ACC Tabs

The ACC includes the following predefined tabs for viewing network activity, threat activity, and blocked activity.

https://t.me/learningnets

Tab	Description
Network Activity	Displays an overview of traffic and user activity on your network including: • Top applications in use • Top users who generate traffic (with a drill down into the bytes, content, threats or URLs accessed by the user) • Most used security rules against which traffic matches occur In addition, you can also view network activity by source or destination zone, region, or IP address, ingress or egress interfaces, and GlobalProtect host information such as the operating systems of the devices most commonly used on the network.
Threat Activity	Displays an overview of the threats on the network, focusing on the top threats: vulnerabilities, spyware, viruses, hosts visiting malicious domains or URLs, top WildFire submissions by file type and application, and applications that use non-standard ports. The Compromised Hosts widget in this tab (the widget is supported on some platforms only), supplements detection with better visualization techniques; it uses the information from the correlated events tab (**Automated Correlation Engine** > **Correlated Events**) to present an aggregated view of compromised hosts on your network by source users/IP addresses and sorted by severity.
Blocked Activity	Focuses on traffic that was prevented from coming into the network. The widgets in this tab allow you to view activity denied by application name, username, threat name, blocked content—files and data that were blocked by a file blocking profile. It also lists the top security rules that were matched on to block threats, content, and URLs.
Tunnel Activity	Displays the activity of tunnel traffic that the firewall inspected based on your tunnel inspection policies. Information includes tunnel usage based on tunnel ID, monitor tag, user, and tunnel protocols such as Generic Routing Encapsulation (GRE), General Packet Radio Service (GPRS) Tunneling Protocol for User Data (GTP-U), and non-encrypted IPSec.
GlobalProtect Activity	Displays an overview of user activity in your GlobalProtect deployment. Information includes the number of users and number of times users connected, the gateways to which users connected, the number of connection failures and the failure reason, a summary of authentication methods and GlobalProtect app versions used, and the number of endpoints that are quarantined. In addition, this tab displays a chart view summary of devices that have been quarantined. Use the toggle at the top of the chart to view the quarantined devices by the actions that caused

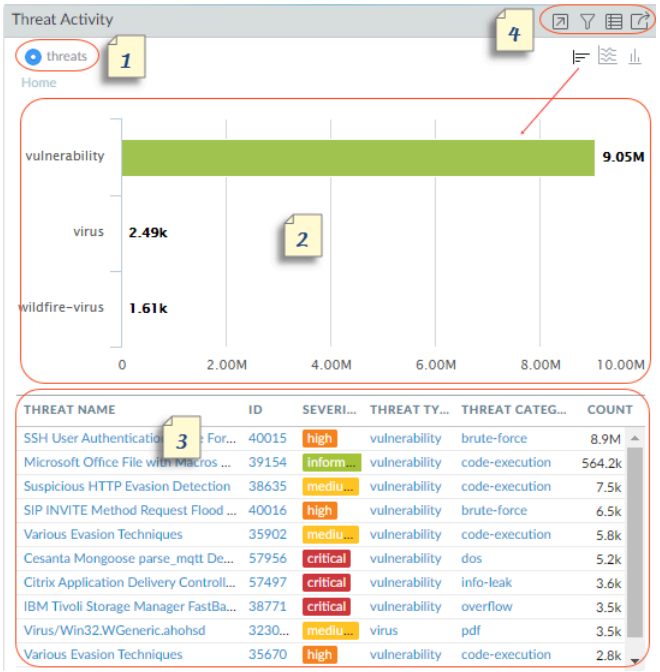
https://t.me/learningnets

Tab	Description
	GlobalProtect to quarantine the device, the reason GlobalProtect quarantined the device, and the location of the quarantined devices.
SSL Activity	Displays an overview of TLS/SSL decryption activity on the firewall. Information includes successful and unsuccessful decryption activity in your network, decryption failure reasons such as protocol, certificate, and version issues, TLS versions, key exchange algorithms, and the amount and type of decrypted and undecrypted traffic. Use the ACC information to evaluate how decryption is working on your network and then use the Decryption Log to drill down into details.

You can also Interact with the ACC to create customized tabs with custom layout and widgets that meet your network monitoring needs, export the tab and share with another administrator.

ACC Widgets

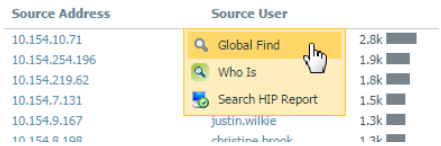
The widgets on each tab are interactive; you can set the ACC Filters and drill down into the details for each table or graph, or customize the widgets included in the tab to focus on the information you need. For details on what each widget displays, see Widget Descriptions.



Widgets		
1	**View**	You can sort the data by bytes, sessions, threats, count, content, URLs, malicious, benign, files,

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Widgets		
		applications, data, profiles, objects, users. The available options vary by widget.
2	**Graph**	The graphical display options are treemap, line graph, horizontal bar graph, stacked area graph, stacked bar graph, and map. The available options vary by widget; the interaction experience also varies with each graph type. For example, the widget for Applications using Non-Standard Ports allows you to choose between a treemap and a line graph. To drill down into the display, click into the graph. The area you click into becomes a filter and allows you to zoom into the selection and view more granular information on the selection.
3	**Table**	The detailed view of the data used to render the graph is provided in a table below the graph. You can interact with the table in several ways: • Click and set a local filter for an attribute in the table. The graph is updated and the table is sorted using the local filter. The information displayed in the graph and the table are always synchronized. • Hover over the attribute in the table and use the options available in the drop-down. 
4	**Actions**	**Maximize view**— Allows you enlarge the widget and view the table in a larger screen space and with more viewable information. **Set up local filters**—Allows you to add ACC Filters to refine the display within the widget. Use these filters to customize the widgets; these customizations are retained between logins. **Jump to logs**—Allows you to directly navigate to the logs (**Monitor** > **Logs** > **<log-type>** tab). The logs are

https://t.me/learningnets

Widgets		
		filtered using the time period for which the graph is rendered. If you have set local and global filters, the log query concatenates the time period and the filters and only displays logs that match the combined filter set.  **Export**—Allows you to export the graph as a PDF. The PDF is downloaded and saved on your computer. It is saved in the Downloads folder associated with your web browser.

Widget Descriptions

Each tab on the ACC includes a different set of widgets.

Widget	Description
Network Activity—Displays an overview of traffic and user activity on your network.	
Application Usage	The table displays the top ten applications used on your network, all the remaining applications used on the network are aggregated and displayed as other. The graph displays all applications by application category, sub category, and application. Use this widget to scan for applications being used on the network, it informs you about the predominant applications using bandwidth, session count, file transfers, triggering the most threats, and accessing URLs. Sort attributes: bytes, sessions, threats, content, URLs Charts available: treemap, area, column, line (the charts vary by the sort by attribute selected)
User Activity	Displays the top ten most active users on the network who have generated the largest volume of traffic and consumed network resources to obtain content. Use this widget to monitor top users on usage sorted on bytes, sessions, threats, content (files and patterns), and URLs visited. Sort attributes: bytes, sessions, threats, content, URLs Charts available: area, column, line (the charts vary by the sort by attribute selected)
Source IP Activity	Displays the top ten IP addresses or hostnames of the devices that have initiated activity on the network. All other devices are aggregated and displayed as other.

https://t.me/learningnets

Widget	Description
	Sort attributes: bytes, sessions, threats, content, URLs Charts available: area, column, line (the charts vary by the sort by attribute selected)
Destination IP Activity	Displays the IP addresses or hostnames of the top ten destinations that were accessed by users on the network. Sort attributes: bytes, sessions, threats, content, URLs Charts available: area, column, line (the charts vary by the sort by attribute selected)
Source Regions	Displays the top ten regions (built-in or custom defined regions) around the world from where users initiated activity on your network. Sort attributes: bytes, sessions, threats, content, URLs Charts available: map, bar
Destination Regions	Displays the top ten destination regions (built-in or custom defined regions) on the world map from where content is being accessed by users on the network. Sort attributes: bytes, sessions, threats, content, URLs Charts available: map, bar
HIP Information	Displays information on the state of the hosts on which the GlobalProtect agent is running; the host system is a GlobalProtect endpoint. This information is sourced from entries in the HIP match log that are generated when the data submitted by the GlobalProtect app matches a HIP object or a HIP profile you have defined on the firewall. If you do not have HIP Match logs, this widget is blank. To learn how to create HIP objects and HIP profiles and use them as policy match criteria, see Configure HIP-Based Policy Enforcement. Sort attributes: profiles, objects, operating systems Charts available: bar
Rule Usage	Displays the top ten rules that have allowed the most traffic on the network. Use this widget to view the most commonly used rules, monitor the usage patterns, and to assess whether the rules are effective in securing your network. Sort attributes: bytes, sessions, threats, content, URLs Charts available: line
Ingress Interfaces	Displays the firewall interfaces that are most used for allowing traffic into the network. Sort attributes: bytes, bytes sent, bytes received

https://t.me/learningnets

Widget	Description
	Charts available: line
Egress Interfaces	Displays the firewall interfaces that are most used by traffic exiting the network. Sort attributes: bytes, bytes sent, bytes received Charts available: line
Source Zones	Displays the zones that are most used for allowing traffic into the network. Sort attributes: bytes, sessions, threats, content, URLs Charts available: line
Destination Zones	Displays the zones that are most used by traffic going outside the network. Sort attributes: bytes, sessions, threats, content, URLs Charts available: line
Threat Activity—Displays an overview of the threats on the network	
Compromised Hosts	Displays the hosts that are likely compromised on your network. This widget summarizes the events from the correlation logs. For each source user/IP address, it includes the correlation object that triggered the match and the match count, which is aggregated from the match evidence collated in the correlated events logs. For details see Use the Automated Correlation Engine. Available on the PA-5200 Series, PA-7000 Series, and Panorama. Sort attributes: severity (by default)
Hosts Visiting Malicious URLs	Displays the frequency with which hosts (IP address/hostnames) on your network have accessed malicious URLs. These URLs are known to be malware based on categorization in PAN-DB. Sort attributes: count Charts available: line
Hosts Resolving Malicious Domains	Displays the top hosts matching DNS signatures; hosts on the network that are attempting to resolve the hostname or domain of a malicious URL. This information is gathered from an analysis of the DNS activity on your network. It utilizes passive DNS monitoring, DNS traffic generated on the network, activity seen in the sandbox if you have configured DNS sinkhole on the firewall, and DNS reports on malicious DNS sources that are available to Palo Alto Networks customers. Sort attributes: count

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Widget	Description
	Charts available: line
Threat Activity	Displays the threats seen on your network. This information is based on signature matches in Antivirus, Anti-Spyware, and Vulnerability Protection profiles and viruses reported by WildFire. Sort attributes: threats Charts available: bar, area, column
WildFire Activity by Application	Displays the applications that generated the most WildFire submissions. This widget uses the malicious and benign verdict from the WildFire Submissions log. Sort attributes: malicious, benign Charts available: bar, line
WildFire Activity by File Type	Displays the threat vector by file type. This widget displays the file types that generated the most WildFire submissions and uses the malicious and benign verdict from the WildFire Submissions log. If this data is unavailable, the widget is empty. Sort attributes: malicious, benign Charts available: bar, line
Applications using Non Standard Ports	Displays the applications that are entering your network on non-standard ports. If you have migrated your firewall rules from a port-based firewall, use this information to craft policy rules that allow traffic only on the default port for the application. Where needed, make an exception to allow traffic on a non-standard port or create a custom application. Sort attributes: bytes, sessions, threats, content, URLs Charts available: treemap, line
Rules Allowing Applications On Non Standard Ports	Displays the security policy rules that allow applications on non-default ports. The graph displays all the rules, while the table displays the top ten rules and aggregates the data from the remaining rules as other. This information helps you identify gaps in network security by allowing you to assess whether an application is hopping ports or sneaking into your network. For example, you can validate whether you have a rule that allows traffic on any port except the default port for the application. Say for example, you have a rule that allow DNS traffic on its *application-default* port (port 53 is the standard port for DNS). This widget will display any rule that allows DNS traffic into your network on any port except port 53. Sort attributes: bytes, sessions, threats, content, URLs

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Widget	Description
	Charts available: treemap, line
Blocked Activity—Focuses on traffic that was prevented from coming into the network	
Blocked Application Activity	Displays the applications that were denied on your network, and allows you to view the threats, content, and URLs that you kept out of your network. Sort attributes: threats, content, URLs Charts available: treemap, area, column
Blocked User Activity	Displays user requests that were blocked by a match on an Antivirus, Anti-spyware, File Blocking or URL Filtering profile attached to Security policy rule. Sort attributes: threats, content, URLs Charts available: bar, area, column
Blocked Threats	Displays the threats that were successfully denied on your network. These threats were matched on antivirus signatures, vulnerability signatures, and DNS signatures available through the dynamic content updates on the firewall. Sort attributes: threats Charts available: bar, area, column
Blocked Content	Displays the files and data that was blocked from entering the network. The content was blocked because security policy denied access based on criteria defined in a File Blocking security profile or a Data Filtering security profile. Sort attributes: files, data Charts available: bar, area, column
Security Policies Blocking Activity	Displays the security policy rules that blocked or restricted traffic into your network. Because this widget displays the threats, content, and URLs that were denied access into your network, you can use it to assess the effectiveness of your policy rules. This widget does not display traffic that blocked because of deny rules that you have defined in policy. Sort attributes: threats, content, URLs Charts available: bar, area, column
GlobalProtect Activity—Displays information of user activity in your GlobalProtect deployment.	

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Widget	Description
Successful GlobalProtect Connection Activity	Displays a chart view of GlobalProtect connection activity over the selected time period. Use the toggle at the top of the chart to switch between connection statistics by users, portals and gateways, and location. Sort attributes: users, portals/gateways, location Charts available: bar, line
Unsuccessful GlobalProtect Connection Activity	Displays a chart view of unsuccessful GlobalProtect connection activity over the selected time period. Use the toggle at the top of the chart to switch between connection statistics by users, portals and gateways, and location. To help you identify and troubleshoot connection issues, you can also view the reasons chart or graph. For this chart, the ACC indicates the error, source user, public IP address and other information to help you identify and resolve the issue quickly. Sort attributes: users, portals/gateways, reasons, location Charts available: bar, line
GlobalProtect Deployment Activity	Displays a chart view summary of your deployment. Use the toggle at the top of the chart to view the distribution of users by authentication method, GlobalProtect app version, and operating system version. Sort attributes: auth method, globalprotect app version, os Charts available: bar, line
GlobalProtect Quarantine Activity	Displays a chart view summary of devices that have been quarantined. Use the toggle at the top of the chart to view the quarantined devices by the actions that caused GlobalProtect to quarantine the device, the reason GlobalProtect quarantined the device, and the location of the quarantined devices. Sort attributes: actions, reason, location Charts available: bar, line
SSL Activity—Displays information about SSL/TLS activity in your network.	
Traffic Activity	Shows SSL/TLS activity compared to non-SSL/TLS activity by total number of sessions or bytes.
SSL/TLS Activity	Shows successful TLS connections by TLS version and application or SNI. This widget helps you understand how much risk you are taking on by allowing weaker TLS protocol versions. Identifying applications and SNIs that use weak protocols enables you to evaluate each one and decide whether you need to allow access to it for business reasons. If you don't need the application for business purposes, you

 ©2022 Palo Alto Networks, Inc.

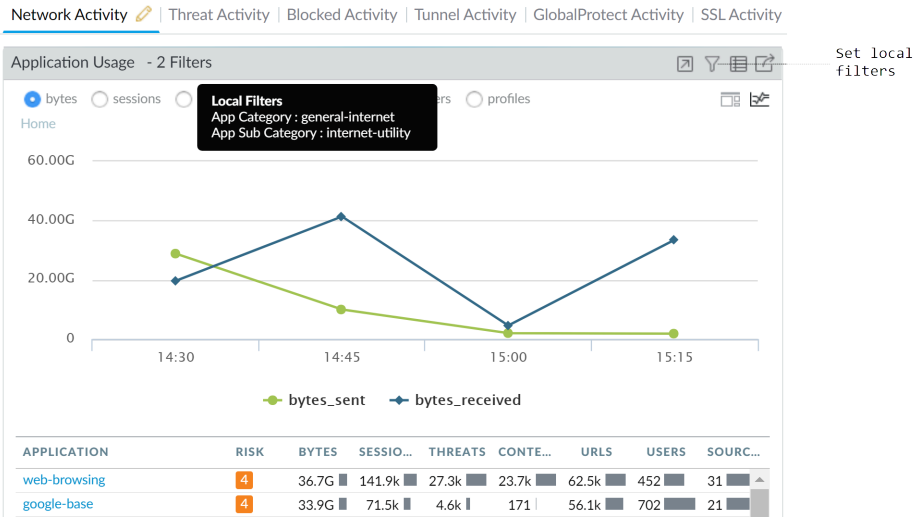
https://t.me/learningnets

Widget	Description
	may want to block the traffic instead of allowing it. Click an application or an SNI to drill down and see detailed information.
Decryption Failure Reasons	Shows the reasons for decryption failures, such as certificate or protocol issues, by SNI. Use this information to detect problems caused by Decryption policy or profile misconfiguration or by traffic that uses weak protocols or algorithms. Click a failure reason to drill down and isolate the number of sessions per SNI or click an SNI to see the failures for that SNI.
Successful TLS Version Activity	Shows the amount of decrypted and non-decrypted traffic by sessions or bytes. Traffic that was not decrypted may be excepted from decryption by policy, policy misconfiguration, or by being on the Decryption Exclusion List (**Device** > **Certificate Management** > **SSL Decryption Exclusion**).
Successful Key Exchange Activity	Shows successful key exchange activity per algorithm, by application or by SNI. Click a key exchange algorithm to see the activity for just that algorithm or click an application or SNI to view the key exchange activity for that application or SNI.

ACC Filters

The graphs and tables on the ACC widgets allow you to use filters to narrow the scope of data that is displayed, so that you can isolate specific attributes and analyze information you want to view in greater detail. The ACC supports the simultaneous use of widget and global filters.

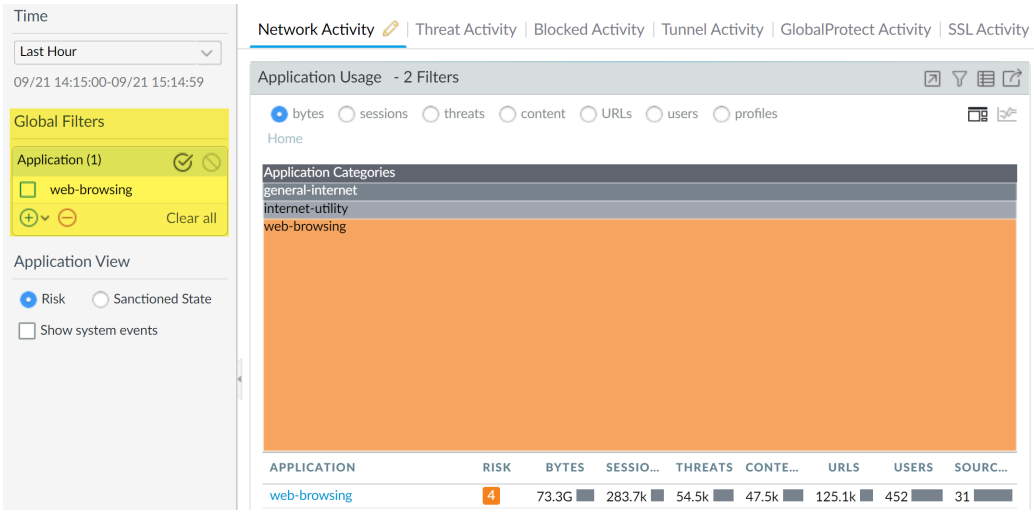
- **Widget Filters**—Apply a widget filter, which is a filter that is *local* to a specific widget. A widget filter allows you to interact with the graph and customize the display so that you can drill down in to the details and access the information you want to monitor on a specific widget. To create a widget filter that is persistent across reboots, you must use the **Set Local Filter** option.



 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

- **Global filters**—Apply global filters across all the tabs in the ACC. A global filter allows you to pivot the display around the details you care about right now and exclude the unrelated information from the current display. For example, to view all events relating to a specific user and application, you can apply the username and the application as a global filter and view only information pertaining to the user and the application through all the tabs and widgets on the ACC. Global filters are not persistent.



You can apply global filters in three ways:

- **Set a global filter from a table**—Select an attribute from a table in any widget and apply the attribute as a global filter.
- **Add a widget filter to a global filter**—Hover over the attribute and click the arrow icon to the right of the attribute. This option allows you to elevate a local filter used in a widget, and apply the attribute globally to update the display across all the tabs on the ACC.
- **Define a global filter**—Define a filter using the **Global Filters** pane on the ACC.

See Interact with the ACC for details on using these filters.

Interact with the ACC

To customize and refine the ACC display, you can add, delete, export and import tabs, add and delete widgets, set local and global filters, and interact with the widgets.

- Add a tab.
  1. Select the + icon along the list of tabs.
  2. Add a **View Name**. This name will be used as the name for the tab. You can add up to five tabs.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

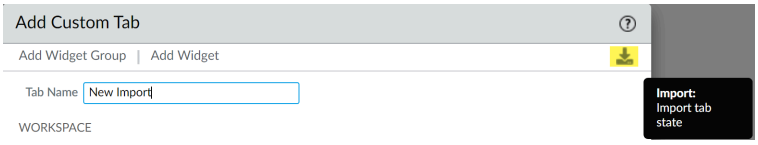
- Edit a tab.

  Select the tab, and click the pencil icon next to the tab name, to edit the tab. For example Threat Activity .

  Editing a tab allows you to add or delete or reset the widgets that are displayed in the tab. You can also change the widget layout in the tab.

  To save the tab as the default tab, select .

- Export and Import tabs.
  1. Select the tab, and click the pencil icon next to the tab name, to edit the tab.
  2. Select the icon to export the current tab as a .txt file. You can share this .txt file with another administrator.
  3. To import the tab as a new tab on another firewall, select the + icon along the list of tabs, and add a name and click the import icon, browse to select the .txt file.



- See what widgets are included in a tab.
  1. Select the tab, and click on the pencil icon to edit it.
  2. Select the **Add Widget** drop-down and verify the widgets that have the check boxes selected.

- Add a widget or a widget group.
  1. Add a new tab or edit a predefined tab.
  2. Select **Add Widget**, and then select the check box that corresponds to the widget you want to add. You can select up to a maximum of 12 widgets.
  3. (Optional) To create a 2-column layout, select **Add Widget Group**. You can drag and drop widgets into the 2-column display. As you drag the widget into the layout, a placeholder will display for you to drop the widget.

     You cannot name a widget group.

- Delete a tab or a widget group/ widget.
  1. To delete a custom tab, select the tab and click the X icon. Custom_threat_user_activity

     You cannot delete a predefined tab.

  2. To delete a widget group/widget, edit the tab and in the workspace section, click the [X] icon on the right. You cannot undo a deletion.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

- Reset the default widgets in a tab.

  On a predefined tab, such as the **Blocked Activity** tab, you can delete one or more widgets. If you want to reset the layout to include the default set of widgets for the tab, edit the tab and click **Reset View**.

- Zoom in on the details in an area, column, or line graph.

  Watch how the zoom-in capability works.

  Click and drag an area in the graph to zoom in. For example, when you zoom into a line graph, it triggers a re-query and the firewall fetches the data for the selected time period. It is not a mere magnification.

- Use the table drop-down to find more information on an attribute.
  1. Hover over an attribute in a table to see the drop-down.
  2. Click into the drop-down to view the available options.
     - **Global Find**—Use Global Find to Search the Firewall or Panorama Management Server for references to the attribute (username/IP address, object name, policy rule name, threat ID, or application name) anywhere in the candidate configuration.
     - **Value**—Displays the details of the threat ID, or application name, or address object.
     - **Who Is**—Performs a domain name (*WHOIS*) lookup for the IP address. The lookup queries databases that store the registered users or assignees of an Internet resource.
     - **Search HIP Report**—Uses the username or IP address to find matches in a HIP Match report.

- Set a widget filter.

  You can also click an attribute in the table (below the graph) to apply it as a widget filter.

  1. Select a widget and click the icon.
  2. Click the icon to add the filters you want to apply.
  3. Click **Apply**. These filters are persistent across reboots.

     The active widget filters are indicated next to the widget name.

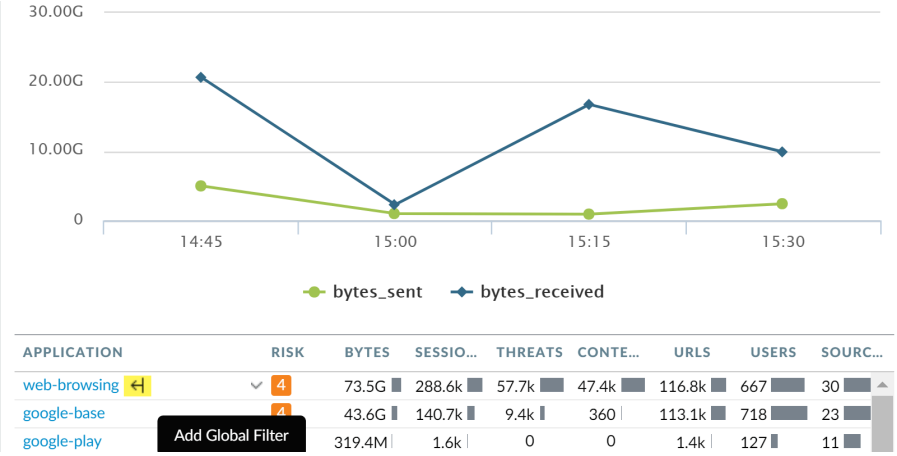
- Negate a widget filter
  1. Click the icon to display the Setup Local Filters dialog.
  2. Add a filter, and then click the negate icon.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

- Set a global filter from a table.

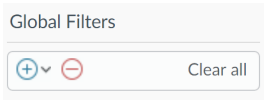
  Hover over an attribute in the table below the chart and click the arrow icon to the right of the attribute.



- Set a global filter using the Global Filters pane.

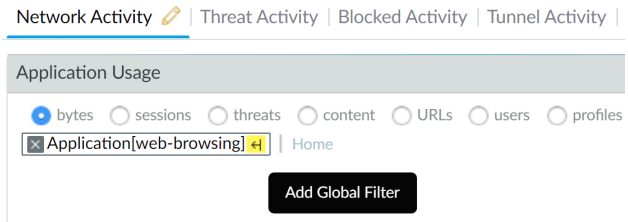
  Watch global filters in action.

  1. Locate the **Global Filters** pane on the left side of the ACC.



  2. Click the icon to view the list of filters you can apply.

- Promote a widget filter to a global filter.
  1. On any table in a widget, click the link for an attribute. This sets the attribute as a widget filter.
  2. To promote the filter to be a global filter, select the arrow to the right of the filter.



- Remove a filter.

  Click the icon to remove a filter.

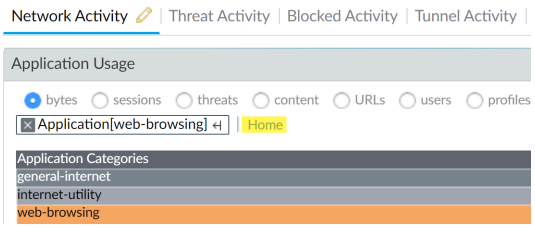
  - For global filters: It is located in the Global Filters pane.
  - For widget filters: Click the icon to display the Setup Local Filters dialog, then select the filter, and click the icon.

https://t.me/learningnets

- Clear all filters.
  - For global filters: Click the **Clear All** button under Global Filters.
  - For widget filters: Select a widget and click the icon. Then click the **Clear All** button in the Setup Local Filters dialog.

- See what filters are in use.
  - For global filters: The number of global filters applied are displayed on the left pane under Global Filters.
  - For widget filters: The number of widget filters applied on a widget are displayed next to the widget name. To view the filters, click the icon.

- Reset the display on a widget.
  - If you set a widget filter or drill into a graph, click the **Home** link to reset the display in the widget.



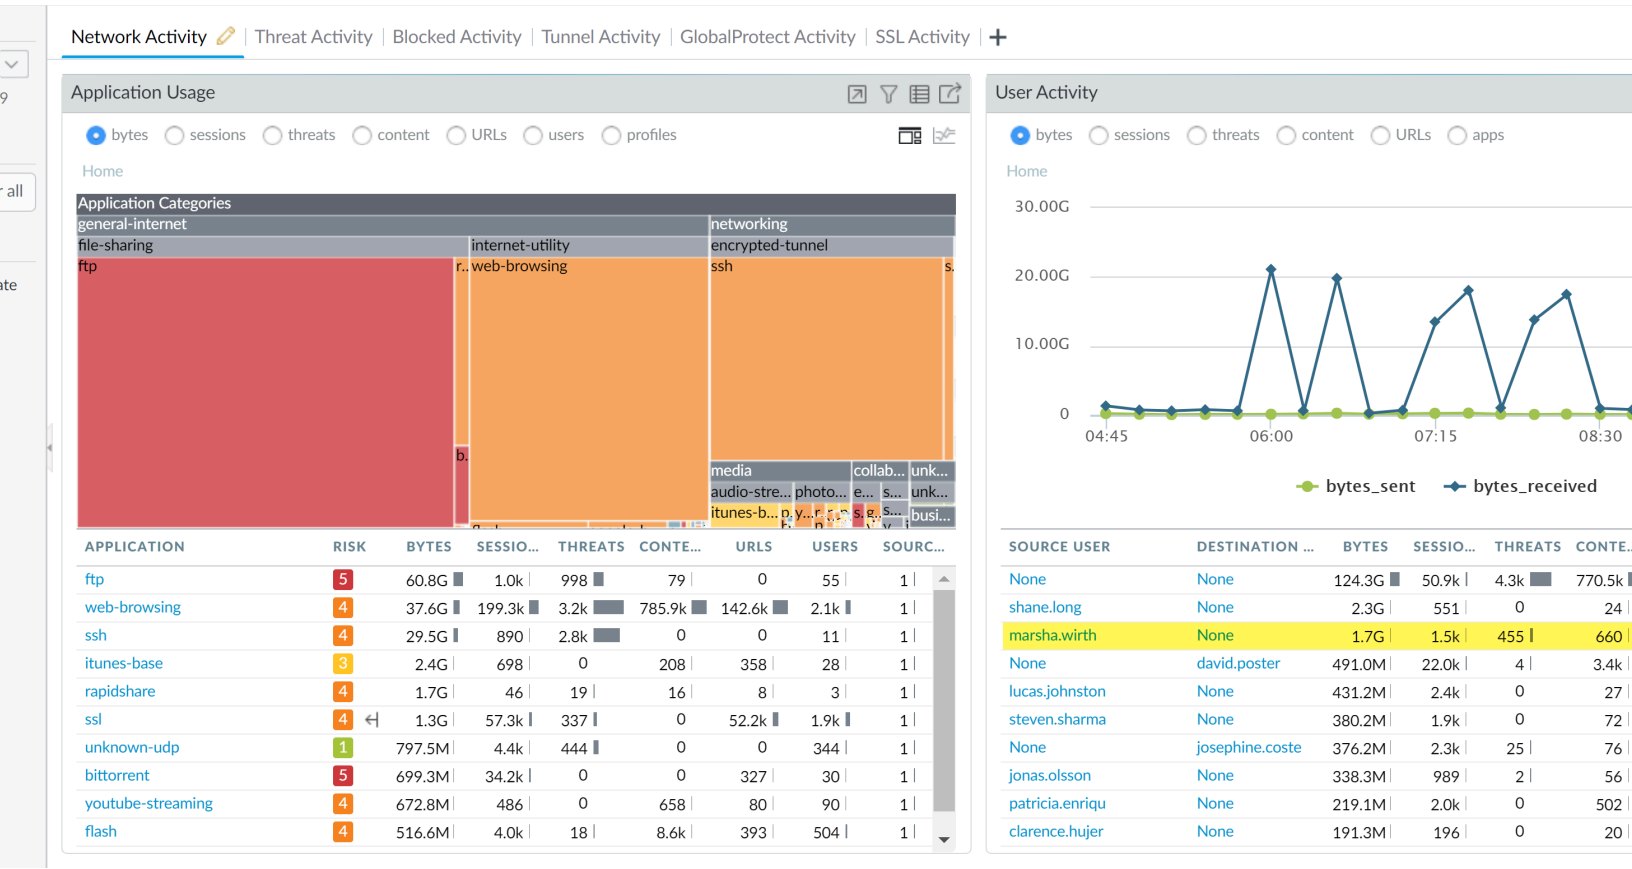
Use Case: ACC—Path of Information Discovery

The ACC has a wealth of information that you can use as a starting point for analyzing network traffic. Let's look at an example on using the ACC to uncover events of interest. This example illustrates how you can use the ACC to ensure that legitimate users can be held accountable for their actions, detect and track unauthorized activity, and detect and diagnose compromised hosts and vulnerable systems on your network.

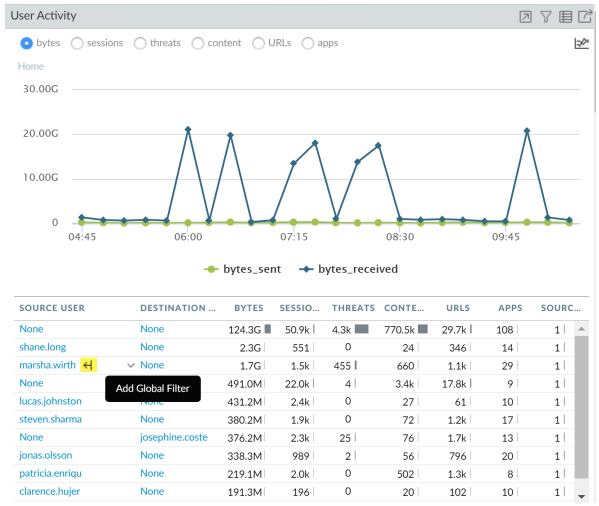
The widgets and filters in the ACC give you the capability to analyze the data and filter the views based on events of interest or concern. You can trace events that pique your interest, directly export a PDF of a tab, access the raw logs, and save a personalized view of the activity that you want to track. These capabilities make it possible for you to monitor activity and develop policies and countermeasures for fortifying your network against malicious activity. In this section, you will Interact with the ACC widgets across different tabs, drill down using widget filters, and pivot the ACC views using global filters, and export a PDF for sharing with incidence response or IT teams.

At first glance, you see the Application Usage and User Activity widgets in the **ACC** > **Network Activity** tab. The User Activity widget shows that user Marsha Wirth has transferred 154 Megabytes of data during the last hour. This volume is nearly six times more than any other user on the network. To see the trend over the past few hours, expand the **Time** period to the **Last 6 Hrs**, and now Marsha's activity has been 1.7 Gigabytes over 1,500 sessions and has triggered 455 threats signatures.

https://t.me/learningnets



Because Marsha has transferred a large volume of data, apply her username as a global filter (ACC Filters) and pivot all the views in the ACC to Marsha's traffic activity.

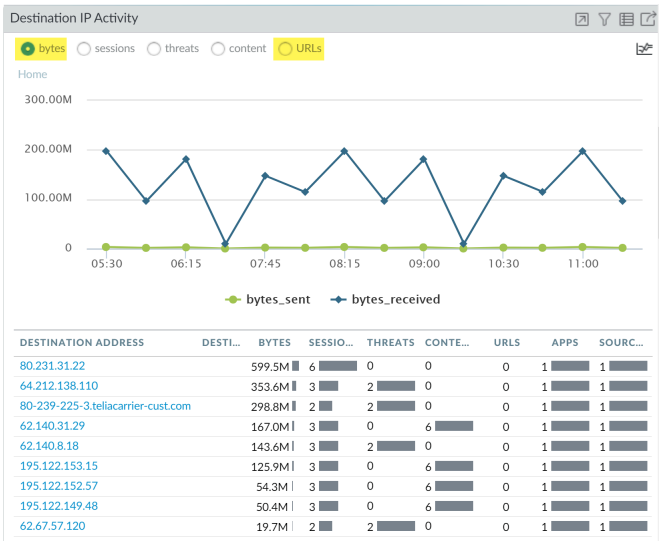


The Application Usage tab now shows that the top application that Martha used was rapidshare, a Swiss-owned file-hosting site that belongs to the file-sharing URL category. For further investigation, add rapidshare as a global filter, and view Marsha's activity in the context of rapidshare.

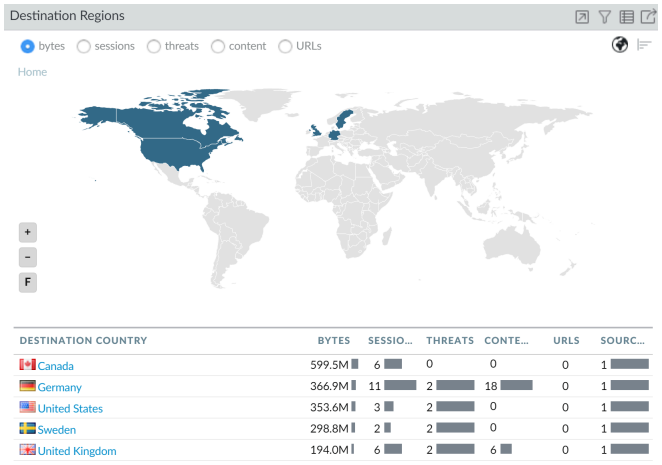
Consider whether you want to sanction rapidshare for company use. Should you allow uploads to this site and do you need a QoS policy to limit bandwidth?

https://t.me/learningnets

To view which IP addresses Marsha has communicated with, check the **Destination IP Activity** widget, and view the data by bytes and by URLs.

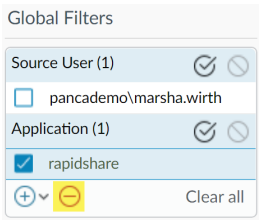


To find out which countries Marsha communicated with, sort on **sessions** in the **Destination Regions** widget.



From this data, you can confirm that Marsha, a user on your network, has established sessions in Canada, Germany, Sweden, United Kingdom, and the United States. She logged 2 threats in her sessions with each destination country.

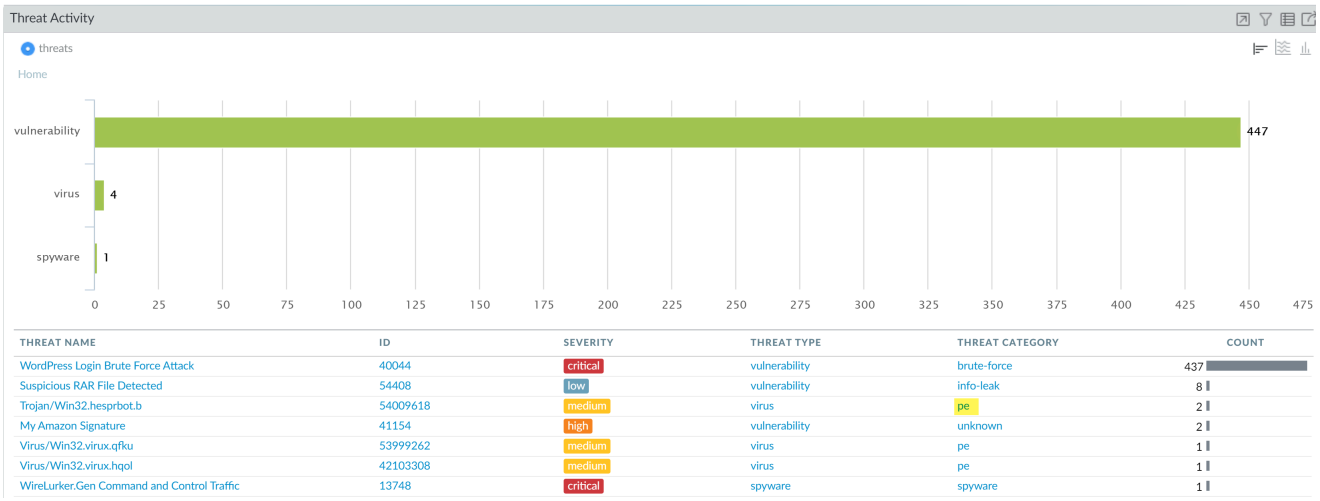
To look at Marsha's activity from a threat perspective, remove the global filter for rapidshare.



In the **Threat Activity** widget on the **Threat Activity** tab, view the threats. The widget displays that her activity had triggered a match for 452 vulnerabilities in the brute force, information

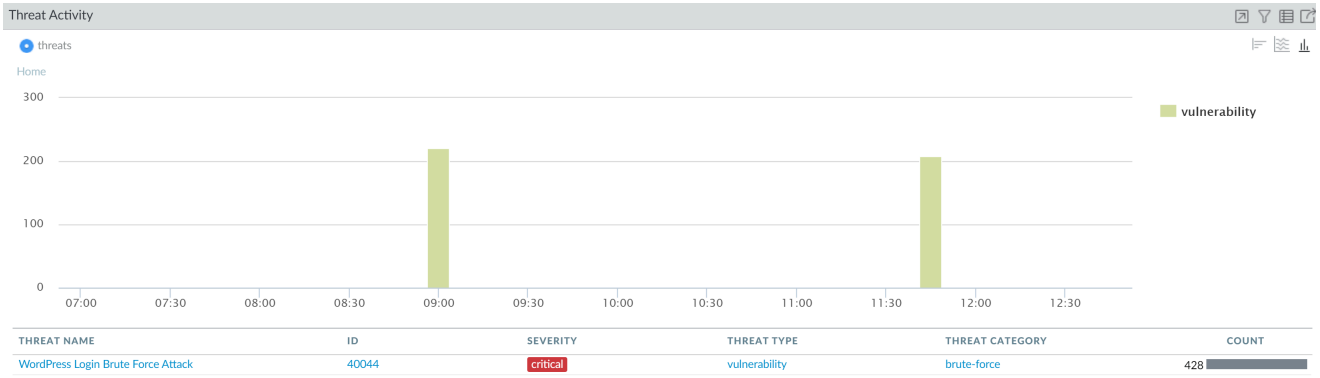
https://t.me/learningnets

leak, portable executable (PE) and spyware threat category. Several of these vulnerabilities are of critical severity.



THREAT NAME	ID	SEVERITY	THREAT TYPE	THREAT CATEGORY	COUNT
WordPress Login Brute Force Attack	40044	critical	vulnerability	brute-force	437
Suspicious RAR File Detected	54408	low	vulnerability	info-leak	8
Trojan/Win32.hesprbot.b	54009618	medium	virus	pe	2
My Amazon Signature	41154	high	vulnerability	unknown	2
Virus/Win32.virux.qfku	53999262	medium	virus	pe	1
Virus/Win32.virux.hqol	42103308	medium	virus	pe	1
WireLurker.Gen Command and Control Traffic	13748	critical	spyware	spyware	1

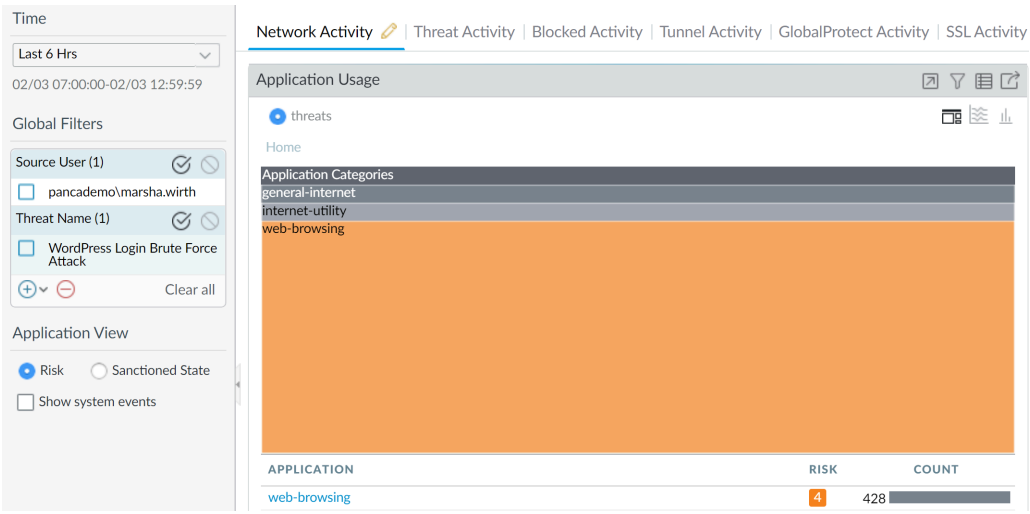
To further drill-down into each vulnerability, click into the graph and narrow the scope of your investigation. Each click automatically applies a local filter on the widget.



THREAT NAME	ID	SEVERITY	THREAT TYPE	THREAT CATEGORY	COUNT
WordPress Login Brute Force Attack	40044	critical	vulnerability	brute-force	428

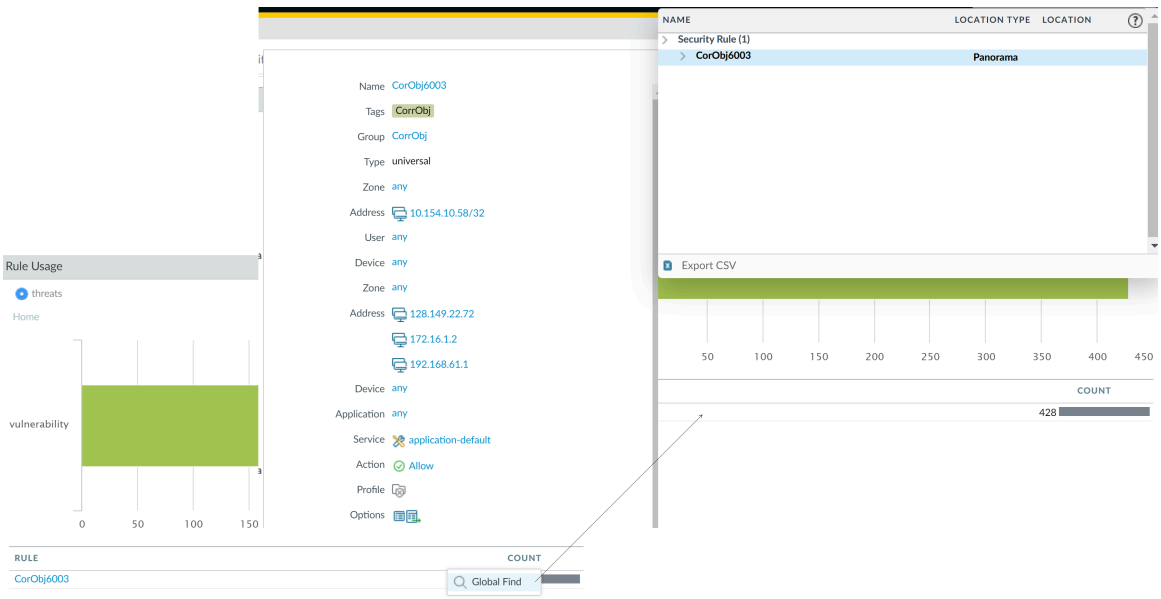
To investigate each threat by name, you can create a global filter for say, **WordPress Login Brute Force Attack**. Then, view the **User Activity widget** in the **Network Activity** tab. The tab is automatically filtered to display threat activity for Marsha (notice the global filters in the screenshot).

https://t.me/learningnets



Notice that this Microsoft code-execution vulnerability was triggered over email, by the imap application. You can now establish that Martha has IE vulnerabilities and email attachment vulnerabilities, and perhaps her computer needs to be patched. You can now either navigate to the **Blocked Threats** widget in the **Blocked Activity** tab to check how many of these vulnerabilities were blocked.

Or, you can check the **Rule Usage** widget on the **Network Activity** tab to discover how many vulnerabilities made it into your network and which security rule allowed this traffic, and navigate directly to the security rule using the **Global Find** capability.

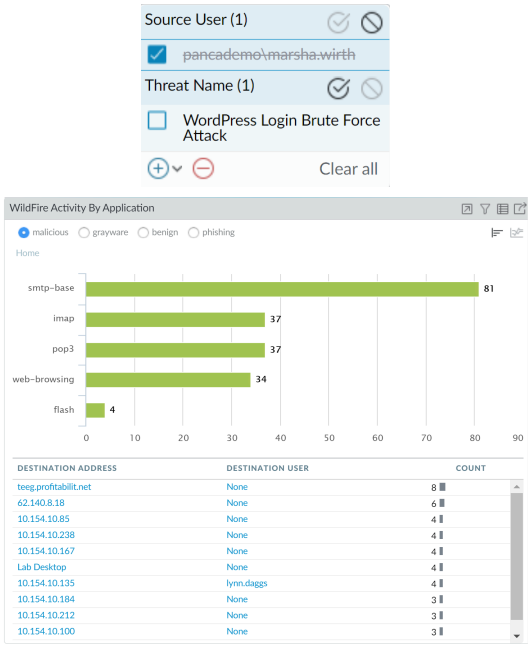


Then, drill into the attackers using web-browsing to attack target destination. Consider modifying the security policy rule to restrict these malicious IP addresses or more narrowly defining which IP addresses can access your network resources.

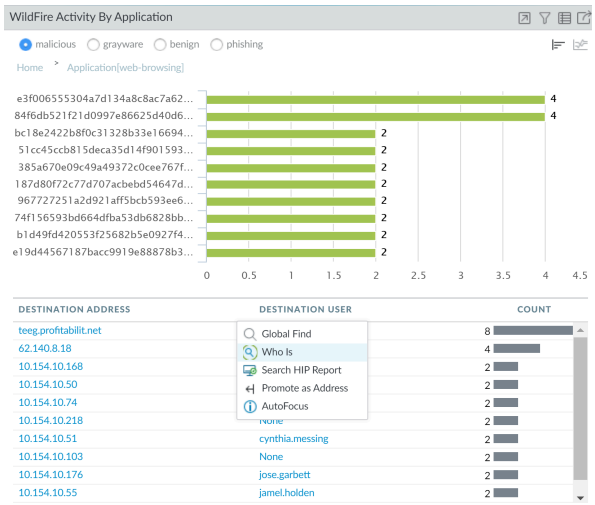
To review if any threats were logged over web-browsing, check Marsha's activity in the **WildFire Activity by Application** widget in the **Threat Activity** tab. You can confirm that Marsha had no malicious activity, but to verify that other no other user was compromised by the web-browsing

https://t.me/learningnets

application, negate Marsha as a global filter and look for other users who triggered threats over web-browsing.



Click into the bar for imap in the graph and drill into the inbound threats associated with the application. To find out who an IP address is registered to, hover over the attacker IP address and select the **Who Is** link in the drop-down.



Because the session count from this IP address is high, check the **Blocked Content** and **Blocked Threats** widgets in the **Blocked Activity** tab for events related to this IP address. The **Blocked Activity** tab allows you to validate whether or not your policy rules are effective in blocking content or threats when a host on your network is compromised.

Use the **Export PDF** capability on the ACC to export the current view (create a snapshot of the data) and send it to an incidence response team. To view the threat logs directly from the widget, you can also click the icon to jump to the logs; the query is generated automatically and only the relevant logs are displayed onscreen (for example in **Monitor** > **Logs** > **Threat Logs**).

https://t.me/learningnets

You have now used the ACC to review network data/trends to find which applications or users are generating the most traffic, and how many application are responsible for the threats seen on the network. You were able to identify which application(s), user(s) generated the traffic, determine whether the application was on the default port, and which policy rule(s) allowed the traffic into the network, and determine whether the threat is spreading laterally on the network. You also identified the destination IP addresses, geo-locations with which hosts on the network are communicating with. Use the conclusions from your investigation to craft goal-oriented policies that can secure users and your network.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Use the App Scope Reports

The App Scope reports provide visibility and analysis tools to help pinpoint problematic behavior, helping you understand changes in application usage and user activity, users and applications that take up most of the network bandwidth, and identify network threats.

With the App Scope reports, you can quickly see if any behavior is unusual or unexpected. Each report provides a dynamic, user-customizable window into the network; hovering the mouse over and clicking either the lines or bars on the charts opens detailed information about the specific application, application category, user, or source on the ACC. The App Scope charts on **Monitor** > **App Scope** give you the ability to:

- Toggle the attributes in the legend to only view chart details that you want to review. The ability to include or exclude data from the chart allows you to change the scale and review details more closely.
- Click into an attribute in a bar chart and drill down to the related sessions in the ACC. Click into an Application name, Application Category, Threat Name, Threat Category, Source IP address or Destination IP address on any bar chart to filter on the attribute and view the related sessions in the ACC.
- Export a chart or map to PDF or as an image. For portability and offline viewing, you can Export charts and maps as PDFs or PNG images.

The following App Scope reports are available:

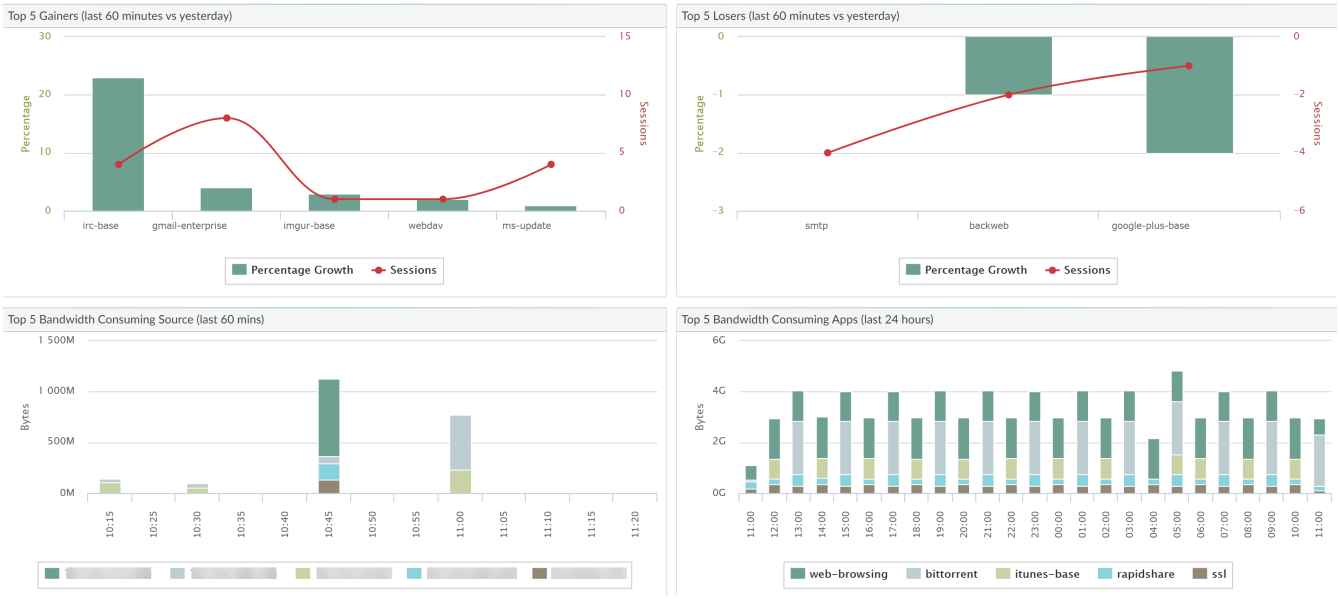
- Summary Report
- Change Monitor Report
- Threat Monitor Report
- Threat Map Report
- Network Monitor Report
- Traffic Map Report

Summary Report

The App Scope Summary report (**Monitor** > **App Scope** > **Summary**) displays charts for the top five gainers, losers, and bandwidth consuming applications, application categories, users, and sources.

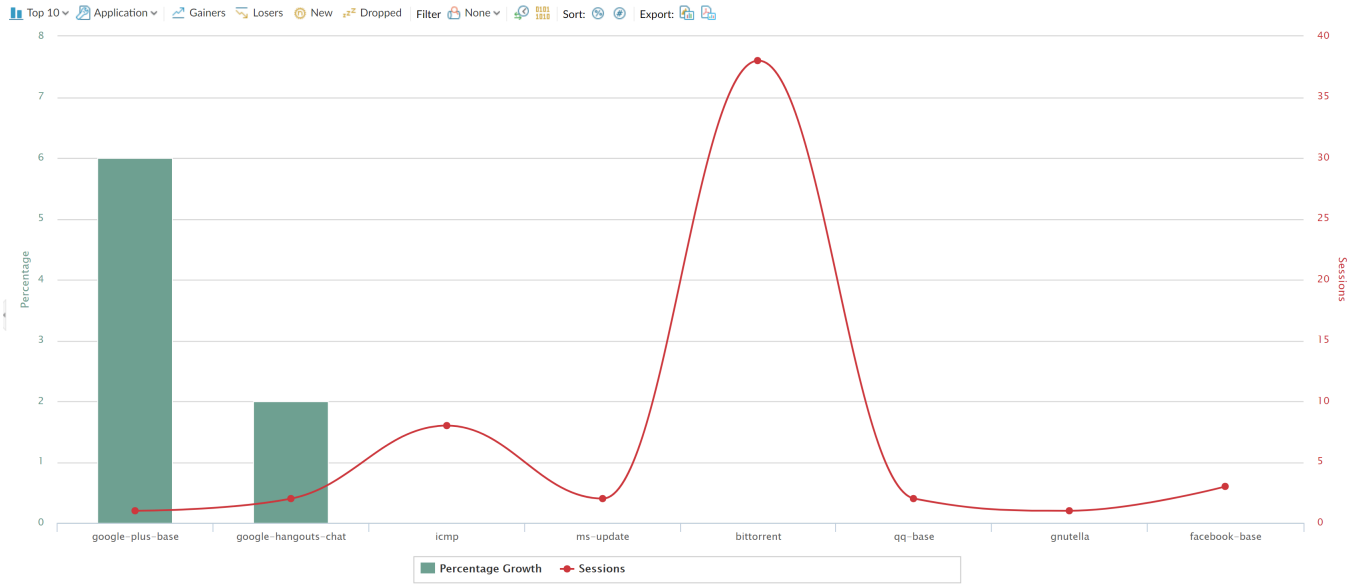
 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets



Change Monitor Report

The App Scope Change Monitor report (**Monitor** > **App Scope** > **Change Monitor**) displays changes over a specified time period. For example, the following chart displays the top applications that gained in use over the last hour as compared with the last 24-hour period. The top applications are determined by session count and sorted by percent.



The Change Monitor Report contains the following buttons and options.

Button	Description
Top 10	Determines the number of records with the highest measurement included in the chart.

https://t.me/learningnets

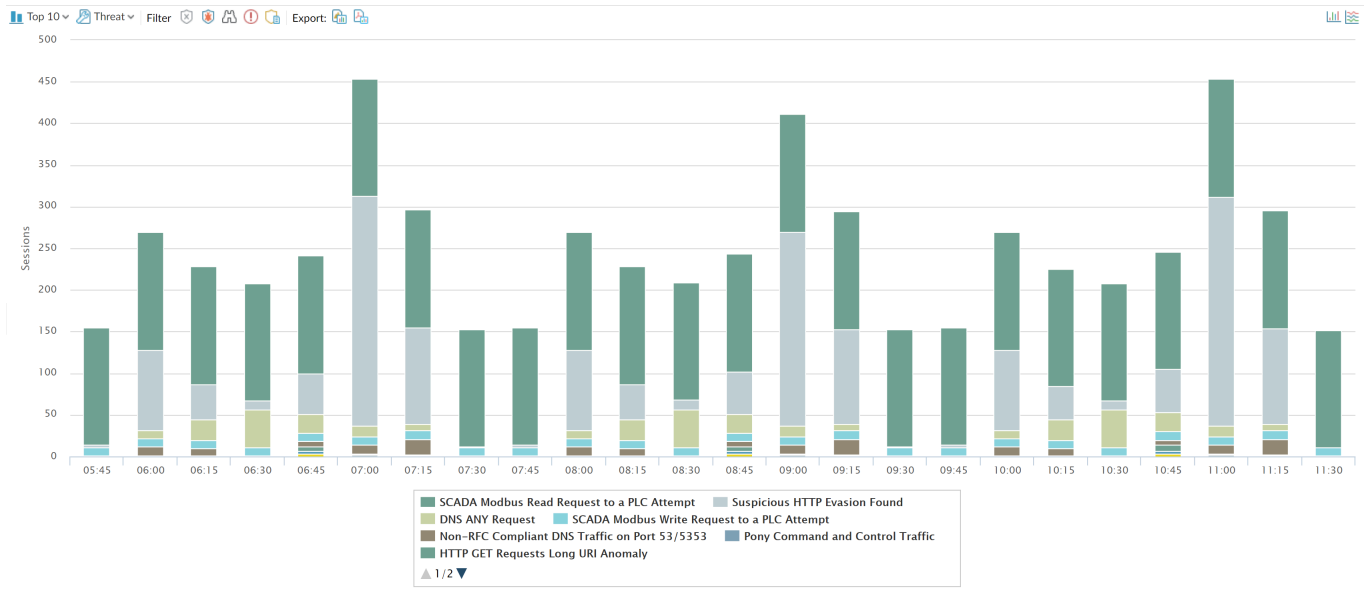
Button	Description
Application	Determines the type of item reported: Application, Application Category, Source, or Destination.
Gainers	Displays measurements of items that have increased over the measured period.
Losers	Displays measurements of items that have decreased over the measured period.
New	Displays measurements of items that were added over the measured period.
Dropped	Displays measurements of items that were discontinued over the measured period.
Filter	Applies a filter to display only the selected item. None displays all entries.
	Determines whether to display session or byte information.
Sort	Determines whether to sort entries by percentage or raw growth.
Export	Exports the graph as a .png image or as a PDF.
Compare	Specifies the period over which the change measurements are taken.

Threat Monitor Report

The App Scope Threat Monitor report (**Monitor** > **App Scope** > **Threat Monitor**) displays a count of the top threats over the selected time period. For example, the following figure shows the top 10 threat types over the last 6 hours.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets



Each threat type is color-coded as indicated in the legend below the chart. The Threat Monitor report contains the following buttons and options.

Button	Description
Top 10	Determines the number of records with the highest measurement included in the chart.
Threats	Determines the type of item measured: Threat, Threat Category, Source, or Destination.
Filter	Applies a filter to display only the selected type of items.
	Determines whether the information is presented in a stacked column chart or a stacked area chart.
Export	Exports the graph as a .png image or as a PDF.
	Specifies the period over which the measurements are taken.

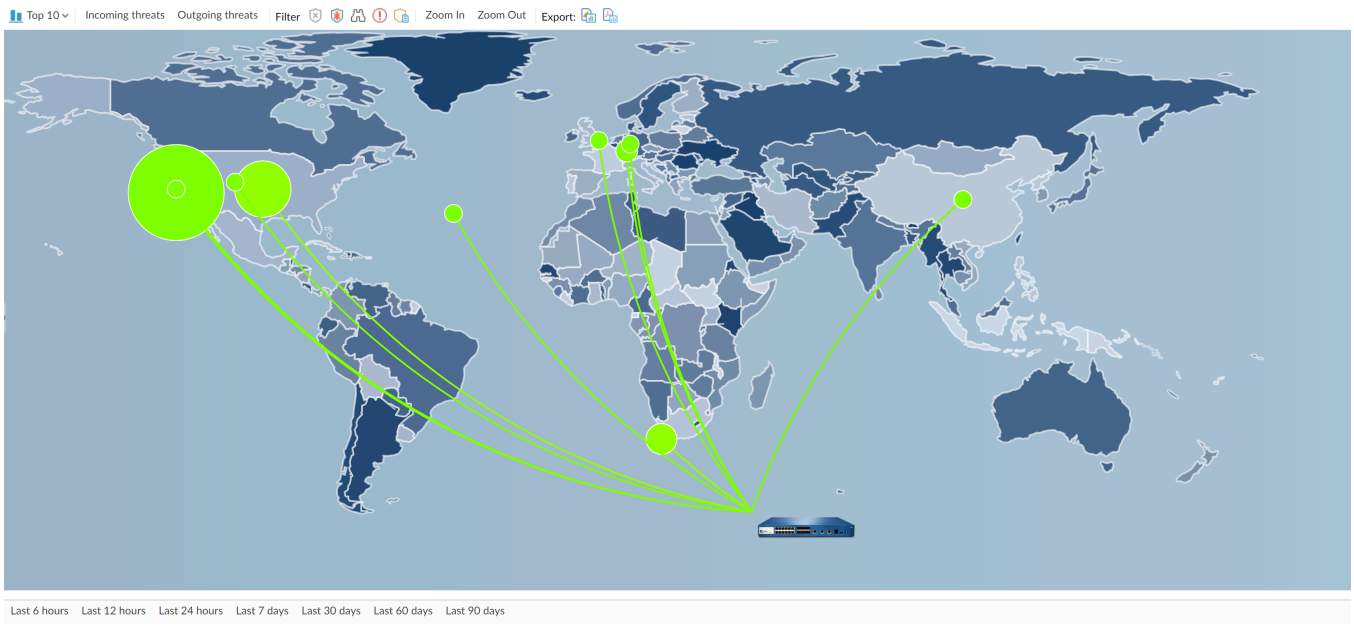
Threat Map Report

The App Scope Threat Map report (**Monitor** > **App Scope** > **Threat Map**) shows a geographical view of threats, including severity. Each threat type is color-coded as indicated in the legend below the chart.

The firewall uses geolocation for creating threat maps. The firewall is placed at the bottom of the threat map screen, if you have not specified the geolocation coordinates (**Device** > **Setup** > **Management**, General Settings section) on the firewall.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets



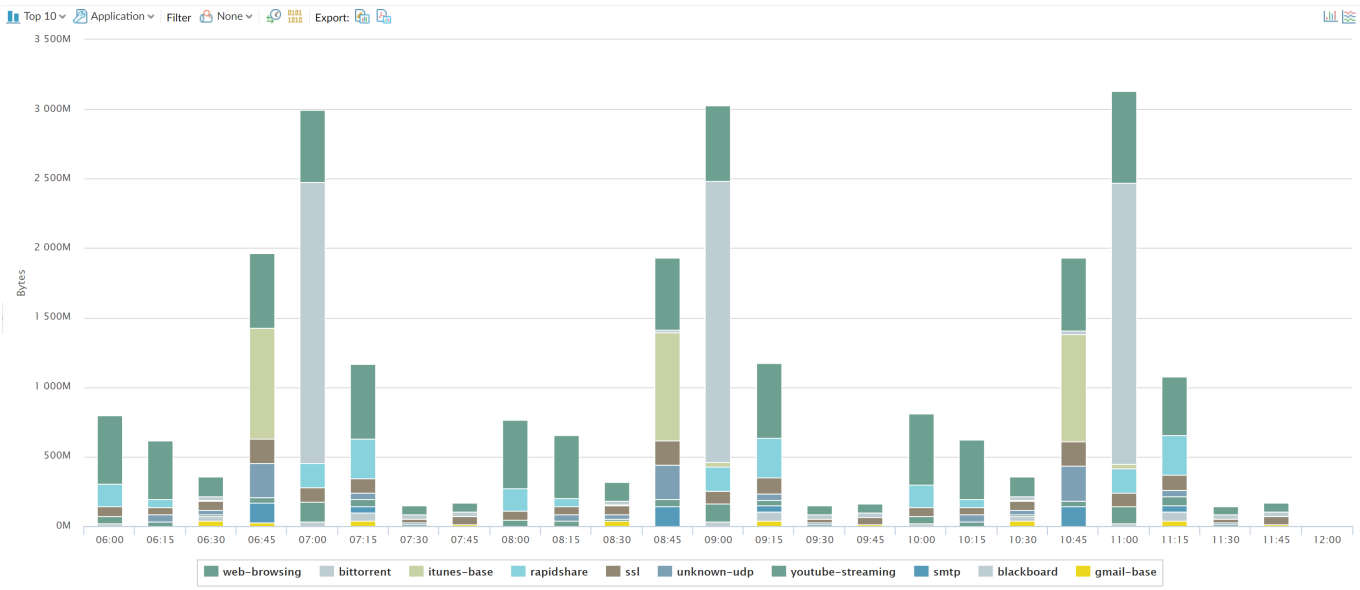
The Threat Map report contains the following buttons and options.

Button	Description
Top 10	Determines the number of records with the highest measurement included in the chart.
Incoming threats	Displays incoming threats.
Outdoing threats	Displays outgoing threats.
Filer	Applies a filter to display only the selected type of items.
Zoom In and Zoom Out	Zoom in and zoom out of the map.
Export	Exports the graph as a .png image or as a PDF.
Last 6 hours Last 12 hours Last 24 hours Last 7 days Last 30 days Last 60 days Last 90 da	Indicates the period over which the measurements are taken.

Network Monitor Report

The App Scope Network Monitor report (**Monitor** > **App Scope** > **Network Monitor**) displays the bandwidth dedicated to different network functions over the specified period of time. Each network function is color-coded as indicated in the legend below the chart. For example, the image below shows application bandwidth for the past 7 days based on session information.

https://t.me/learningnets



The Network Monitor report contains the following buttons and options.

Button	Description
Top 10	Determines the number of records with the highest measurement included in the chart.
Application	Determines the type of item reported: Application, Application Category, Source, or Destination.
Filter	Applies a filter to display only the selected item. **None** displays all entries.
	Determines whether to display session or byte information.
Export	Exports the graph as a .png image or as a PDF.
	Determines whether the information is presented in a stacked column chart or a stacked area chart.
Last 6 hours Last 12 hours Last 24 hours Last 7 days Last 30 days Last 60 days Last 90 da	Indicates the period over which the change measurements are taken.

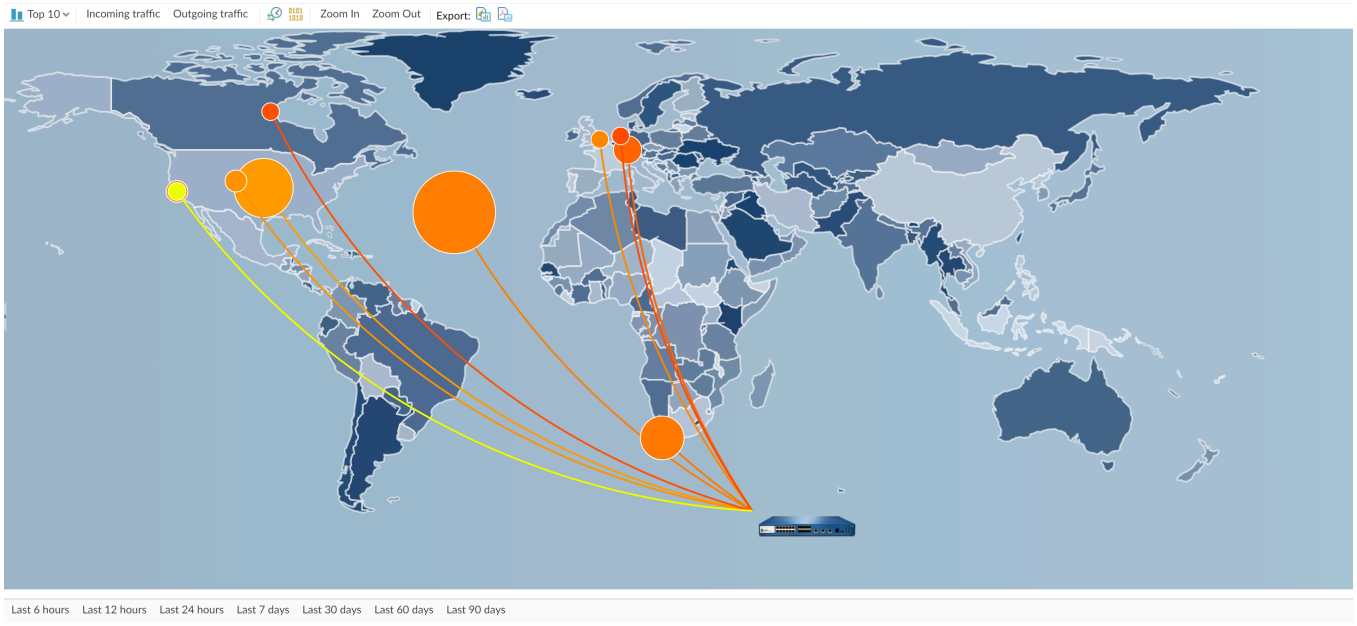
Traffic Map Report

The App Scope Traffic Map (**Monitor** > **App Scope** > **Traffic Map**) report shows a geographical view of traffic flows according to sessions or flows.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

The firewall uses geolocation for creating traffic maps. The firewall is placed at the bottom of the traffic map screen, if you have not specified the geolocation coordinates (**Device** > **Setup** > **Management,** General Settings section) on the firewall.



Each traffic type is color-coded as indicated in the legend below the chart. The Traffic Map report contains the following buttons and options.

Buttons	Description
Top 10	Determines the number of records with the highest measurement included in the chart.
Incoming threats	Displays incoming threats.
Outgoing threats	Displays outgoing threats.
	Determines whether to display session or byte information.
Zoom In and Zoom Out	Zoom in and zoom out of the map.
Export	Exports the graph as a .png image or as a PDF.
Last 6 hours Last 12 hours Last 24 hours Last 7 days Last 30 days Last 60 days Last 90 days	Indicates the period over which the change measurements are taken.

https://t.me/learningnets

Use the Automated Correlation Engine

The automated correlation engine is an analytics tool that uses the logs on the firewall to detect actionable events on your network. The engine correlates a series of related threat events that, when combined, indicate a likely compromised host on your network or some other higher level conclusion. It pinpoints areas of risk, such as compromised hosts on the network, allows you to assess the risk and take action to prevent exploitation of network resources. The automated correlation engine uses *correlation objects* to analyze the logs for patterns and when a match occurs, it generates a *correlated event*.

The following models support the automated correlation engine:

- *Panorama—M-Series appliances and virtual appliances*
- *PA-7000 Series firewalls*
- *PA-5450 firewall*
- *PA-5200 Series firewalls*
- *PA-3200 Series firewalls*

- Automated Correlation Engine Concepts
- View the Correlated Objects
- Interpret Correlated Events
- Use the Compromised Hosts Widget in the ACC

Automated Correlation Engine Concepts

The automated correlation engine uses *correlation objects* to analyze the logs for patterns and when a match occurs, it generates a *correlated event*.

- Correlation Object
- Correlated Events

Correlation Object

A correlation object is a definition file that specifies patterns to match against, the data sources to use for the lookups, and time period within which to look for these patterns. A pattern is a boolean structure of conditions that queries the following data sources (or logs) on the firewall: application statistics, traffic, traffic summary, threat summary, threat, data filtering, and URL filtering. Each pattern has a severity rating, and a threshold for the number of times the pattern match must occur within a defined time limit to indicate malicious activity. When the match conditions are met, a correlated event is logged.

A correlation object can connect isolated network events and look for patterns that indicate a more significant event. These objects identify suspicious traffic patterns and network anomalies, including suspicious IP activity, known command-and-control activity, known vulnerability exploits, or botnet activity that, when correlated, indicate with a high probability that a host on the network has been compromised. Correlation objects are defined and developed by the Palo Alto Networks Threat Research team, and are delivered with the weekly dynamic updates to

https://t.me/learningnets

the firewall and Panorama. To obtain new correlation objects, the firewall must have a Threat Prevention license. Panorama requires a support license to get the updates.

The patterns defined in a correlation object can be static or dynamic. Correlated objects that include patterns observed in WildFire are dynamic, and can correlate malware patterns detected by WildFire with command-and-control activity initiated by a host that was targeted with the malware on your network or activity seen by a Traps protected endpoint on Panorama. For example, when a host submits a file to the WildFire cloud and the verdict is malicious, the correlation object looks for other hosts or clients on the network that exhibit the same behavior seen in the cloud. If the malware sample had performed a DNS query and browsed to a malware domain, the correlation object will parse the logs for a similar event. When the activity on a host matches the analysis in the cloud, a high severity correlated event is logged.

Correlated Events

A correlated event is logged when the patterns and thresholds defined in a correlation object match the traffic patterns on your network. To Interpret Correlated Events and to view a graphical display of the events, see Use the Compromised Hosts Widget in the ACC.

View the Correlated Objects

You can view the correlation objects that are currently available on the firewall.

STEP 1 | Select **Monitor** > **Automated Correlation Engine** > **Correlation Objects**. All the objects in the list are enabled by default.

TITLE	CATEGORY	STATE	DESCRIPTION
Multiple User from One Endpoint MFA Credential Theft	credential-theft-abuse	active	This correlation object detects multiple account abuse from a possibly compromised endpoint
WildFire C2	compromised-host	active	This correlation object detects hosts that have exhibited command-and-control (C2) network behavior corresponding to malware detected by WildFire elsewhere on your network.
WildFire and Traps ESM Correlated C2	compromised-host	active	This correlation object detects hosts that have received malware detected by WildFire or executed malware as seen by Traps, and have also exhibited command- and-control (C2) network behavior corresponding to the detected malware.
Single Account and Endpoint MFA Credential Theft	credential-theft-abuse	active	This correlation object detects activity from a possibly compromised user account from a single endpoint
Compromise Activity Sequence	compromised-host	active	This correlation object detects a host involved in a sequence of activity indicating remote compromise, starting with scanning or probing activity, progressing to exploitation, and concluding with network contact to a known malicious domain.
Exploit Kit Activity	compromised-host	active	This object detects probable exploit kit activity targeted at a host on the network. Exploit kits are identified by a vulnerability exploit or exploit kit landing page signature, combined with either a malware download signature or a known command-and-control signature.
Single Account 1 FA Multiple Endpoints Credential Timeouts	credential-theft-abuse	active	This correlation object detects timed out attempts of first factor authentications from multiple endpoints using a single user account
Beacon Detection	compromised-host	active	This correlation object detects likely compromised hosts based on activity that resembles command-and-control (C2) beaconing, such as repeated visits to recently registered domains or dynamic DNS domains, repeated file downloads from the same location, generation of unknown traffic, etc.
Single Account and Endpoint MFA Credential Timeout	credential-theft-abuse	active	This correlation object detects timedout MFA authentication attempts from a single endpoint using single account
Multiple Endpoint MFA Credential Timeout Abuse	credential-theft-abuse	active	This correlation object detects timed out second factor authentications from multiple endpoints using a single user account
Multiple Endpoint MFA Credential Abuse	credential-theft-abuse	active	This correlation object detects activity from multiple endpoints using a single user account
Exploit Kit Delivering XOR obfuscated malware	compromised-host	active	This correlation object detects exclusive-or (XOR) obfuscated malware downloaded to a host. XOR obfuscation is a technique to evade detection by encrypting portions of a file in order to hide malicious code. This correlation object specifically identifies XOR obfuscated malware that is delivered to the host by an exploit kit. While the Exploit Kit Activity object detects exploit kits combined with either a malware download signature or a known command-and-control signature, this object is provided to specifically detect an event where XOR obfuscation malware inserted on a host by an exploit kit and to distinguish such an event from other exploit kit activities.
Single Account 1 FA Credential Abuse	credential-theft-abuse	active	This correlation object detects timed out first factor authentications from an endpoint using a single user account

STEP 2 | View the details on each correlation object. Each object provides the following information:

- **Name** and **Title**—The name and title indicate the type of activity that the correlation object detects. The name column is hidden from view, by default. To view the definition of the object, unhide the column and click the name link.
- **ID**— A unique number that identifies the correlation object; this column is also hidden by default. The IDs are in the 6000 series.
- **Category**—A classification of the kind of threat or harm posed to the network, user, or host. For now, all the objects identify compromised hosts on the network.
- **State**—Indicates whether the correlation object is enabled (active) or disabled (inactive). All the objects in the list are enabled by default, and are hence active. Because these objects

https://t.me/learningnets

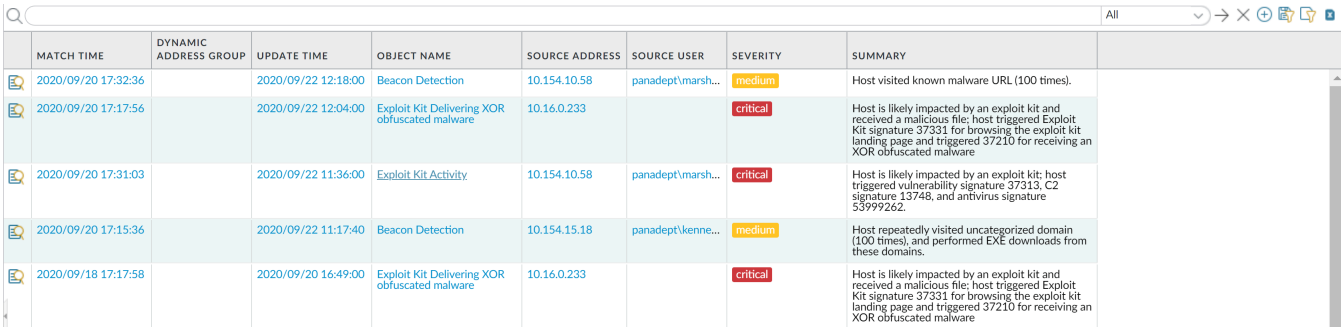
are based on threat intelligence data and are defined by the Palo Alto Networks Threat Research team, keep the objects active in order to track and detect malicious activity on your network.

- **Description**—Specifies the match conditions for which the firewall or Panorama will analyze logs. It describes the sequence of conditions that are matched on to identify acceleration or escalation of malicious activity or suspicious host behavior. For example, the **Compromise Lifecycle** object detects a host involved in a complete attack lifecycle in a three-step escalation that starts with scanning or probing activity, progressing to exploitation, and concluding with network contact to a known malicious domain.

For more information, see Automated Correlation Engine Concepts and Use the Automated Correlation Engine.

Interpret Correlated Events

You can view and analyze the logs generated for each correlated event in the **Monitor** > **Automated Correlation Engine** > **Correlated Events** tab.



Correlated Events includes the following details:

Field	Description
Match Time	The time the correlation object triggered a match.
Update Time	The time when the event was last updated with evidence on the match. As the firewall collects evidence on pattern or sequence of events defined in a correlation object, the time stamp on the correlated event log is updated.
Object Name	The name of the correlation object that triggered the match.
Source Address	The IP address of the user/device on your network from which the traffic originated.
Source User	The user and user group information from the directory server, if User-ID is enabled.
Severity	A rating that indicates the urgency and impact of the match. The severity level indicates the extent of damage or escalation pattern, and the frequency of occurrence. Because correlation objects are

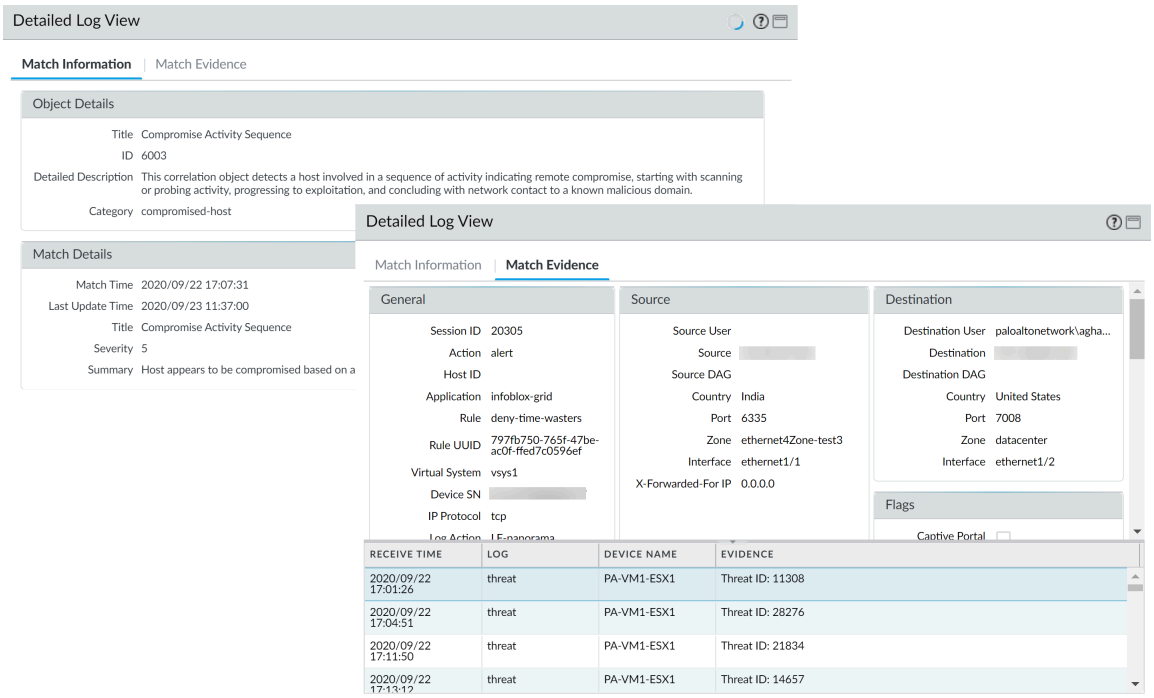
https://t.me/learningnets

Field	Description
To configure the firewall or Panorama to send alerts using email, SNMP or syslog messages for a desired severity level, see Use External Services for Monitoring.	primarily for detecting threats, the correlated events typically relate to identifying compromised hosts on the network and the severity implies the following: • **Critical**—Confirms that a host has been compromised based on correlated events that indicate an escalation pattern. For example, a critical event is logged when a host that received a file with a malicious verdict by WildFire exhibits the same command-and-control activity that was observed in the WildFire sandbox for that malicious file. • **High**—Indicates that a host is very likely compromised based on a correlation between multiple threat events, such as malware detected anywhere on the network that matches the command-and-control activity generated by a particular host. • **Medium**—Indicates that a host is likely compromised based on the detection of one or multiple suspicious events, such as repeated visits to known malicious URLs, which suggests a scripted command-and-control activity. • **Low**—Indicates that a host is possibly compromised based on the detection of one or multiple suspicious events, such as a visit to a malicious URL or a dynamic DNS domain. • **Informational**—Detects an event that may be useful in aggregate for identifying suspicious activity, but the event is not necessarily significant on its own.
Summary	A description that summarizes the evidence gathered on the correlated event.

Click the icon to see the detailed log view, which includes all the evidence on a match:

 ©2022 Palo Alto Networks, Inc.

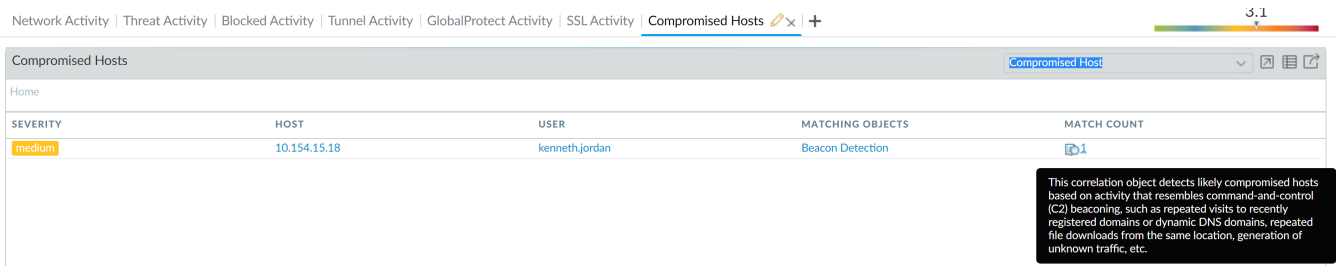
https://t.me/learningnets



Tab	Description
Match Information	Object Details: Presents information on the Correlation Object that triggered the match.
	Match Details: A summary of the match details that includes the match time, last update time on the match evidence, severity of the event, and an event summary.
Match Evidence	Presents all the evidence that corroborates the correlated event. It lists detailed information on the evidence collected for each session.

Use the Compromised Hosts Widget in the ACC

The compromised hosts widget on **ACC** > **Threat Activity**, aggregates the Correlated Events and sorts them by severity. It displays the source IP address/user who triggered the event, the correlation object that was matched and the number of times the object was matched. Use the match count link to jump to the match evidence details.



For more details, see Use the Automated Correlation Engine and Use the Application Command Center.

https://t.me/learningnets

Take Packet Captures

All Palo Alto Networks firewalls allow you to take packet captures (pcaps) of traffic that traverses the management interface and network interfaces on the firewall. When taking packet captures on the dataplane, you may need to Disable Hardware Offload to ensure that the firewall captures all traffic.

Packet capture can be very CPU intensive and can degrade firewall performance. Only use this feature when necessary and make sure you turn it off after you have collected the required packets.

- Types of Packet Captures
- Disable Hardware Offload
- Take a Custom Packet Capture
- Take a Threat Packet Capture
- Take an Application Packet Capture
- Take a Packet Capture on the Management Interface

Types of Packet Captures

There are different types of packet captures you can enable, depending on what you need to do:

- **Custom Packet Capture**—The firewall captures packets for all traffic or for specific traffic based on filters that you define. For example, you can configure the firewall to only capture packets to and from a specific source and destination IP address or port. You then use the packet captures for troubleshooting network#related issues or for gathering application attributes to enable you to write custom application signatures or to request an application signature from Palo Alto Networks. See Take a Custom Packet Capture.
- **Threat Packet Capture**—The firewall captures packets when it detects a virus, spyware, or vulnerability. You enable this feature in Antivirus, Anti-Spyware, and Vulnerability Protection security profiles. A link to view or export the packet captures will appear in the second column of the Threat log. These packet captures provide context around a threat to help you determine if an attack is successful or to learn more about the methods used by an attacker. You can also submit this type of pcap to Palo Alto Networks to have a threat re-analyzed if you feel it's a false-positive or false-negative. See Take a Threat Packet Capture.
- **Application Packet Capture**—The firewall captures packets based on a specific application and filters that you define. A link to view or export the packet captures will appear in the second column of the Traffic logs for traffic that matches the packet capture rule. See Take an Application Packet Capture.
- **Management Interface Packet Capture**—The firewall captures packets on the management interface (MGT). The packet captures are useful when troubleshooting services that traverse the interface, such as firewall management authentication to External Authentication Services, software and content updates, log forwarding, communication with SNMP servers, and authentication requests for GlobalProtect and Authentication Portal. See Take a Packet Capture on the Management Interface.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

- **GTP Event Packet Capture**—The firewall captures a single GTP event, such as GTP-in-GTP, end user IP spoofing, and abnormal GTP messages, to make GTP troubleshooting easier for mobile network operators. Enable packet capture in a Mobile Network Protection profile.

Disable Hardware Offload

Packet captures for traffic passing through the network data ports on a Palo Alto Networks firewall are performed by the dataplane CPU. To capture traffic that passes through the management interface, you must Take a Packet Capture on the Management Interface, in which case the packet capture is performed on the management plane.

When a packet capture is performed on the dataplane, the packet capture filter is used differently by the ingress stage, compared to the firewall, drop, and egress capture stages. The ingress stage uses the packet capture filter to copy individual packets that match the filter to the capture file. Packets that fail packet-parsing checks are dropped before being captured. The firewall, drop, and egress capture stages use the same packet capture filter to mark all new sessions that match the filter. Because each session, as recorded in the session tables, identifies both client-to-server and server-to-client connections, any traffic, in either direction, that matches to the flagged session will be copied to the firewall-stage and transmit-stage capture files. Likewise, any dropped traffic (post receive stage) in either direction that matches to a flagged session will be copied to the drop-stage capture file.

On firewall models that include a network processor, traffic that meets certain pre-determined criteria by Palo Alto Networks may be offloaded for handling by the network processor. Such offloaded traffic will not reach the dataplane CPU and will, therefore, not be captured. To capture offloaded traffic, you must use the CLI to turn off the hardware offload feature.

Common types of traffic that may be offloaded include non-decrypted SSL and SSH traffic (which being encrypted cannot be usefully inspected beyond the initial SSL/SSH session setup), network protocols (such as OSPF, BGP, RIP), and traffic that matches an application-override policy. Some types of traffic will never be offloaded, such as ARP, all non-IP traffic, IPSec, and VPN sessions. Individual SYN, FIN, and RST packets, even for session traffic that has been offloaded, will never be offloaded, and will always be passed through to the dataplane CPU, once recognized as such by the network processor.

Hardware offload is supported on the following firewalls: PA-3200 Series, PA-5200 Series, and PA-7000 Series firewall.

Disabling hardware offload may increase the dataplane CPU usage. If dataplane CPU usage is already high, you may want to schedule a maintenance window before disabling hardware offload.

STEP 1 | Disable hardware offload by running the following CLI command:

```
admin@PA-7050>set session offload no
```

STEP 2 | After the firewall captures the required traffic, enable hardware offload by running the following CLI command:

```
admin@PA-7050>set session offload yes
```

 ©2022 Palo Alto Networks, Inc.

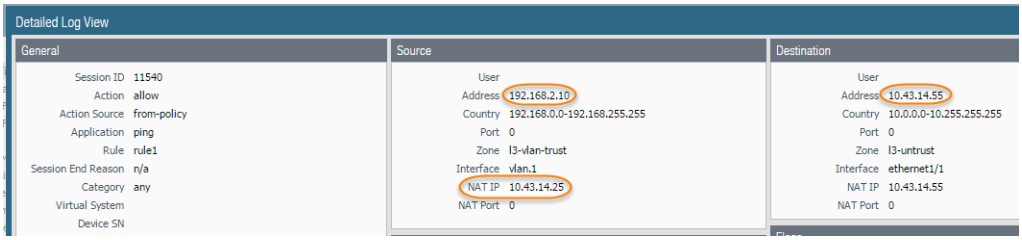
https://t.me/learningnets

Take a Custom Packet Capture

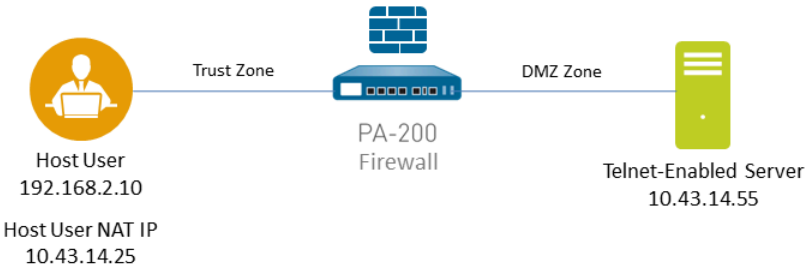
Custom packet captures allow you to define the traffic that the firewall will capture. To ensure that you capture all traffic, you may need to Disable Hardware Offload.

STEP 1 | Before you start a packet capture, identify the attributes of the traffic that you want to capture.

For example, to determine the source IP address, source NAT IP address, and the destination IP address for traffic between two systems, perform a ping from the source system to the to the destination system. After the ping is complete, go to **Monitor** > **Traffic** and locate the traffic log for the two systems. Click the **Detailed Log View** icon located in the first column of the log and note the source address, source NAT IP, and the destination address.



The following example shows how to use a packet capture to troubleshoot a Telnet connectivity issue from a user in the Trust zone to a server in the DMZ zone.



https://t.me/learningnets

STEP 2 | Set packet capture filters, so the firewall only captures traffic you are interested in.

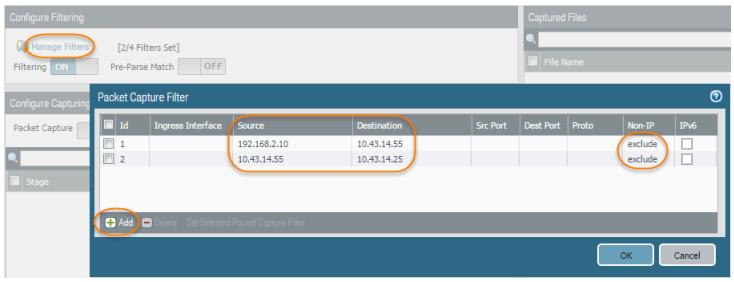
Using filters makes it easier for you to locate the information you need in the packet capture and will reduce the processing power required by the firewall to take the packet capture. To capture all traffic, do not define filters and leave the filter option off.

For example, if you configured NAT on the firewall, you will need to apply two filters. The first one filters on the pre-NAT source IP address to the destination IP address and the second one filters traffic from the destination server to the source NAT IP address.

1. Select **Monitor** > **Packet Capture**.
2. Click **Clear All Settings** at the bottom of the window to clear any existing capture settings.
3. Click **Manage Filters** and click **Add**.
4. Select **Id** 1 and in the **Source** field enter the source IP address you are interested in and in the **Destination** field enter a destination IP address.

   For example, enter the source IP address **`192.168.2.10`** and the destination IP address **`10.43.14.55`**. To further filter the capture, set **Non-IP** to **exclude** non-IP traffic, such as broadcast traffic.
5. **Add** the second filter and select **Id** 2.

   For example, in the **Source** field enter **`10.43.14.55`** and in the **Destination** field enter **`10.43.14.25`**. In the **Non-IP** drop-down menu select **exclude**.



6. Click **OK**.

STEP 3 | Set **Filtering** to **On**.

 ©2022 Palo Alto Networks, Inc.

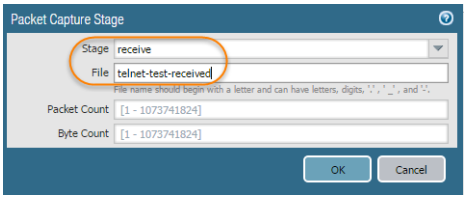
https://t.me/learningnets

STEP 4 | Specify the traffic stage(s) that trigger the packet capture and the filename(s) to use to store the captured content. For a definition of each stage, click the **Help** icon on the packet capture page.

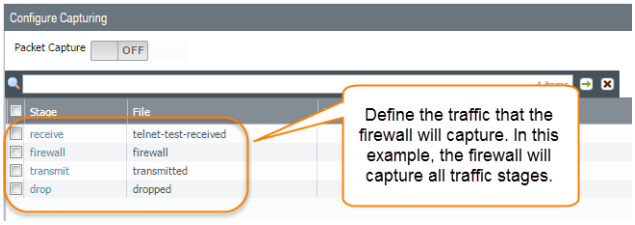
For example, to configure all packet capture stages and define a filename for each stage, perform the following procedure:

1. **Add** a **Stage** to the packet capture configuration and define a **File** name for the resulting packet capture.

   For example, select **receive** as the **Stage** and set the **File** name to telnet-test-received.

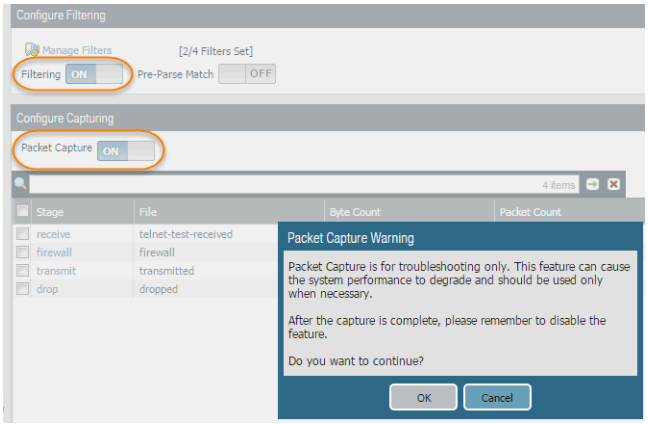


2. Continue to **Add** each **Stage** you want to capture (**receive, firewall**, **transmit**, and **drop**) and set a unique **File** name for each stage.



STEP 5 | Set **Packet Capture** to **ON**.

The firewall or appliance warns you that system performance can be degraded; acknowledge the warning by clicking **OK**. If you define filters, the packet capture should have little impact on performance, but you should always turn **Off** packet capture after the firewall captures the data that you want to analyze.



STEP 6 | Generate traffic that matches the filters that you defined.

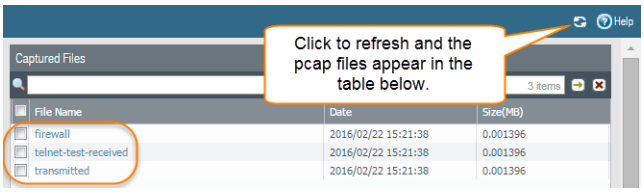
For this example, generate traffic from the source system to the Telnet-enabled server by running the following command from the source system (192.168.2.10):

```
telnet 10.43.14.55
```

 ©2022 Palo Alto Networks, Inc.

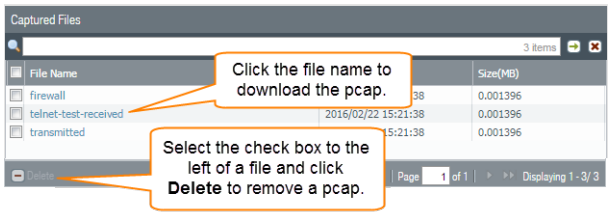
https://t.me/learningnets

STEP 7 | Turn packet capture **OFF** and then click the refresh icon to see the packet capture files.



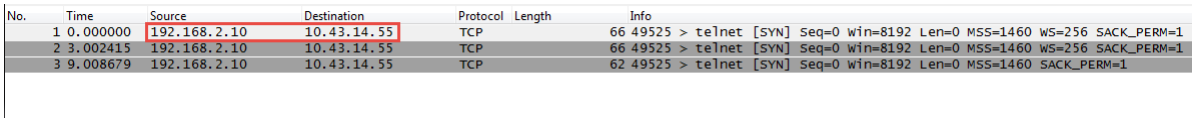
Notice that in this case, there were no dropped packets, so the firewall did not create a file for the drop stage.

STEP 8 | Download the packet captures by clicking the filename in the File Name column.



STEP 9 | View the packet capture files using a network packet analyzer.

In this example, the received.pcap packet capture shows a failed Telnet session from the source system at 192.168.2.10 to the Telnet-enabled server at 10.43.14.55. The source system sent the Telnet request to the server, but the server did not respond. In this example, the server may not have Telnet enabled, so check the server.



STEP 10 | Enable the Telnet service on the destination server (10.43.14.55) and turn on packet capture to take a new packet capture.

STEP 11 | Generate traffic that will trigger the packet capture.

Run the Telnet session again from the source system to the Telnet-enabled server

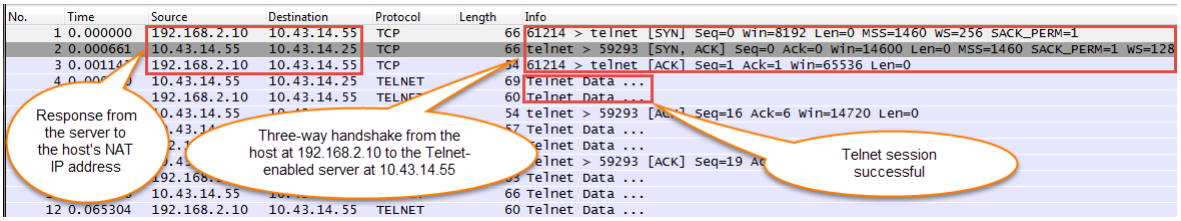
```
telnet 10.43.14.55
```


https://t.me/learningnets

STEP 12 | Download and open the received.pcap file and view it using a network packet analyzer.

The following packet capture now shows a successful Telnet session from the host user at 192.168.2.10 to the Telnet-enabled server at 10.43.14.55.

You also see the NAT address 10.43.14.25. When the server responds, it does so to the NAT address. You can see the session is successful as indicated by the three-way handshake between the host and the server and then you see Telnet data.



Take a Threat Packet Capture

To configure the firewall to take a packet capture (pcap) when it detects a threat, enable packet capture on Antivirus, Anti-Spyware, and Vulnerability Protection security profiles.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 1 | Enable the packet capture option in the security profile.

Some security profiles allow you to define a single-packet capture or an extended-capture. If you choose extended-capture, define the capture length. This will allow the firewall to capture more packets to provide additional context related to the threat.

If the action for a given threat is allow, the firewall does not trigger a Threat log and does not capture packets. If the action is alert, you can set the packet capture to single-packet or extended-capture. All blocking actions (drop, block, and reset actions) capture a single packet. The content package on the device determines the default action.

1. Select **Objects** > **Security Profiles** and enable the packet capture option for the supported profiles as follows:
   - **Antivirus**—Select a custom antivirus profile and in the **Antivirus** tab select the **Packet Capture** check box.
   - **Anti-Spyware**—Select a custom Anti-Spyware profile, click the **DNS Signatures** tab and in the **Packet Capture** drop-down, select **single-packet** or **extended-capture**.
   - **Vulnerability Protection**—Select a custom Vulnerability Protection profile and in the **Rules** tab, click **Add** to add a new rule, or select an existing rule. Set **Packet Capture** to **single-packet** or **extended-capture**.

   *If the profile has signature exceptions defined, click the **Exceptions** tab and in the **Packet Capture** column for a signature, set **single-packet** or **extended-capture**.*
2. (Optional) If you selected **extended-capture** for any of the profiles, define the extended packet capture length.
   1. Select **Device** > **Setup** > **Content-ID** and edit the Content-ID Settings.
   2. In the **Extended Packet Capture Length (packets)** section, specify the number of packets that the firewall will capture (range is 1-50; default is 5).
   3. Click **OK**.

STEP 2 | Add the security profile (with packet capture enabled) to a Security Policy rule.

1. Select **Policies** > **Security** and select a rule.
2. Select the **Actions** tab.
3. In the Profile Settings section, select a profile that has packet capture enabled.

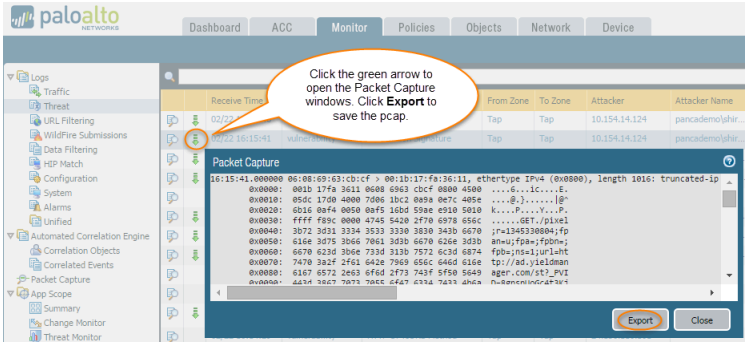
   For example, click the **Antivirus** drop-down and select a profile that has packet capture enabled.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 3 | View/export the packet capture from the Threat logs.

1. Select **Monitor** > **Logs** > **Threat**.
2. In the log entry that you are interested in, click the green packet capture icon in the second column. View the packet capture directly or **Export** it to your system.



Take an Application Packet Capture

The following topics describe two ways that you can configure the firewall to take application packet captures:

- Take a Packet Capture for Unknown Applications
- Take a Custom Application Packet Capture

Take a Packet Capture for Unknown Applications

Palo Alto Networks firewalls automatically generate a packet capture for sessions that contain an application that the firewall cannot identify. Typically, the only applications that are classified as unknown traffic—tcp, udp, or non-syn-tcp—are commercially available applications that do not yet have App-ID signatures, are internal or custom applications on your network, or potential threats. You can use these packet captures to gather more context related to the unknown application or use the information to analyze the traffic for potential threats. You can also Manage Custom or Unknown Applications by controlling them through security policy or by writing a custom application signature and then creating a security rule based on the custom signature. If the application is a commercial application, you can submit the packet capture to Palo Alto Networks to have an App-ID signature created.

STEP 1 | Verify that unknown application packet capture is enabled (this option is enabled by default).

1. To view the unknown application capture setting, run the following CLI command:

```
admin@PA-220>show running application setting | match “Unknown
 capture”
```

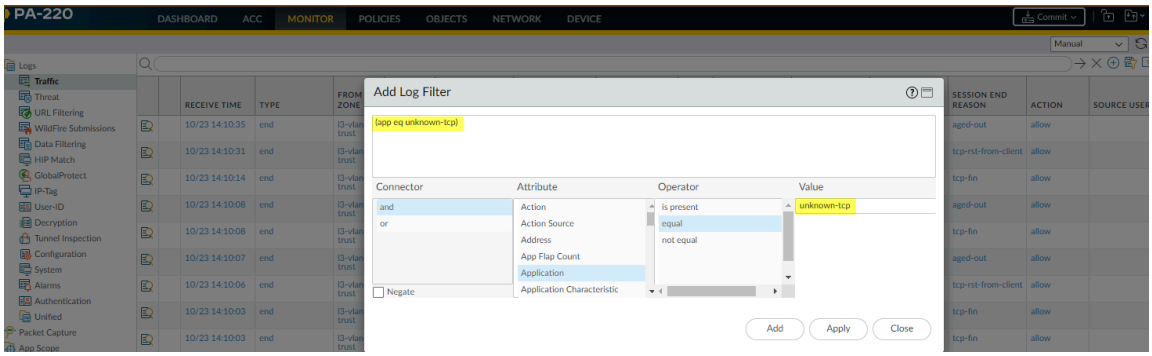
2. If the unknown capture setting option is off, enable it:

```
admin@PA-220>set application dump-unknown yes
```

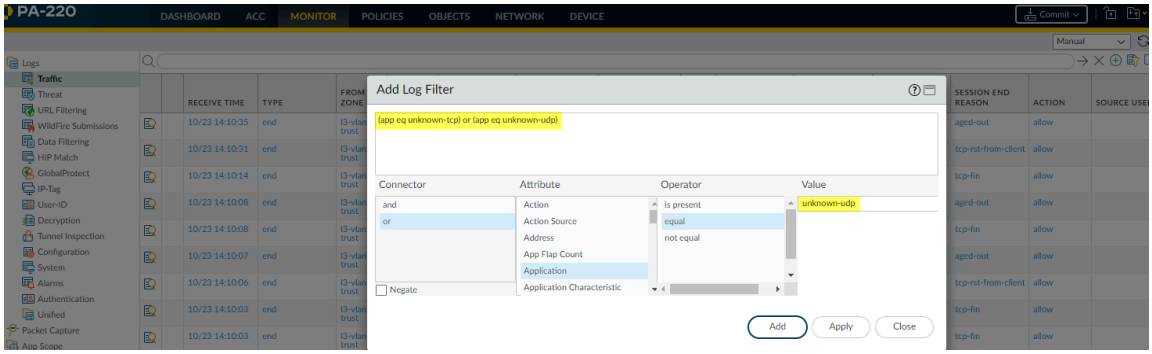

https://t.me/learningnets

STEP 2 | Locate unknown TCP and UDP applications by filtering the traffic logs.

1. Select **Monitor** > **Logs** > **Traffic**.
2. Click **Add Filter**, create the unknown TCP portion of the filter (**Connector** = “and”, **Attribute** = “Application”, **Operator** = “equal”, and enter “unknown-tcp” as the **Value**), and then click **Add** to add the query to the filter.

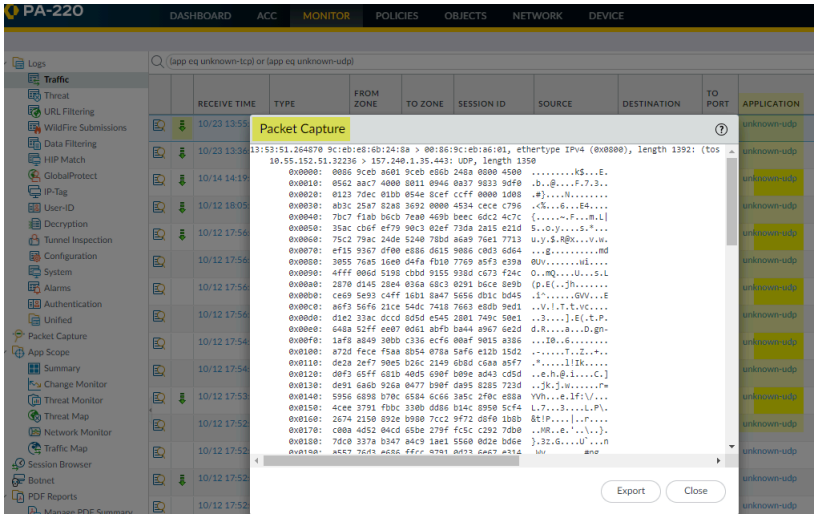


3. Create the unknown UDP portion of the filter (**Connector** = “or”, **Attribute** = “Application”, **Operator** = “equal”, and enter “unknown-udp” as the **Value**), and then click **Add** to add the query to the filter.



4. Click **Apply** to place the filter in the log screen query field.

STEP 3 | Click the **Apply Filter** arrow next to the query field to run the filter and then click the packet capture icon to view the packet capture or **Export** it to your local system.



https://t.me/learningnets

Take a Custom Application Packet Capture

You can configure a Palo Alto Networks firewall to take a packet capture based on an application name and filters that you define. You can then use the packet capture to troubleshoot issues with controlling an application. When configuring an application packet capture, you must use the application name defined in the App-ID database. You can view a list of all App-ID applications using Applipedia or from the web interface on the firewall in **Objects** > **Applications**.

STEP 1 | Using a terminal emulation application, such as PuTTY, launch an SSH session to the firewall.

STEP 2 | Turn on the application packet capture and define filters.

```
admin@PA-220>set application dump on application <application-name>
 rule <rule-name>
```

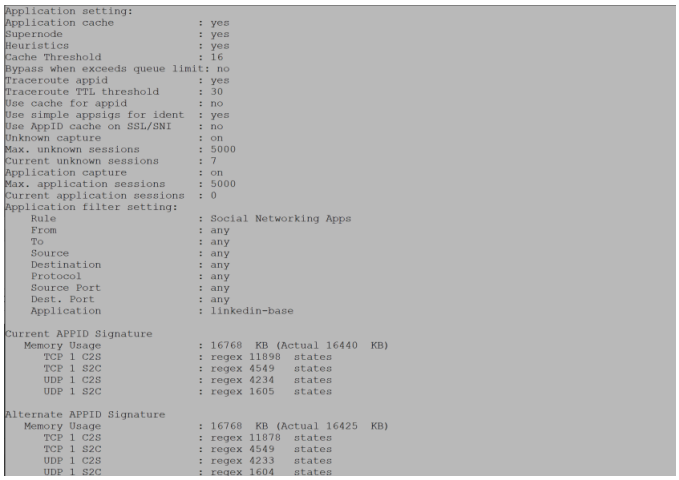
For example, to capture packets for the linkedin-base application that matches the security rule named Social Networking Apps, run the following CLI command:

```
admin@PA-220>set application dump on application linkedin-base rule
 "Social Networking Apps"
```

You can also apply other filters, such as source IP address and destination IP address.

STEP 3 | View the packet capture output to ensure that the correct filters are applied. The output displays after you enable the packet capture.

The following output confirms that application capture filtering is now based on the linkedin-base application for traffic that matches the Social Networking Apps rule.



```
Application setting:
Application cache              : yes
Supernode                      : yes
Heuristics                     : yes
Cache Threshold                : 16
Bypass when exceeds queue limit: no
Traceroute appid               : yes
Traceroute TTL threshold       : 30
Use cache for appid            : no
Use simple appsigs for ident   : yes
Use AppID cache on SSL/SNI     : no
Unknown capture                : on
Max. unknown sessions          : 5000
Current unknown sessions       : 7
Application capture            : on
Max. application sessions      : 5000
Current application sessions   : 0
Application filter setting:
    Rule                       : Social Networking Apps
    From                       : any
    To                         : any
    Source                     : any
    Destination                : any
    Protocol                   : any
    Source Port                : any
    Dest. Port                 : any
    Application                : linkedin-base

Current APPID Signature
   Memory Usage                : 16768  KB (Actual 16440  KB)
      TCP 1 C2S                : regex 11898  states
      TCP 1 S2C                : regex 4549   states
      UDP 1 C2S                : regex 4234   states
      UDP 1 S2C                : regex 1605   states

Alternate APPID Signature
   Memory Usage                : 16768  KB (Actual 16425  KB)
      TCP 1 C2S                : regex 11878  states
      TCP 1 S2C                : regex 4549   states
      UDP 1 C2S                : regex 4233   states
      UDP 1 S2C                : regex 1604   states
```

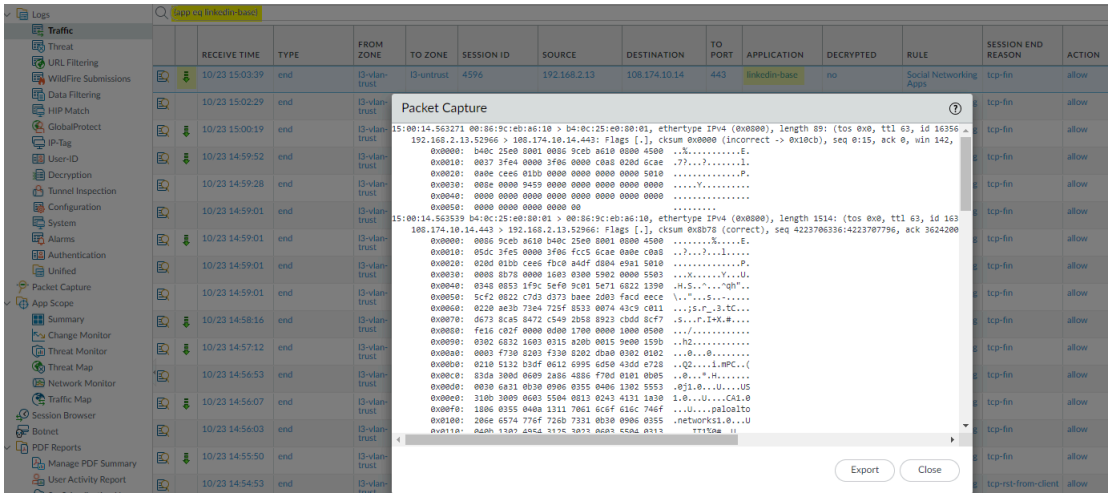
STEP 4 | Access linkedin.com from a web browser and perform some LinkedIn tasks to generate LinkedIn traffic, and then run the following CLI command to turn off application packet capture:

```
admin@PA-220>set application dump off
```


https://t.me/learningnets

STEP 5 | View/export the packet capture.

1. Log in to the web interface on the firewall and select **Monitor** > **Logs** > **Traffic**.
2. In the log entry that you are interested in, click the green packet capture icon .
3. View the packet capture directly or **Export** it to your computer. The following screen capture shows the linkedin-base packet capture.



Take a Packet Capture on the Management Interface

The **tcpdump** CLI command enables you to capture packets that traverse the management interface (MGT) on a Palo Alto Networks firewall.

*Each platform has a default number of bytes that **tcpdump** captures. The PA-220 firewalls capture 68 bytes of data from each packet and anything over that is truncated. The PA-7000 Series firewalls and VM-Series firewalls capture 96 bytes of data from each packet. To define the number of packets that **tcpdump** will capture, use the **snaplen** (snap length) option (range 0-65535). Setting the **snaplen** to 0 will cause the firewall to use the maximum length required to capture whole packets.*

STEP 1 | Using a terminal emulation application, such as PuTTY, launch an SSH session to the firewall.

https://t.me/learningnets

STEP 2 | To start a packet capture on the MGT interface, run the following command:

```
admin@PA-220>tcpdump filter “<filter-option> <IP-address>” snaplen
 length
```

For example, to capture the traffic that is generated when and administrator authenticates to the firewall using RADIUS, filter on the destination IP address of the RADIUS server (10.5.104.99 in this example):

```
admin@PA-220>tcpdump filter “dst 10.5.104.99” snaplen 0
```

You can also filter on src (source IP address), host, net, and you can exclude content. For example, to filter on a subnet and exclude all SCP, SFTP, and SSH traffic (which uses port 22), run the following command:

```
admin@PA-220>tcpdump filter “net 10.5.104.0/24 and not port 22”
 snaplen 0
```

*Each time **tcpdump** takes a packet capture, it stores the content in a file named mgmt.pcap. This file is overwritten each time you run **tcpdump**.*

STEP 3 | After the traffic you are interested in has traversed the MGT interface, press Ctrl + C to stop the capture.

STEP 4 | View the packet capture by running the following command:

```
admin@PA-220> view-pcap mgmt-pcap mgmt.pcap
```

The following output shows the packet capture from the MGT port (10.5.104.98) to the RADIUS server (10.5.104.99):

```
09:55:29.139394 IP 10.5.104.98.43063 > 10.5.104.99.radius: RADIUS,
 Access Request (1), id: 0x00 length: 89
 09:55:29.144354 arp reply 10.5.104.98 is-at 00:25:90:23:94:98 (oui
 Unknown)
 09:55:29.379290 IP 10.5.104.98.43063 > 10.5.104.99.radius: RADIUS,
 Access Request (1), id: 0x00 length: 70
 09:55:34.379262 arp who-has 10.5.104.99 tell 10.5.104.98
```

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 5 | (Optional) Export the packet capture from the firewall using SCP (or TFTP). For example, to export the packet capture using SCP, run the following command:

```
admin@PA-220>scp export mgmt-pcap from mgmt.pcap
 to <username@host:path>
```

For example, to export the pcap to an SCP enabled server at 10.5.5.20 to a temp folder named temp-SCP, run the following CLI command:

```
admin@PA-220>scp export mgmt-pcap from mgmt.pcap to
 admin@10.5.5.20:c:/temp-SCP
```

Enter the login name and password for the account on the SCP server to enable the firewall to copy the packet capture to the c:\temp-SCP folder on the SCP-enabled.

STEP 6 | You can now view the packet capture files using a network packet analyzer, such as Wireshark.

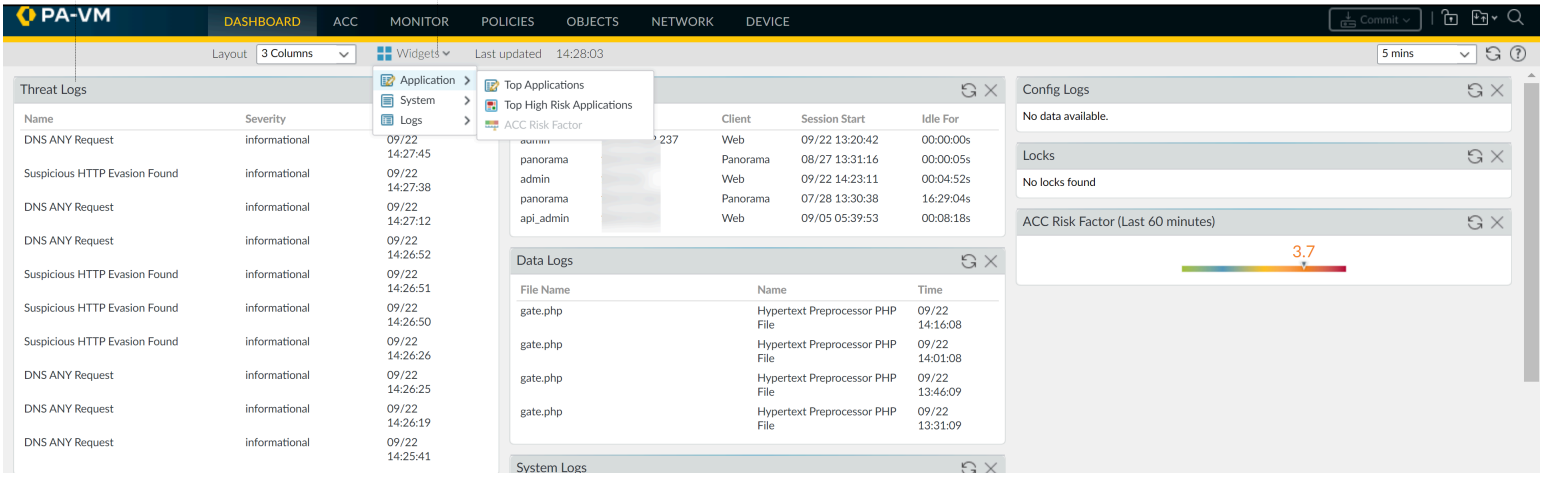
 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Monitor Applications and Threats

All Palo Alto Networks next-generation firewalls come equipped with the App-ID technology, which identifies the applications traversing your network, irrespective of protocol, encryption, or evasive tactic. You can then Use the Application Command Center to monitor the applications. The ACC graphically summarizes the data from a variety of log databases to highlight the applications traversing your network, who is using them, and their potential security impact. ACC is dynamically updated, using the continuous traffic classification that App-ID performs; if an application changes ports or behavior, App-ID continues to see the traffic, displaying the results in ACC. Additional visibility into URL categories, threats, and data provides a complete and well-rounded picture of network activity. With ACC, you can very quickly learn more about the traffic traversing the network and then translate that information into a more informed security policy

You can also Use the Dashboard to monitor the network.



Review the Content Delivery Network Infrastructure to check whether logged events on the firewall pose a security risk. The AutoFocus intelligence summary shows the prevalence of properties, activities, or behaviors associated with logs in your network and on a global scale, as well as the WildFire verdict and AutoFocus tags linked to them. With an active AutoFocus subscription, you can use this information to create customized AutoFocus Alerts that track specific threats on your network.

https://t.me/learningnets

View and Manage Logs

A log is an automatically generated, time-stamped file that provides an audit trail for system events on the firewall or network traffic events that the firewall monitors. Log entries contain *artifacts*, which are properties, activities, or behaviors associated with the logged event, such as the application type or the IP address of an attacker. Each log type records information for a separate event type. For example, the firewall generates a Threat log to record traffic that matches a spyware, vulnerability, or virus signature or a DoS attack that matches the thresholds configured for a port scan or host sweep activity on the firewall.

- Log Types and Severity Levels
- View Logs
- Filter Logs
- Export Logs
- Configure Log Storage Quotas and Expiration Periods
- Schedule Log Exports to an SCP or FTP Server

Log Types and Severity Levels

You can see the following log types in the **Monitor** > **Logs** pages.

- Traffic Logs
- Threat Logs
- URL Filtering Logs
- WildFire Submissions Logs
- Data Filtering Logs
- Correlation Logs
- Tunnel Inspection Logs
- Config Logs
- System Logs
- HIP Match Logs
- GlobalProtect Logs
- IP-Tag Logs
- User-ID Logs
- Decryption Logs
- Alarms Logs
- Authentication Logs
- Unified Logs

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Traffic Logs

Traffic logs display an entry for the start and end of each session. Each entry includes the following information: date and time; source and destination zones, source and destination dynamic address groups, addresses and ports; application name; security rule applied to the traffic flow; rule action (allow, deny, or drop); ingress and egress interface; number of bytes; and session end reason.

A dynamic address group only appears in a log if the rule the traffic matches includes a dynamic address group. If an IP address appears in more than one dynamic address group, the firewall displays up to five dynamic address groups in logs along with the source IP address

The Type column indicates whether the entry is for the start or end of the session. The Action column indicates whether the firewall allowed, denied, or dropped the session. A drop indicates the security rule that blocked the traffic specified any application, while a deny indicates the rule identified a specific application. If the firewall drops traffic before identifying the application, such as when a rule drops all traffic for a specific service, the Application column displays not-applicable.

Click beside an entry to view additional details about the session, such as whether an ICMP entry aggregates multiple sessions between the same source and destination (in which case the Count column value is greater than one).

When the Decryption log introduced in PAN-OS 10.1 is disabled, the firewall sends HTTP/2 logs as Traffic logs. However, when the Decryption logs are enabled, the firewall sends HTTP/2 logs as Tunnel Inspection logs (when Decryption logs are disabled, HTTP/2 logs are sent as Traffic logs), so you need to check the Tunnel Inspection logs instead of the Traffic logs for HTTP/2 events.

Threat Logs

Threat logs display entries when traffic matches one of the Security Profiles attached to a security rule on the firewall. Each entry includes the following information: date and time; type of threat (such as virus or spyware); threat description or URL (Name column); source and destination zones, addresses, source and destination dynamic address groups, and ports; application name; alarm action (such as allow or block); and severity level.

A dynamic address group only appears in a log if the rule the traffic matches includes a dynamic address group. If an IP address appears in more than one dynamic address group, the firewall displays up to five dynamic address groups in logs along with the source IP address

To see more details on individual Threat log entries:

- Click beside a threat entry to view details such as whether the entry aggregates multiple threats of the same type between the same source and destination (in which case the Count column value is greater than one).
- If you configured the firewall to Take Packet Captures, click beside an entry to access the captured packets.

The following table summarizes the Threat severity levels:

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Severity	Description
Critical	Serious threats, such as those that affect default installations of widely deployed software, result in root compromise of servers, and the exploit code is widely available to attackers. The attacker usually does not need any special authentication credentials or knowledge about the individual victims and the target does not need to be manipulated into performing any special functions.
High	Threats that have the ability to become critical but have mitigating factors; for example, they may be difficult to exploit, do not result in elevated privileges, or do not have a large victim pool. WildFire Submissions log entries with a malicious verdict and an action set to allow are logged as High.
Medium	Minor threats in which impact is minimized, such as DoS attacks that do not compromise the target or exploits that require an attacker to reside on the same LAN as the victim, affect only non-standard configurations or obscure applications, or provide very limited access. • Threat log entries with a malicious verdict and an action of block or alert, based on the existing WildFire signature severity, are logged as Medium.
Low	Warning-level threats that have very little impact on an organization's infrastructure. They usually require local or physical system access and may often result in victim privacy or DoS issues and information leakage. • Data Filtering profile matches are logged as Low. • WildFire Submissions log entries with a grayware verdict and any action are logged as Low.
Informational	Suspicious events that do not pose an immediate threat, but that are reported to call attention to deeper problems that could possibly exist. • URL Filtering log entries are logged as Informational. • WildFire Submissions log entries with a benign verdict and any action are logged as Informational. • WildFire Submissions log entries with any verdict and an action set to block and forward are logged as Informational. • Log entries with any verdict and an action set to block are logged as Informational.

URL Filtering Logs

URL Filtering logs display entries for traffic that matches either URL Filtering profiles attached to Security policy rules or URL categories used as match criteria in Security policy rules. For example, the firewall generates a log if a rule blocks access to specific websites and URL categories or if you configured a rule to generate an alert when a user accesses a website.

https://t.me/learningnets

WildFire Submissions Logs

The firewall forwards samples (files and emails links) to the WildFire cloud for analysis based on WildFire Analysis profiles settings (**Objects** > **Security Profiles** > **WildFire Analysis**). The firewall generates WildFire Submissions log entries for each sample it forwards after WildFire completes static and dynamic analysis of the sample. WildFire Submissions log entries include the firewall Action for the sample (allow or block), the WildFire verdict for the submitted sample, and the severity level of the sample.

The following table summarizes the WildFire verdicts:

Verdict	Description
Benign	Indicates that the entry received a WildFire analysis verdict of benign. Files categorized as benign are safe and do not exhibit malicious behavior.
Grayware	Indicates that the entry received a WildFire analysis verdict of grayware. Files categorized as grayware do not pose a direct security threat, but might display otherwise obtrusive behavior. Grayware can include, adware, spyware, and Browser Helper Objects (BHOs).
Phishing	Indicates that WildFire assigned a link an analysis verdict of phishing. A phishing verdict indicates that the site to which the link directs users displayed credential phishing activity.
Malicious	Indicates that the entry received a WildFire analysis verdict of malicious. Samples categorized as malicious are can pose a security threat. Malware can include viruses, C2 (command-and-control), worms, Trojans, Remote Access Tools (RATs), rootkits, and botnets. For samples that are identified as malware, the WildFire cloud generates and distributes a signature to prevent against future exposure. *C2 samples are classified as C2 in the WildFire analysis report and other Palo Alto Networks products that rely on WildFire analysis data; however, that verdict is translated and categorized as malicious by the firewall.*

Data Filtering Logs

Data Filtering logs display entries for the security rules that help prevent sensitive information such as credit card numbers from leaving the area that the firewall protects. See Data Filtering for information on defining Data Filtering profiles.

This log type also shows information for File Blocking Profiles. For example, if a rule blocks .exe files, the log shows the blocked files.

Correlation Logs

The firewall logs a correlated event when the patterns and thresholds defined in a Correlation Object match the traffic patterns on your network. To Interpret Correlated Events and view a graphical display of the events, see Use the Compromised Hosts Widget in the ACC.

https://t.me/learningnets

The following table summarizes the Correlation log severity levels:

Severity	Description
Critical	Confirms that a host has been compromised based on correlated events that indicate an escalation pattern. For example, a critical event is logged when a host that received a file with a malicious verdict by WildFire, exhibits the same command-and control activity that was observed in the WildFire sandbox for that malicious file.
High	Indicates that a host is very likely compromised based on a correlation between multiple threat events, such as malware detected anywhere on the network that matches the command and control activity being generated from a particular host.
Medium	Indicates that a host is likely compromised based on the detection of one or multiple suspicious events, such as repeated visits to known malicious URLs that suggests a scripted command-and-control activity.
Low	Indicates that a host is possibly compromised based on the detection of one or multiple suspicious events, such as a visit to a malicious URL or a dynamic DNS domain.
Informational	Detects an event that may be useful in aggregate for identifying suspicious activity; each event is not necessarily significant on its own.

Tunnel Inspection Logs

Tunnel inspection logs are like traffic logs for tunnel sessions; they display entries of non-encrypted tunnel sessions. To prevent double counting, the firewall saves only the inner flows in traffic logs, and sends tunnel sessions to the tunnel inspection logs. The tunnel inspection log entries include Receive Time (date and time the log was received), the tunnel ID, monitor tag, session ID, the Security rule applied to the tunnel session, number of bytes in the session, parent session ID (session ID for the tunnel session), source address, source user and source zone, destination address, destination user, and destination zone.

When the Decryption logs introduced in PAN-OS 10.1 are enabled, the firewall sends HTTP/2 logs as Tunnel Inspection logs (when Decryption logs are disabled, HTTP/2 logs are sent as Traffic logs), so you need to check the Tunnel Inspection logs instead of the Traffic logs for HTTP/2 events. In this case, you must also enable Tunnel Content Inspection *to obtain the App-ID for HTTP/2 traffic.*

Click the Detailed Log view to see details for an entry, such as the tunnel protocol used, and the flag indicating whether the tunnel content was inspected or not. Only a session that has a parent session will have the Tunnel Inspected flag set, which means the session is in a tunnel-in-tunnel (two levels of encapsulation). The first outer header of a tunnel will not have the Tunnel Inspected flag set.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Config Logs

Config logs display entries for changes to the firewall configuration. Each entry includes the date and time, the administrator username, the IP address from where the administrator made the change, the type of client (Web, CLI, or Panorama), the type of command executed, the command status (succeeded or failed), the configuration path, and the values before and after the change.

System Logs

System logs display entries for each system event on the firewall. Each entry includes the date and time, event severity, and event description. The following table summarizes the System log severity levels. For a partial list of System log messages and their corresponding severity levels, refer to System Log Events.

Severity	Description
Critical	Hardware failures, including high availability (HA) failover and link failures.
High	Serious issues, including dropped connections with external devices, such as LDAP and RADIUS servers.
Medium	Mid-level notifications, such as antivirus package upgrades.
Low	Minor severity notifications, such as user password changes.
Informational	Log in/log off, administrator name or password change, any configuration change, and all other events not covered by the other severity levels.

HIP Match Logs

The GlobalProtect Host Information Profile (HIP) matching enables you to collect information about the security status of the end devices accessing your network (such as whether they have disk encryption enabled). The firewall can allow or deny access to a specific host based on adherence to the HIP-based security rules you define. HIP Match logs display traffic flows that match a HIP Object or HIP Profile that you configured for the rules.

GlobalProtect Logs

GlobalProtect logs display the following logs related to GlobalProtect:

- GlobalProtect system logs.

  GlobalProtect authentication event logs remain in **Monitor** > **Logs** > **System**; however, the **Auth Method** column of the GlobalProtect logs display the authentication method used for logins.
- LSVPN/satellite events.
- GlobalProtect portal and gateway logs.
- Clientless VPN logs.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

IP-Tag Logs

IP-tag logs display how and when a source IP address is registered or unregistered on the firewall and what tag the firewall applied to the address. Additionally, each log entry displays the configured timeout (when configured) and the source of the IP address-to-tag mapping information, such as User-ID agent VM information sources and auto-tagging. See how to Register IP Address and Tags Dynamically for more information.

User-ID Logs

User-ID logs display information about IP address-to-username mappings and Authentication Timestamps, such as the sources of the mapping information and the times when users authenticated. You can use this information to help troubleshoot User-ID and authentication issues. For example, if the firewall is applying the wrong policy rule for a user, you can view the logs to verify whether that user is mapped to the correct IP address and whether the group associations are correct.

Decryption Logs

Decryption Logs display entries for unsuccessful TLS handshakes by default and can display entries for successful TLS handshakes if you enable them in Decryption policy. If you enable entries for successful handshakes, ensure that you have the system resources (log space) for the logs.

Decryption logs include a vast amount of information to help you Troubleshoot and Monitor Decryption and then resolve issues. There are 62 columns of different types of information you can enable in the logs, and you can select any individual log (, the magnifying glass) and see the details in a single Detail view. You can view certificate, cipher suite, and error information such as: subject common name, issuer common name, root common name, root status, certificate key type and size, certificate start and end date, certificate serial number, certificate fingerprint, TLS version, key exchange algorithm, encryption algorithm, negotiated EC curve, authentication algorithm, SNI, proxy type, errors information (cipher, HSM, resource, resume, protocol, feature, certificate, version), and error indexes (codes that you can look up to get more error information).

Alarms Logs

An alarm is a firewall-generated message indicating that the number of events of a particular type (for example, encryption and decryption failures) has exceeded the threshold configured for that event type. To enable alarms and configure alarm thresholds, select **Device** > **Log Settings** and edit the Alarm Settings.

When generating an alarm, the firewall creates an Alarm log and opens the System Alarms dialog to display the alarm. After you **Close** the dialog, you can reopen it anytime by clicking **Alarms** ( ) at the bottom of the web interface. To prevent the firewall from automatically opening the dialog for a particular alarm, select the alarm in the Unacknowledged Alarms list and **Acknowledge** the alarm.

Authentication Logs

Authentication logs display information about authentication events that occur when end users try to access network resources for which access is controlled by Authentication Policy rules. You can use this information to help troubleshoot access issues and to adjust your Authentication

https://t.me/learningnets

policy as needed. In conjunction with correlation objects, you can also use Authentication logs to identify suspicious activity on your network, such as brute force attacks.

Optionally, you can configure Authentication rules to log timeout events. These timeouts relate to the period when a user need authenticate for a resource only once but can access it repeatedly. Seeing information about the timeouts helps you decide if and how to adjust them (for details, see Authentication Timestamps).

System logs record authentication events relating to GlobalProtect and to administrator access to the web interface.

Unified Logs

Unified logs are entries from the Traffic, Threat, URL Filtering, WildFire Submissions, and Data Filtering logs displayed in a single view. Unified log view enables you to investigate and filter the latest entries from different log types in one place, instead of searching through each log type separately. Click Effective Queries ( ) in the filter area to select which log types will display entries in Unified log view.

The Unified log view displays only entries from logs that you have permission to see. For example, an administrator who does not have permission to view WildFire Submissions logs will not see WildFire Submissions log entries when viewing Unified logs. Administrative Role Types define these permissions.

When you Set Up Remote Search in AutoFocus to perform a targeted search on the firewall, the search results are displayed in Unified log view.

View Logs

You can view the different log types on the firewall in a tabular format. The firewall locally stores all log files and automatically generates Configuration and System logs by default. To learn more about the security rules that trigger the creation of entries for the other types of logs, see Log Types and Severity Levels.

To configure the firewall to forward logs as syslog messages, email notifications, or Simple Network Management Protocol (SNMP) traps, Use External Services for Monitoring.

STEP 1 | Select a log type to view.

1. Select **Monitor** > **Logs**.
2. Select a log type from the list.

   The firewall displays only the logs you have permission to see. For example, if your administrative account does not have permission to view WildFire Submissions logs, the firewall does not display that log type when you access the logs pages. Administrative Role Types define the permissions.

STEP 2 | (Optional) Customize the log column display.

1. Click the arrow to the right of any column header, and select **Columns**.
2. Select columns to display from the list. The log updates automatically to match your selections.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 3 | View additional details about log entries.

- Click the spyglass ( ) for a specific log entry. The Detailed Log View has more information about the source and destination of the session, as well as a list of sessions related to the log entry.
- (Threat log only) Click next to an entry to access local packet captures of the threat. To enable local packet captures, see Take Packet Captures.
- (Traffic, Threat, URL Filtering, WildFire Submissions, Data Filtering, and Unified logs only) View AutoFocus threat data for a log entry.
  1. Enable AutoFocus Threat Intelligence.

     *Enable AutoFocus in Panorama to view AutoFocus threat data for all Panorama log entries, including those from firewalls that are not connected to AutoFocus and/or are running PAN-OS 7.0 and earlier release versions (**Panorama** > **Setup** > **Management** > **AutoFocus**).*
  2. Hover over an IP address, URL, user agent, threat name (subtype: virus and wildfire-virus only), filename, or SHA-256 hash.
  3. Click the drop-down ( ) and select **AutoFocus**.
  4. Content Delivery Network Infrastructure.

Next Steps...

- Filter Logs.
- Export Logs.
- Configure Log Storage Quotas and Expiration Periods.

Filter Logs

Each log has a filter area that allows you to set a criteria for which log entries to display. The ability to filter logs is useful for focusing on events on your firewall that possess particular properties or attributes. Filter logs by artifacts that are associated with individual log entries.

For example, filtering by the rule UUID makes it easier to pinpoint the specific rule you want to locate, even among many similarly-named rules. If your ruleset is very large and contains many rules, using the rule's UUID as a filter spotlights the particular rule you need to find without having to navigate through pages of results.

STEP 1 | (Unified logs only) Select the log types to include in the Unified log display.

1. Click Effective Queries ( ).
2. Select one or more log types from the list (**traffic**, **threat**, **url**, **data**, and **wildfire**).
3. Click **OK**. The Unified log updates to show only entries from the log types you have selected.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 2 | Add a filter to the filter field.

> *If the value of the artifact matches the operator (such as **has** or **in**), enclose the value in quotation marks to avoid a syntax error. For example, if you filter by destination country and use IN as a value to specify INDIA, enter the filter as **(dstloc eq "IN")**.*

- Click one or more artifacts (such as the application type associated with traffic and the IP address of an attacker) in a log entry. For example, click the Source **10.0.0.25** and Application **web-browsing** of a log entry to display only entries that contain both artifacts in the log (AND search).
- To specify artifacts to add to the filter field, click Add Filter (⊕).
- To add a previously saved filter, click Load Filter ( ).

STEP 3 | Apply the filter to the log.

Click Apply Filter (→). The log will refresh to display only log entries that match the current filter.

STEP 4 | (Optional) Save frequently used filters.

1. Click Save Filter ( ).
2. Enter a **Name** for the filter.
3. Click **OK**. You can view your saved filters by clicking Load Filter ( ).

Next Steps...

- View Logs.
- Export Logs.

Export Logs

You can export the contents of a log type to a comma-separated value (CSV) formatted report. By default, the report contains up to 2,000 rows of log entries.

STEP 1 | Set the number of rows to display in the report.

1. Select **Device** > **Setup** > **Management**, then edit the Logging and Reporting Settings.
2. Click the **Log Export and Reporting** tab.
3. Edit the number of **Max Rows in CSV Export** (up to 1048576 rows).
4. Click **OK**.

STEP 2 | Download the log.

1. Click Export to CSV ( ). A progress bar showing the status of the download appears.
2. When the download is complete, click **Download file** to save a copy of the log to your local folder. For descriptions of the column headers in a downloaded log, refer to Syslog Field Descriptions.

Next Step...

Schedule Log Exports to an SCP or FTP Server.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Configure Log Storage Quotas and Expiration Periods

The firewall automatically deletes logs that exceed the expiration period. When the firewall reaches the storage quota for a log type, it automatically deletes older logs of that type to create space even if you don't set an expiration period.

*If you want to manually delete logs, select **Device** > **Log Settings** and, in the Manage Logs section, click the links to clear logs by type.*

STEP 1 | Select **Device** > **Setup** > **Management** and edit the Logging and Reporting Settings.

STEP 2 | Select **Log Storage** and enter a **Quota (%)** for each log type. When you change a percentage value, the dialog refreshes to display the corresponding absolute value (Quota GB/MB column).

STEP 3 | Enter the **Max Days** (expiration period) for each log type (range is 1-2,000). The fields are blank by default, which means the logs never expire.

The firewall synchronizes expiration periods across high availability (HA) pairs. Because only the active HA peer generates logs, the passive peer has no logs to delete unless failover occurs and it starts generating logs.

STEP 4 | Click **OK** and **Commit**.

Schedule Log Exports to an SCP or FTP Server

You can schedule exports of Traffic, Threat, URL Filtering, Data Filtering, HIP Match, and WildFire Submission logs to a Secure Copy (SCP) server or File Transfer Protocol (FTP) server. Perform this task for each log type you want to export.

You can use Secure Copy (SCP) commands from the CLI to export the entire log database to an SCP server and import it to another firewall. Because the log database is too large for an export or import to be practical on the following platforms, they do not support these options: PA-7000 Series firewalls (all PAN-OS releases), Panorama virtual appliance running Panorama 6.0 or later releases, and Panorama M-Series appliances (all Panorama releases).

STEP 1 | Select **Device** > **Scheduled Log Export** and click **Add**.

STEP 2 | Enter a **Name** for the scheduled log export and **Enable** it.

STEP 3 | Select the **Log Type** to export.

STEP 4 | Select the daily **Scheduled Export Start Time**. The options are in 15-minute increments for a 24-hour clock (00:00 - 23:59).

STEP 5 | Select the **Protocol** to export the logs: **SCP** (secure) or **FTP**.

STEP 6 | Enter the **Hostname** or IP address of the server.

STEP 7 | Enter the **Port** number. By default, FTP uses port 21 and SCP uses port 22.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 8 | Enter the **Path** or directory in which to save the exported logs.

STEP 9 | Enter the **Username** and, if necessary, the **Password** (and **Confirm Password**) to access the server.

STEP 10 | **(**FTP only**)** Select **Enable FTP Passive Mode** if you want to use FTP passive mode, in which the firewall initiates a data connection with the FTP server. By default, the firewall uses FTP active mode, in which the FTP server initiates a data connection with the firewall. Choose the mode based on what your FTP server supports and on your network requirements.

STEP 11 | **(**SCP only**)** Click **Test SCP server connection**. Before establishing a connection, the firewall must accept the host key for the SCP server.



If you use a Panorama template to configure the log export schedule, you must perform this step after committing the template configuration to the firewalls. After the template commit, log in to each firewall, open the log export schedule, and click ***Test SCP server connection****.*

STEP 12 | Click **OK** and **Commit**.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Monitor Block List

There are two ways you can cause the firewall to place an IP address on the block list:

- Configure a Vulnerability Protection profile with a rule to Block IP connections and apply the profile to a Security policy, which you apply to a zone.
- Configure a DoS Protection policy rule with the Protect action and a Classified DoS Protection profile, which specifies a maximum rate of connections per second allowed. When incoming packets match the DoS Protection policy and exceed the Max Rate, and if you specified a Block Duration and a Classified policy rule to include source IP address, the firewall puts the offending source IP address on the block list.

In the cases described above, the firewall automatically blocks that traffic in hardware before those packets use CPU or packet buffer resources. If attack traffic exceeds the blocking capacity of the hardware, the firewall uses IP blocking mechanisms in software to block the traffic.

The firewall automatically creates a hardware block list entry based on your Vulnerability Protection profile or DoS Protection policy rule; the source address from the rule is the source IP address in the hardware block list.

Entries on the block list indicate in the Type column whether they were blocked by hardware (hw) or software (sw). The bottom of the screen displays:

- Count of **Total Blocked IPs** out of the number of blocked IP addresses the firewall supports.
- Percentage of the block list that the firewall has used.

To view details about an address on the block list, hover over a Source IP address and click the down arrow link. Click the Who Is link, which displays the Network Solutions Who Is feature, providing information about the address.

For information on configuring a Vulnerability Protection profile, see Customize the Action and Trigger Conditions for a Brute Force Signature. For more information on block list and DoS Protection profiles, see DoS Protection Against Flooding of New Sessions.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

View and Manage Reports

The reporting capabilities on the firewall allow you to keep a pulse on your network, validate your policies, and focus your efforts on maintaining network security for keeping your users safe and productive.

- Report Types
- View Reports
- Configure the Expiration Period and Run Time for Reports
- Disable Predefined Reports
- Custom Reports
- Generate Custom Reports
- Generate Botnet Reports
- Generate the SaaS Application Usage Report
- Manage PDF Summary Reports
- Generate User/Group Activity Reports
- Manage Report Groups
- Schedule Reports for Email Delivery
- Manage Report Storage Capacity

Report Types

The firewall includes predefined reports that you can use as-is, or you can build custom reports that meet your needs for specific data and actionable tasks, or you can combine predefined and custom reports to compile information you need. The firewall provides the following types of reports:

- **Predefined Reports**—Allow you to view a quick summary of the traffic on your network. A suite of predefined reports are available in four categories—Applications, Traffic, Threat, and URL Filtering. See View Reports.
- **User or Group Activity Reports**—Allow you to schedule or create an on-demand report on the application use and URL activity for a specific user or for a user group. The report includes the URL categories and an estimated browse time calculation for individual users. See Generate User/Group Activity Reports.
- **Custom Reports**—Create and schedule custom reports that show exactly the information you want to see by filtering on conditions and columns to include. You can also include query builders for more specific drill down on report data. See Generate Custom Reports.
- **PDF Summary Reports**—Aggregate up to 18 predefined or custom reports/graphs from Threat, Application, Trend, Traffic, and URL Filtering categories into one PDF document. See Manage PDF Summary Reports.
- **Botnet Reports**—Allow you to use behavior-based mechanisms to identify potential botnet-infected hosts in the network. See Generate Botnet Reports.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

- **Report Groups**—Combine custom and predefined reports into report groups and compile a single PDF that is emailed to one or more recipients. See Manage Report Groups.

Reports can be generated on demand, on a recurring schedule, and can be scheduled for email delivery.

View Reports

The firewall provides an assortment of over 40 predefined reports that it generates every day. You can view these reports directly on the firewall. You can also view custom reports and summary reports.

About 200 MB of storage is allocated for saving reports on the firewall. This limit can be reconfigured for PA-7000 series and PA-5200 series firewalls only. For all other firewall models, you can Configure the Expiration Period and Run Time for Reports to allow the firewall to delete reports that exceed the period. Keep in mind that when the firewall reaches its storage limit, it automatically deletes older reports to create space even if you don't set an expiration period. Another way to conserve system resources on the firewall is to Disable Predefined Reports. For long-term retention of reports, you can export the reports (as described below) or Schedule Reports for Email Delivery.

Unlike other reports, you can't save User/Group Activity reports on the firewall. You must Generate User/Group Activity Reports on demand or schedule them for email delivery.

STEP 1 | (VM-50, VM-50 Lite, and PA-200 firewalls only) Enable generation of predefined reports.

By default, predefined reports are disabled on VM-50, VM-50 Lite, and PA-200 firewalls to save resources.

1. Select **Device** > **Setup** > **Management** and edit **Logging and Reporting**.
2. Select **Pre-Defined Reports** and enable (check) **Pre-Defined Reports**.
3. Check (enable) the predefined reports you want to generate and click **OK**
4. **Commit** your configuration changes.
5. Access the firewall CLI to enable predefined reports.

   This step is required for local predefined reports and predefined reports pushed from a Panorama™ management server.

   ```
   admin> debug predefined-default enable
   ```

STEP 2 | Select **Monitor** > **Reports**.

The reports are grouped into sections (types) on the right-hand side of the page: **Custom Reports**, **Application Reports**, **Traffic Reports**, **Threat Reports**, **URL Filtering Reports**, and **PDF Summary Reports**.

STEP 3 | Select a report to view. The reports page then displays the report for the previous day.

To view reports for other days, select a date in the calendar at the bottom right of the page and select a report. If you select a report in another section, the date selection resets to the current date.

https://t.me/learningnets

STEP 4 | To view a report offline, you can export the report to PDF, CSV or to XML formats. Click **Export to PDF**, **Export to CSV**, or **Export to XML** at the bottom of the page, then print or save the file.

Configure the Expiration Period and Run Time for Reports

The expiration period and run time are global settings that apply to all Report Types. After running new reports, the firewall automatically deletes reports that exceed the expiration period.

STEP 1 | Select **Device** > **Setup** > **Management**, edit the Logging and Reporting Settings, and select the **Log Export and Reporting** tab.

STEP 2 | Set the **Report Runtime** to an hour in the 24-hour clock schedule (default is 02:00; range is 00:00 [midnight] to 23:00).

STEP 3 | Enter the **Report Expiration Period** in days (default is no expiration; range is 1 is 2,000).

*You can't change the storage that the firewall allocates for saving reports: it is predefined at about 200 MB. When the firewall reaches the storage maximum, it automatically deletes older reports to create space even if you don't set a **Report Expiration Period**.*

STEP 4 | Click **OK** and **Commit**.

Disable Predefined Reports

The firewall includes about 40 predefined reports that it automatically generates daily. If you do not use some or all of these, you can disable selected reports to conserve system resources on the firewall.

Make sure that no report group or PDF summary report includes the predefined reports you will disable. Otherwise, the firewall will render the PDF summary report or report group without any data.

STEP 1 | Select **Device** > **Setup** > **Management** and edit the Logging and Reporting Settings.

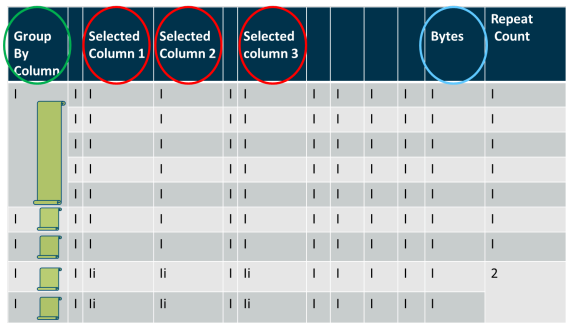
STEP 2 | Select the **Pre-Defined Reports** tab and clear the check box for each report you want to disable. To disable all predefined reports, click **Deselect All**.

STEP 3 | Click **OK** and **Commit**.

Custom Reports

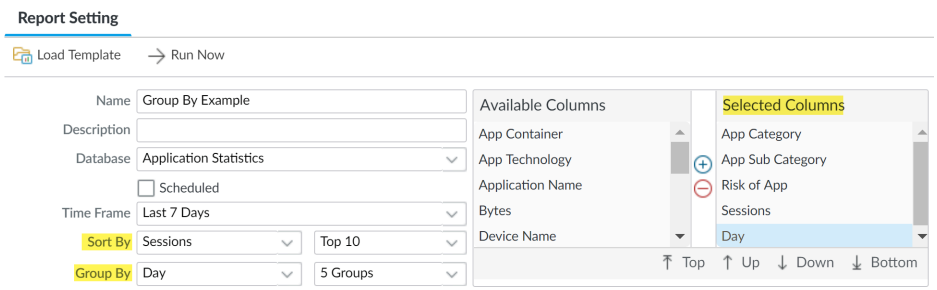
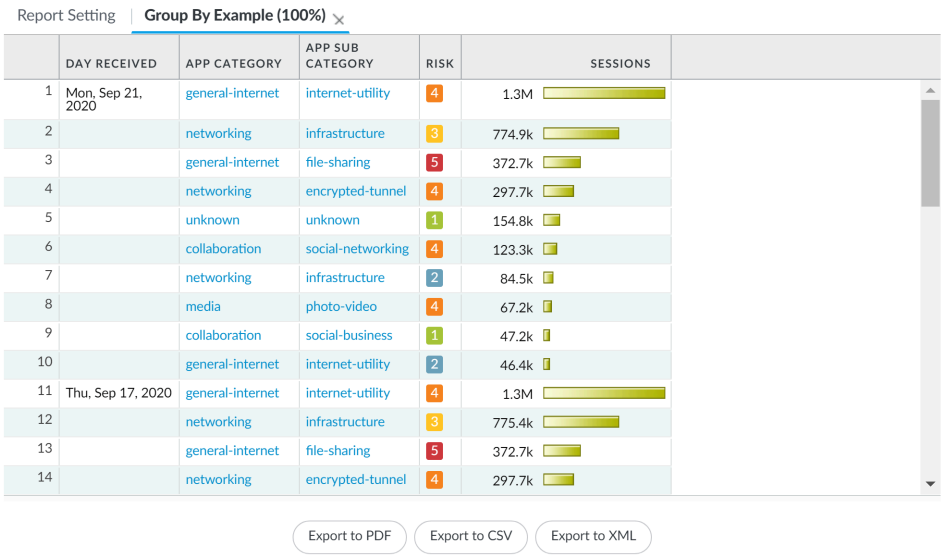
In order to create purposeful custom reports, you must consider the attributes or key pieces of information that you want to retrieve and analyze, such as threats, as well as the best way to categorize the information, such as grouping by rule UUID, which will allow you to see the rule that applies to each threat type. This consideration guides you in making the following selections in a custom report:

https://t.me/learningnets

Selection	Description
Database	You can base the report on one of the following database types: • **Summary databases**—These databases are available for Application Statistics, Traffic, Threat, URL Filtering, and Tunnel Inspection logs. The firewall aggregates the detailed logs at 15-minute intervals. To enable faster response time when generating reports, the firewall condenses the data: duplicate sessions are grouped and incremented with a repeat counter, and some attributes (columns) are excluded from the summary. • **Detailed logs**—These databases itemize the logs and list all the attributes (columns) for each log entry. *Reports based on detailed logs take much longer to run and are not recommended unless absolutely necessary.*
Attributes	The columns that you want to use as the match criteria. The attributes are the columns that are available for selection in a report. From the list of **Available Columns**, you can add the selection criteria for matching data and for aggregating the details (the **Selected Columns**).
Sort By/ Group By	The **Sort By** and the **Group By** criteria allow you to organize/segment the data in the report; the sorting and grouping attributes available vary based on the selected data source. The Sort By option specifies the attribute that is used for aggregation. If you do not select an attribute to sort by, the report will return the first N number of results without any aggregation. The Group By option allows you to select an attribute and use it as an anchor for grouping data; all the data in the report is then presented in a set of top 5, 10, 25 or 50 groups. For example, when you select Hour as the Group By selection and want the top 25 groups for a 24-hr time period, the results of the report will be generated on an hourly basis over a 24-hr period. The first column in the report will be the hour and the next set of columns will be the rest of your selected report columns.
	The following example illustrates how the **Selected Columns** and **Sort By**/**Group By** criteria work together when generating reports: 

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Selection	Description
	The columns circled in red (above) depict the columns selected, which are the attributes that you match against for generating the report. Each log entry from the data source is parsed and these columns are matched on. If multiple sessions have the same values for the selected columns, the sessions are aggregated and the repeat count (or sessions) is incremented. The column circled in blue indicates the chosen sort order. When the sort order (**Sort By**) is specified, the data is sorted (and aggregated) by the selected attribute. The column circled in green indicates the **Group By** selection, which serves as an anchor for the report. The **Group By** column is used as a match criteria to filter for the top N groups. Then, for each of the top N groups, the report enumerates the values for all the other selected columns.
	For example, if a report has the following selections:  The output will display as follows:  The report is anchored by **Day** and sorted by **Sessions**. It lists the 5 days (**5 Groups**) with maximum traffic in the **Last 7 Days** time frame. The

https://t.me/learningnets

Selection	Description
	data is enumerated by the **Top 5** sessions for each day for the selected columns—**App Category**, **App Subcategory** and **Risk**.
Time Frame	The date range for which you want to analyze data. You can define a custom range or select a time period ranging from the last 15 minutes to the last 30 days. The reports can be run on demand or scheduled to run at a daily or weekly cadence.
Query Builder	The query builder allows you to define specific queries to further refine the selected attributes. It allows you see just what you want in your report using **and** and **or** operators and a match criteria, and then include or exclude data that matches or negates the query in the report. Queries enable you to generate a more focused collation of information in a report.

Generate Custom Reports

You can configure custom reports that the firewall generates immediately (on demand) or on schedule (each night). To understand the selections available to create a purposeful custom report, see Custom Reports.

After the firewall has generated a scheduled custom report, you risk invalidating the past results of that report if you modify its configuration to change its future output. If you need to modify a scheduled report configuration, the best practice is to create a new report.

STEP 1 | Select **Monitor** > **Manage Custom Reports**.

STEP 2 | Click **Add** and then enter a **Name** for the report.

*To base a report on an predefined template, click **Load Template** and choose the template. You can then edit the template and save it as a custom report.*

STEP 3 | Select the **Database** to use for the report.

Each time you create a custom report, a log view report is automatically created. This report show the logs that were used to build the custom report. The log view report uses the same name as the custom report, but appends the phrase (Log View) to the report name.

When creating a report group, you can include the log view report with the custom report. For more information, see Manage Report Groups.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 4 | Select the **Scheduled** check box to run the report each night. The report is then available for viewing in the **Reports** column on the side.

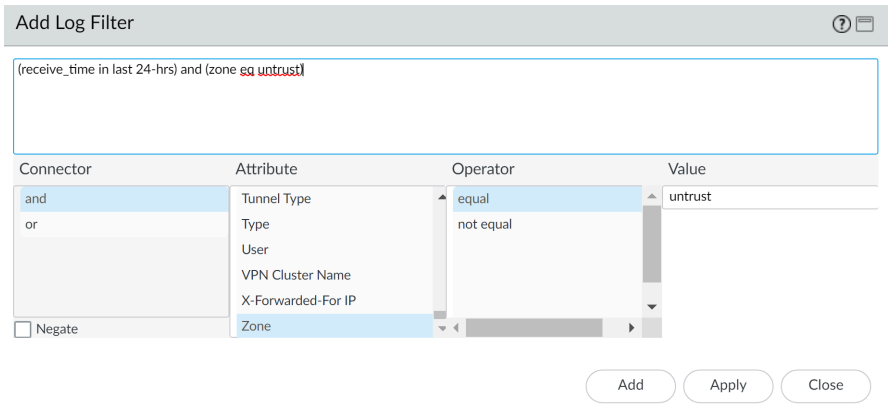
> *To generate a scheduled custom report using logs stored in Cortex Data Lake on the Panorama™ management server, Cloud Service plugin 1.8 or later release must be installed on Panorama.*

STEP 5 | Define the filtering criteria. Select the **Time Frame**, the **Sort By** order, **Group By** preference, and select the columns that must display in the report.

STEP 6 | (Optional) Select the **Query Builder** attributes if you want to further refine the selection criteria. To build a report query, specify the following and click **Add**. Repeat as needed to construct the full query.

- **Connector**—Choose the connector (and/or) to precede the expression you are adding.
- **Negate**—Select the check box to interpret the query as a negation. If, for example, you choose to match entries in the last 24 hours and/or are originating from the untrust zone, the negate option causes a match on entries that are not in the past 24 hours and/or are not from the untrust zone.
- **Attribute**—Choose a data element. The available options depend on the choice of database.
- **Operator**—Choose the criterion to determine whether the attribute applies (such as =). The available options depend on the choice of database.
- **Value**—Specify the attribute value to match.

For example, the following figure (based on the `Traffic Log` database) shows a query that matches if the Traffic log entry was received in the past 24 hours and is from the untrust zone.



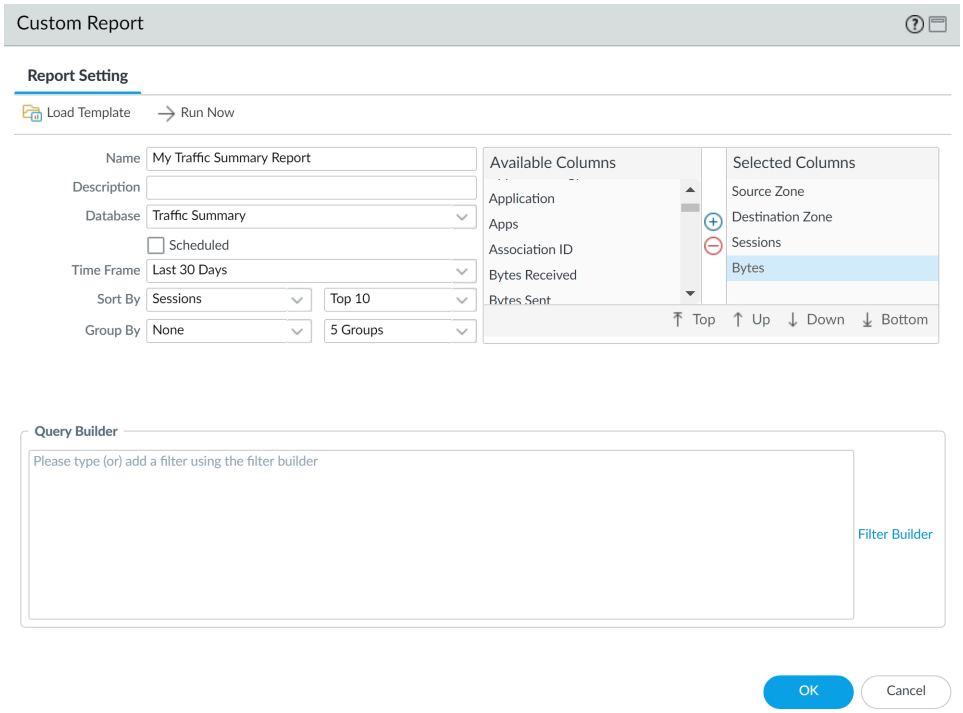
STEP 7 | To test the report settings, select **Run Now**. Modify the settings as required to change the information that is displayed in the report.

https://t.me/learningnets

STEP 8 | Click **OK** to save the custom report.

Examples of Custom Reports

If you want to set up a simple report in which you use the traffic summary database from the last 30 days, and sort the data by the top 10 sessions and these sessions are grouped into 5 groups by day of the week. You would set up the custom report to look like this:



And the PDF output for the report would look as follows:

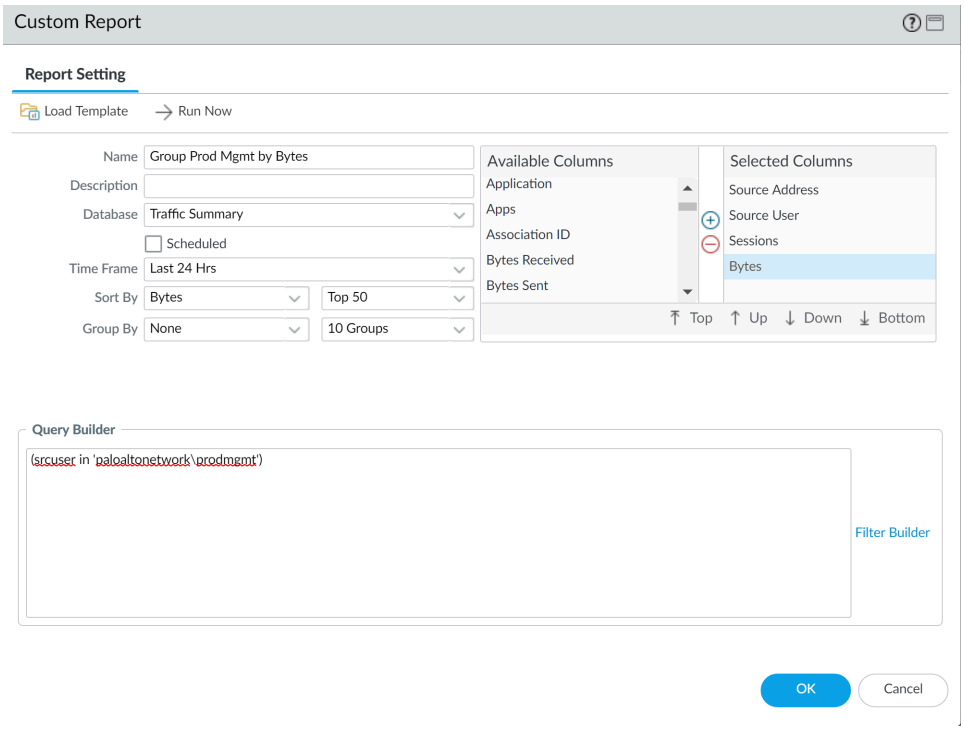
My Traffic Summary Report

ca1demo.paloaltonetworks.com : 2016/01/25 10:34:39 - 2016/02/24 10:34:38

Source Zone	Destination Zone	App Category	Application	Sessions	Bytes
Tap	Tap	general-internet	web-browsing	74.54 M	2.47 T
Tap	Tap	networking	dns	52.03 M	28.93 G
Tap	Tap	networking	ssl	18.01 M	678.13 G
Tap	Tap	general-internet	bittorrent	9.80 M	1.62 T
Tap	Tap	general-internet	google-base	4.48 M	168.99 G
Tap	Tap	unknown	insufficient-data	4.45 M	31.30 G
Tap	Tap	collaboration	facebook-base	4.09 M	99.14 G
Tap	Tap	networking	ntp	4.07 M	3.29 G
Tap	Tap	collaboration	blackboard	2.84 M	186 G
Tap	Tap	collaboration	smtp	1.92 M	172.57 G
Tap	Tap	networking	icmp	1.36 M	320.49 M
Tap	Tap	general-internet	gnutella	1.17 M	17.84 G
Tap	Tap	collaboration	myspace-base	1.10 M	35.22 G
Tap	Tap	general-internet	ping	1.06 M	86.21 M
Tap	Tap	general-internet	flash	1.01 M	168.14 G

Now, if you want to use the query builder to generate a custom report that represents the top consumers of network resources within a user group, you would set up the report to look like this:

https://t.me/learningnets



The report would display the top users in the product management user group sorted by bytes.

Generate Botnet Reports

The botnet report enables you to use heuristic and behavior-based mechanisms to identify potential malware- or botnet-infected hosts in your network. To evaluate botnet activity and infected hosts, the firewall correlates user and network activity data in Threat, URL, and Data Filtering logs with the list of malware URLs in PAN-DB, known dynamic DNS domain providers, and domains registered within the last 30 days. You can configure the report to identify hosts that visited those sites, as well as hosts that communicated with Internet Relay Chat (IRC) servers or that used unknown applications. Malware often use dynamic DNS to avoid IP blocking, while IRC servers often use bots for automated functions.

The firewall requires Threat Prevention and URL Filtering licenses to use the botnet report. You can Use the Automated Correlation Engine to monitor suspicious activities based on additional indicators besides those that the botnet report uses. However, the botnet report is the only tool that uses newly registered domains as an indicator.

- Configure a Botnet Report
- Interpret Botnet Report Output

Configure a Botnet Report

You can schedule a botnet report or run it on demand. The firewall generates scheduled botnet reports every 24 hours because behavior-based detection requires correlating traffic across multiple logs over that timeframe.

https://t.me/learningnets

STEP 1 | Define the types of traffic that indicate possible botnet activity.

1. Select **Monitor** > **Botnet** and click **Configuration** on the right side of the page.
2. **Enable** and define the **Count** for each type of HTTP Traffic that the report will include.

   The **Count** values represent the minimum number of events of each traffic type that must occur for the report to list the associated host with a higher confidence score (higher likelihood of botnet infection). If the number of events is less than the **Count**, the report will display a lower confidence score or (for certain traffic types) won't display an entry for the host. For example, if you set the **Count** to three for **Malware URL visit**, then hosts that visit three or more known malware URLs will have higher scores than hosts that visit less than three. For details, see Interpret Botnet Report Output.
3. Define the thresholds that determine whether the report will include hosts associated with traffic involving Unknown TCP or Unknown UDP applications.
4. Select the **IRC** check box to include traffic involving IRC servers.
5. Click **OK** to save the report configuration.

STEP 2 | Schedule the report or run it on demand.

1. Click **Report Setting** on the right side of the page.
2. Select a time interval for the report in the **Test Run Time Frame** drop-down.
3. Select the **No. of Rows** to include in the report.
4. **(**Optional**) Add** queries to the Query Builder to filter the report output by attributes such as source/destination IP addresses, users, or zones.

   For example, if you know in advance that traffic initiated from the IP address 10.3.3.15 contains no potential botnet activity, add **`not (addr.src in 10.0.1.35)`** as a query to exclude that host from the report output. For details, see Interpret Botnet Report Output.
5. Select **Scheduled** to run the report daily or click **Run Now** to run the report immediately.
6. Click **OK** and **Commit**.

Interpret Botnet Report Output

The botnet report displays a line for each host that is associated with traffic you defined as suspicious when configuring the report. For each host, the report displays a confidence score of 1 to 5 to indicate the likelihood of botnet infection, where 5 indicates the highest likelihood. The scores correspond to threat severity levels: 1 is informational, 2 is low, 3 is medium, 4 is high, and 5 is critical. The firewall bases the scores on:

- **Traffic type**—Certain HTTP traffic types are more likely to involve botnet activity. For example, the report assigns a higher confidence to hosts that visit known malware URLs than to hosts that browse to IP domains instead of URLs, assuming you defined both those activities as suspicious.
- **Number of events**—Hosts that are associated with a higher number of suspicious events will have higher confidence scores based on the thresholds (**Count** values) you define when you Configure a Botnet Report.
- **Executable downloads**—The report assigns a higher confidence to hosts that download executable files. Executable files are a part of many infections and, when combined with the

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

other types of suspicious traffic, can help you prioritize your investigations of compromised hosts.

When reviewing the report output, you might find that the sources the firewall uses to evaluate botnet activity (for example, the list of malware URLs in PAN-DB) have gaps. You might also find that these sources identify traffic that you consider safe. To compensate in both cases, you can add query filters when you Configure a Botnet Report.

Generate the SaaS Application Usage Report

The SaaS Application Usage PDF report is a two-part report that allows you to easily explore SaaS application activity by risk and sanction state. A sanctioned application is an application that you formally approve for use on your network. A SaaS application is an application that has the characteristic SaaS=yes in the applications details page in **Objects** > **Applications**, all other applications are considered as non-SaaS. To indicate that you have sanctioned a SaaS or non-SaaS application, you must tag it with the predefined tag named Sanctioned. The firewall and Panorama consider any application without this predefined tag as unsanctioned for use on the network.

- The first part of the report presents the key findings for the SaaS applications on your network during the reporting period with a comparison of the sanctioned versus unsanctioned applications and lists the top applications based on sanction state by usage, compliance, and data transfers. To help you identify and explore the extent of high risk application usage, the applications with risky characteristics section of the report lists the SaaS applications with the following unfavorable hosting characteristics: certifications achieved, past data breaches, support for IP-based restrictions, financial viability, and terms of service. You can also view a comparison of sanctioned versus unsanctioned SaaS applications by total number of applications used on your network, bandwidth consumed by these applications, the number of users using these applications, top user groups that use the largest number of SaaS applications, and the top user groups that transfer the largest volume of data through sanctioned and unsanctioned SaaS applications. This first part of the report also highlights the top SaaS application subcategories listed in order by maximum number of applications used, the number of users, and the amount of data (bytes) transferred in each application subcategory.
- The second part of the report focuses on the detailed browsing information for SaaS and non-SaaS applications for each application subcategory listed in the first-part of the report. For each application in a subcategory, it also includes information about the top users who transferred data, the top blocked or alerted file types, and the top threats for each application. In addition, this section of the report tallies samples for each application that the firewall submitted for WildFire analysis, and the number of samples determined to be benign and malicious.

Use the insights from this report to consolidate the list of business-critical and approved SaaS applications and to enforce policies for controlling unsanctioned and risky applications that pose unnecessary risks for malware propagation and data leaks.



The predefined SaaS application usage report is still available as a daily View Reports that lists the top 100 SaaS applications (which means applications with the SaaS application characteristic, SaaS=yes) running on your network on a given day. This report does not give visibility into applications you have designated as sanctioned, but rather gives visibility into all of the SaaS applications in use on your network.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

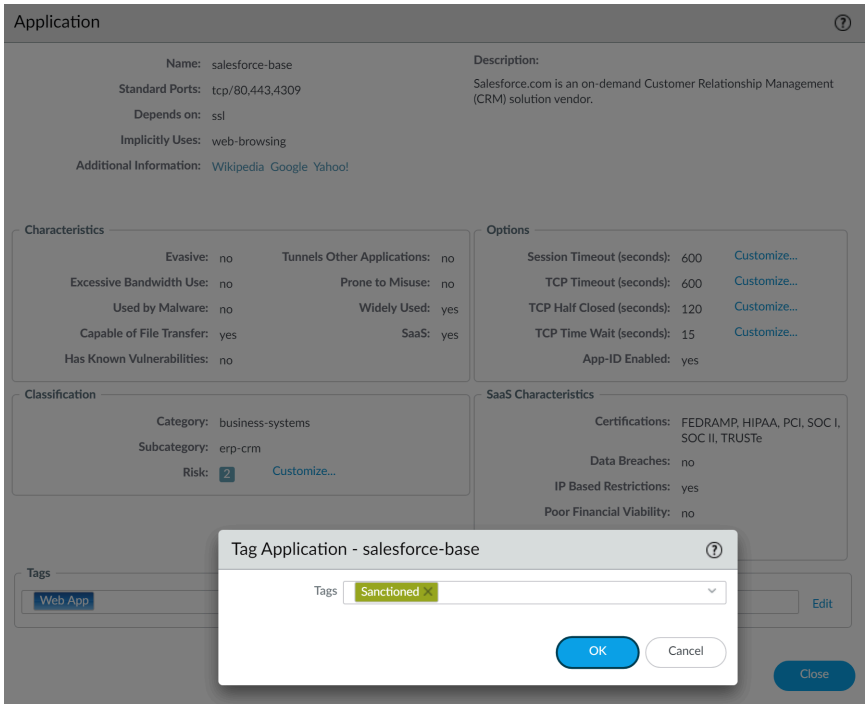
STEP 1 | Tag applications that you approve for use on your network as Sanctioned.

For generating an accurate and informative report, you need to tag the sanctioned applications consistently across firewalls with multiple virtual systems, and across firewalls that belong to a device group on Panorama. If the same application is tagged as sanctioned in one virtual system and is not sanctioned in another or, on Panorama, if an application is unsanctioned in a parent device group but is tagged as sanctioned in a child device group (or vice versa), the SaaS Application Usage report will report the application as partially sanctioned and will have overlapping results.

Example: If Box is sanctioned on vsys1 and Google Drive is sanctioned on vsys2, Google Drive users in vsys1 will be counted as users of an unsanctioned SaaS application and Box users in vsys2 will be counted as users of an unsanctioned SaaS application. The key finding in the report will highlight that a total of two unique SaaS applications are discovered on the network with two sanctioned applications and two unsanctioned applications.

1. Select **Objects** > **Applications**.
2. Click the application **Name** to edit an application and select **Edit** in the Tag section.
3. Select **Sanctioned** from the **Tags** drop-down.

   You must use the predefined **Sanctioned** tag (Sanctioned). If you use any other tag to indicate that you sanctioned an application, the firewall will fail to recognize the tag and the report will be inaccurate.



4. Click **OK** and **Close** to exit all open dialogs.

https://t.me/learningnets

STEP 2 | Configure the SaaS Application Usage report.

1. Select **Monitor** > **PDF Reports** > **SaaS Application Usage**.
2. Click **Add**, enter a **Name**, and select a **Time Period** for the report (default is **Last 7 Days**).

   *By default, the report includes detailed information on the top SaaS and non-SaaS application subcategories, which can make the report large by page count and file size. Clear the **Include detailed application category information in report** check box if you want to reduce the file size and restrict the page count to 10 pages.*
3. Select whether you want the report to **Include logs from**:

   *In PAN-OS 10.0.2 and later releases, reports generated from logs in the Cortex Data Lake only support including logs from the **Selected Zone**.*

   - **All User Groups and Zones**—The report includes data on all security zones and user groups available in the logs.

     If you want to include specific user groups in the report, select **Include user group information in the report** and click the **manage groups** link to select the groups you want to include. You must add between one and up to a maximum of 25 user groups, so that the firewall or Panorama can filter the logs for the selected user groups. If you do select the groups to include, the report will aggregate all user groups in to one group called Others.
   - **Selected Zone**—The report filters data for the specified security zone, and includes data on that zone only.

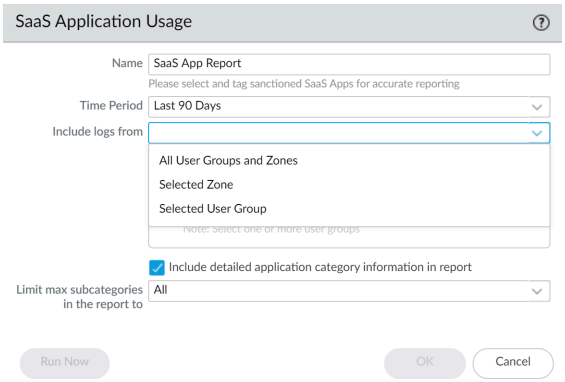
     If you want to include specific user groups in the report, select **Include user group information in the report** and click the **manage groups for selected zone** link to select the user groups within this zone that you want to include in the report. You must add between one and up to a maximum of 25 user groups, so that the firewall or Panorama can filter the logs for the selected user groups within the security zone. If

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

you do select the groups to include, the report will aggregate all user groups in to one group called Others.

- **Selected User Group**—The report filters data for the specified user group only, and includes SaaS application usage information for the selected user group only.



4. Select whether you want to include all the application subcategories in the report (the default) or **Limit the max subcategories in the report** to the top 10, 15, 20 or 25 categories (default is all subcategories).
5. Click **Run Now** to generate the report on-demand for the last 7-day and the last 30-day time period. Make sure that the pop-up blocker is disabled on your browser because the report opens in a new tab.
6. Click **OK** to save your changes.

STEP 3 | Schedule Reports for Email Delivery.

The last 90-days report must be scheduled for email delivery.

On the PA-220R and the PA-800 Series firewalls, the SaaS Application Usage report is not sent as a PDF attachment in the email. Instead, the email includes a link that you must click to open the report in a web browser.

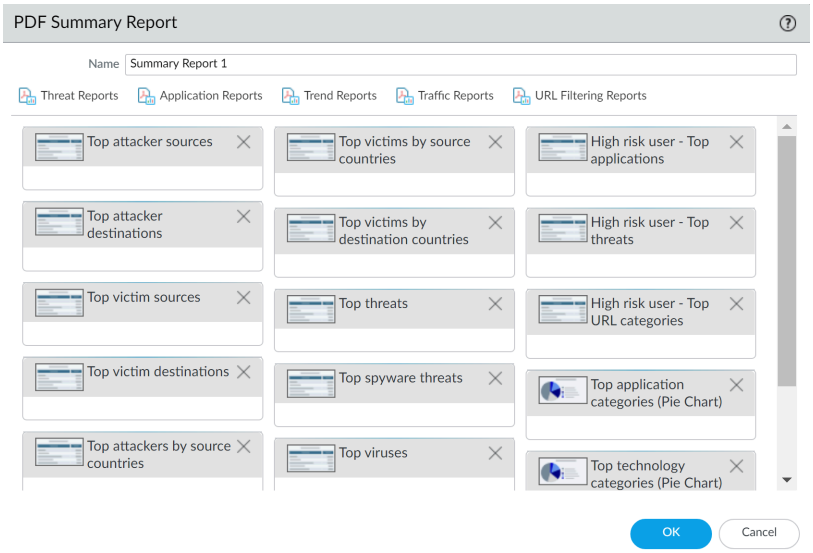
Manage PDF Summary Reports

PDF summary reports contain information compiled from existing reports, based on data for the top 5 in each category (instead of top 50). They also contain trend charts that are not available in other reports.

https://t.me/learningnets

STEP 1 | Set up a **PDF Summary Report**.

1. Select **Monitor** > **PDF Reports** > **Manage PDF Summary**.
2. Click **Add** and then enter a **Name** for the report.
3. Use the drop-down for each report group and select one or more of the elements to design the PDF Summary Report. You can include a maximum of 18 report elements.



*Selecting **Top Threats** is displayed as `top-attacks` in the Predefined Widgets column for the PDF Summary Report.*

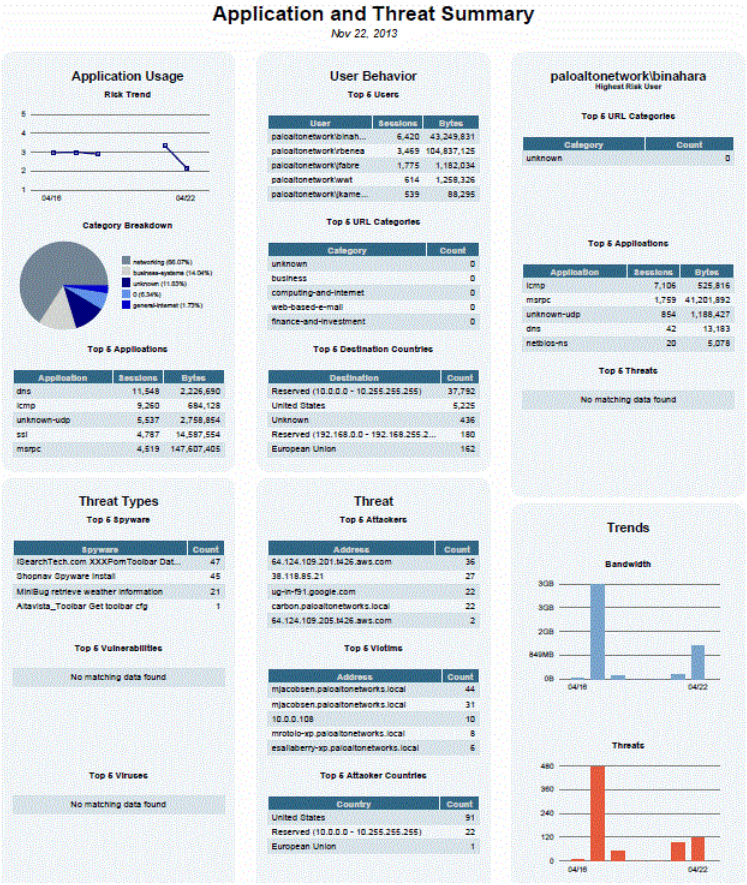
- To remove an element from the report, click the **x** icon or clear the selection from the drop-down for the appropriate report group.
- To rearrange the reports, drag and drop the element icons to another area of the report.

4. Click **OK** to save the report.
5. **Commit** the changes.

https://t.me/learningnets

STEP 2 | View the report.

To download and view the PDF Summary Report, see View Reports.



The following summary sections refer to the following PDF Summary Report elements:

- ***Top 5 Attacks**—Refers to the **Top threats** element.*
- ***Top 5 Threats**—Refers to the **High risk user - Top threats** element.*
- *Top Threats report—Refers to the full list of threats from the **Top threats** element.*

Generate User/Group Activity Reports

User/Group Activity reports summarize the web activity of individual users or user groups. Both reports include the same information except for the **Browsing Summary by URL Category** and **Browse time calculations**, which only the User Activity report includes.

You must configure User-ID on the firewall to access the list of users and user groups.

https://t.me/learningnets

STEP 1 | Configure the browse times and number of logs for User/Group Activity reports.

Required only if you want to change the default values.

1. Select **Device** > **Setup** > **Management**, edit the Logging and Reporting Settings, and select the **Log Export and Reporting** tab.
2. For the **Max Rows in User Activity Report**, enter the maximum number of rows that the detailed user activity report supports (range is 1-1048576, default is 5000). This determines the number of logs that the report analyzes.
3. Enter the **Average Browse Time** in seconds that you estimate users should take to browse a web page (range is 0-300, default is 60). Any request made after the average browse time elapses is considered a new browsing activity. The calculation uses Log Only the Page a User Visits (logged in the URL Filtering logs) as the basis and ignores any new web pages that are loaded between the time of the first request (start time) and the average browse time. For example, if you set the **Average Browse Time** to two minutes and a user opens a web page and views that page for five minutes, the browse time for that page will still be two minutes. This is done because the firewall can't determine how long a user views a given page. The average browse time calculation ignores sites categorized as web advertisements and content delivery networks.
4. For the **Page Load Threshold**, enter the estimated time in seconds for page elements to load on the page (default is 20). Any requests that occur between the first page load and the page load threshold are assumed to be elements of the page. Any requests that occur outside of the page load threshold are assumed to be the user clicking a link within the page.
5. Click **OK** to save your changes.

STEP 2 | Generate the User/Group Activity report.

1. Select **Monitor** > **PDF Reports** > **User Activity Report**.
2. Click **Add** and then enter a **Name** for the report.
3. Create the report:
   - User Activity Report—Select **User** and enter the **Username** or **IP address** (IPv4 or IPv6) of the user.
   - Group Activity Report—Select **Group** and select the **Group Name** of the user group.
4. Select the **Time Period** for the report.
5. (Optional) Select the **Include Detailed Browsing** check box (default is cleared) to include detailed URL logs in the report.

   The detailed browsing information can include a large volume of logs (thousands of logs) for the selected user or user group and can make the report very large.
6. To run the report on demand, click **Run Now**.
7. To save the report configuration, click **OK**. You can't save the output of User/Group Activity reports on the firewall. To schedule the report for email delivery, see Schedule Reports for Email Delivery.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

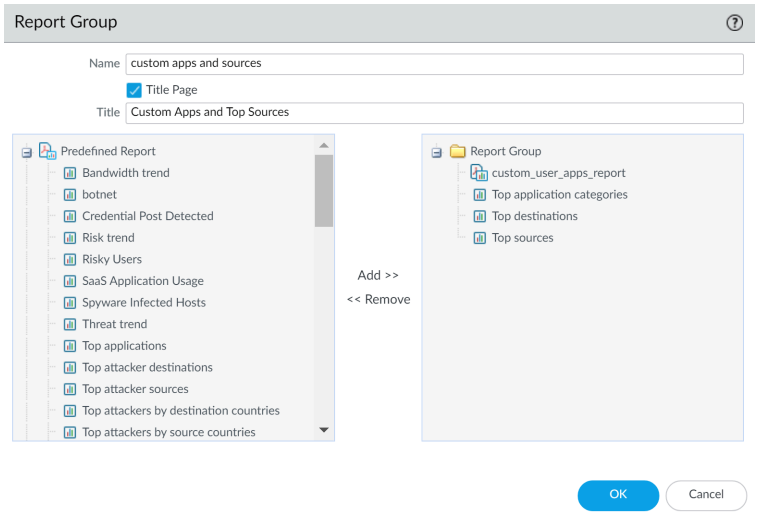
Manage Report Groups

Report groups allow you to create sets of reports that the system can compile and send as a single aggregate PDF report with an optional title page and all the constituent reports included.

Set up report groups.

You must set up a **Report Group** to email report(s).

1. Create an Email server profile.
2. Define the **Report Group**. A report group can compile predefined reports, PDF Summary reports, custom reports, and Log View report into a single PDF.
   1. Select **Monitor** > **Report Group**.
   2. Click **Add** and then enter a **Name** for the report group.
   3. (Optional) Select **Title Page** and add a **Title** for the PDF output.
   4. Select reports from the left column and click **Add** to move each report to the report group on the right.



   The **Log View** report is a report type that is automatically created each time you create a custom report and uses the same name as the custom report. This report will show the logs that were used to build the contents of the custom report.

   To include the log view data, when creating a report group, add your custom report under the **Custom Reports** list and then add the log view report by selecting the matching report name from the **Log View** list. The report will include the custom report data and the log data that was used to create the custom report.
   5. Click **OK** to save the settings.
   6. To use the report group, see Schedule Reports for Email Delivery.

Schedule Reports for Email Delivery

Reports can be scheduled for daily delivery or delivered weekly on a specified day. Scheduled reports are executed starting at 2:00 AM, and email delivery starts after all scheduled reports have been generated.

https://t.me/learningnets

STEP 1 | Select **Monitor** > **PDF Reports** > **Email Scheduler** and click **Add**.

STEP 2 | Enter a **Name** to identify the schedule.

STEP 3 | Select the **Report Group** for email delivery. To set up a report group; see Manage Report Groups.

STEP 4 | For the **Email Profile**, select an Email server profile to use for delivering the reports, or click the **Email Profile** link to Create an Email server profile.

STEP 5 | Select the frequency at which to generate and send the report in **Recurrence**.

STEP 6 | The **Override Email Addresses** field allows you to send this report exclusively to the specified recipients. When you add recipients to the field, the firewall does not send the report to the recipients configured in the Email server profile. Use this option for those occasions when the report is for the attention of someone other than the administrators or recipients defined in the Email server profile.

STEP 7 | Click **OK** and **Commit**.

Manage Report Storage Capacity

By default, firewalls contain 200MB of dedicated storage for reports generated by the firewall. In some instances, especially for PA-7000 series and PA-5200 series firewalls, you may need to increase the capacity of available report storage space in order to successfully generate new reports.

STEP 1 | Access the firewall CLI.

STEP 2 | Confirm the current report storage capacity of the firewall:

The command output displays the report storage size in bytes. For this procedure, the firewall has the default 200MB report storage capacity.

```
admin@ISP-CONDOR-B(active)> request report-storage-size show

209715200
```

STEP 3 | Verify you have sufficient storage across the firewall to allocate toward expanding the report storage capacity:

```
admin> show system disk-space
```

```
admin@ISP-CONDOR-B(active)> show system disk-space

Filesystem      Size  Used Avail Use% Mounted on
/dev/root        12G  8.9G  2.0G  83% /
none            7.9G   52K  7.9G   1% /dev
/dev/sda5        16G  8.5G  5.9G  59% /opt/pancfg
/dev/sda6        12G  5.8G  5.0G  54% /opt/panrepo
tmpfs           7.9G  247M  7.6G   4% /dev/shm
/dev/sda8        22G  8.7G   12G  43% /opt/panlogs
tmpfs            12M     0   12M   0% /opt/pancfg/mgmt/lcaas/ssl/private
```

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 4 | Increase the report storage capacity as needed:

For example, we are increasing the report storage size to 1GB.

```
admin> request report-storage-size set size <0-4>
```

```
admin@ISP-CONDOR-B(active)> request report-storage-size set size 1

cfg.report-storage-size-gb: 1
```

STEP 5 | Verify that the report storage capacity is increased to the amount set in the previous step:

```
admin> request report-storage-size show
```

```
admin@ISP-CONDOR-B(active)> request report-storage-size show

1073741824
```

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

View Policy Rule Usage

View the number of times a Security, NAT, QoS, policy-based forwarding (PBF), Decryption, Tunnel Inspection, Application Override, Authentication, or DoS protection rule matches traffic to help keep your firewall policies up to date as your environment and security needs change. To prevent attackers from exploiting over-provisioned access, such as when a server is decommissioned or when you no longer need temporary access to a service, use the policy rule hit count data to identify and remove unused rules.

Policy rule usage data enables you to validate rule additions and rule changes and to monitor the time frame when a rule was used. For example, when you migrate port-based rules to app-based rules, you create an app-based rule above the port-based rule and check for any traffic that matches the port-based rule. After migration, the hit count data helps you determine whether it is safe to remove the port-based rule by confirming whether traffic is matching the app-based rule instead of the port-based rule. The policy rule hit count helps you determine whether a rule is effective for access enforcement.

You can reset the rule hit count data to validate an existing rule or to gauge rule usage within a specified period of time. Policy rule hit count data is not stored on the firewall or Panorama so that data is no longer available after you reset (clear) the hit count.

After filtering your policy rulebase, administrators can take action to delete, disable, enable, and tag policy rules directly from the policy optimizer. For example, you can filter for unused rules and then tag them for review to determine whether they can be safely deleted or kept in the rulebase. By enabling administrators to take action directly from the policy optimizer, you reduce the management overhead required to further assist in simplifying your rule lifecycle management and ensure that your firewalls are not over-provisioned.

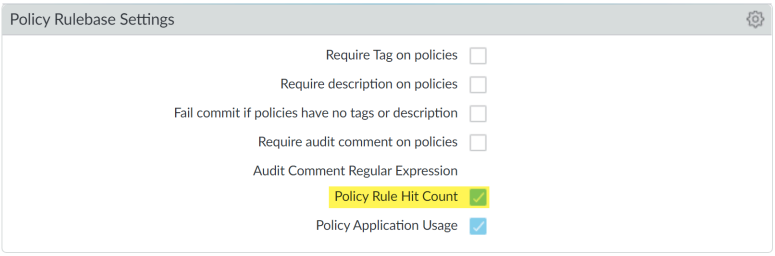
The rule hit count data is not synchronized across firewalls in a high availability (HA) deployment so you need to log in to each firewall to view the policy rule hit count data for each firewall or use Panorama to view information on the HA firewall peers.

Policy rule usage data is also useful when using Security Policy Rule Optimization to determine which rules to migrate or clean up first.

STEP 1 | Launch the Web Interface.

STEP 2 | Verify that **Policy Rule Hit Count** is enabled.

1. Navigate to Policy Rulebase Settings (**Device** > **Setup** > **Management**).
2. Verify that **Policy Rule Hit Count** is enabled.



STEP 3 | Select **Policies**.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 4 | View the policy rule usage for each policy rule:

- **Hit Count**—The number of times traffic matched the criteria you defined in the policy rule. Persists through reboot, dataplane restarts, and upgrades unless you manually reset or rename the rule.
- **Last Hit**—The most recent timestamp for when traffic matched the rule.
- **First Hit**—The first instance when traffic was matched to this rule.
- **Modified**—The date and time the policy rule was last modified.
- **Created**—The date and time the policy rule was created.

If the rule was created when Panorama was running PAN-OS 8.1 and the Policy Rule Hit Count setting is enabled, the First Hit date and time is used as the Created date and time on upgrade to PAN-OS 9.0. If the rule was created in PAN-OS 8.1 when the Policy Rule Hit Count setting was disabled or if the rule was created when Panorama was running PAN-OS 8.0 or an earlier release, the Created date for the rule will be the date and time you successfully upgraded Panorama to PAN-OS 9.0

		Source			Rule Usage				
NAME	T...	Z...	A...	U...	HIT COUNT	LAST HIT	FIRST HIT	MODIFIED	CREATED
video	n...	a...	a...	a...	2424328	2020-09-22 11:33:00	2019-07-30 10:12:57	2020-07-27 13:27:16	2019-07-30 09:50
Video Streaming	n...	a...	a...	a...	14337228	2020-09-22 16:26:58	2019-07-30 10:12:57	2020-07-27 13:27:16	2019-07-30 09:50
cavenger	n...	a...	a...	a...	321760616	2020-09-22 16:27:10	2019-07-30 10:12:57	2020-07-27 13:27:16	2019-07-30 09:50
Web Traffic	n...	a...	a...	a...	1509584361	2020-09-22 16:27:10	2019-07-30 10:12:02	2020-07-27 13:27:16	2019-07-30 09:50
iperf	n...	a...	a...	a...	5	2019-10-15 14:54:31	2019-10-11 13:08:28	2020-07-27 13:27:16	2019-07-30 09:50

STEP 5 | In the Policy Optimizer dialog, view the **Rule Usage** filter.

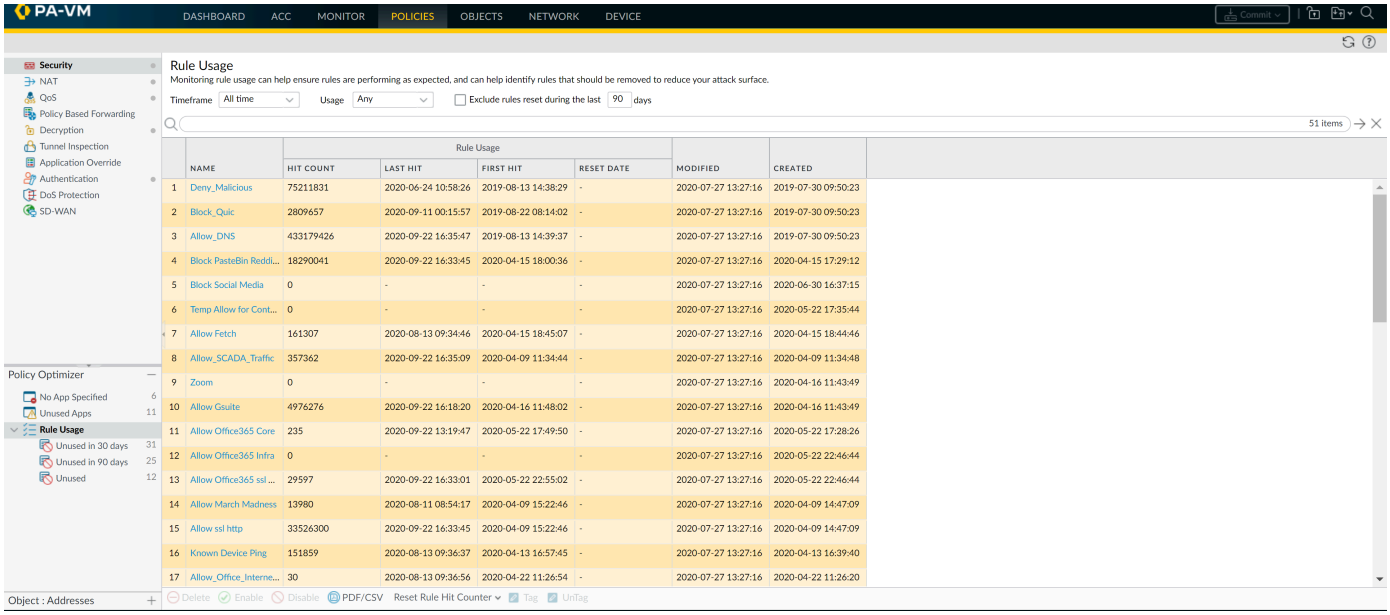
STEP 6 | Filter rules in the selected rulebase.

Use the rule usage filter to evaluate the rule usage within a specified period of time. For example, filter the selected rulebase for Unused rules within the last 30 days. You can also evaluate rule usage with other rule attributes, such as the Created and Modified dates, which enables you to filter for the correct set of rules to review. You can use this data to help manage your rule lifecycle and to determine if a rule needs to be removed to reduce your network attack surface.

1. Select the **Timeframe** you want to filter on or specify a **Custom** time frame.
2. Select the rule **Usage** on which to filter.
3. (Optional) If you have reset the rule usage data for any rules, check for **Exclude rules reset during the last <*number of days*> days** and decide when to exclude a rule based

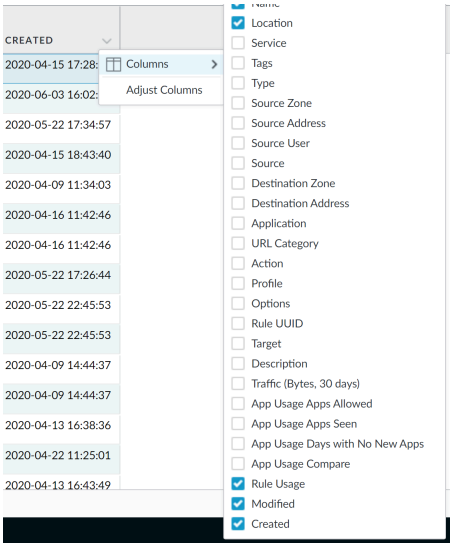
https://t.me/learningnets

on the number of days you specify since the rule was reset. Only rules that were reset before your specified number of days are included in the filtered results.



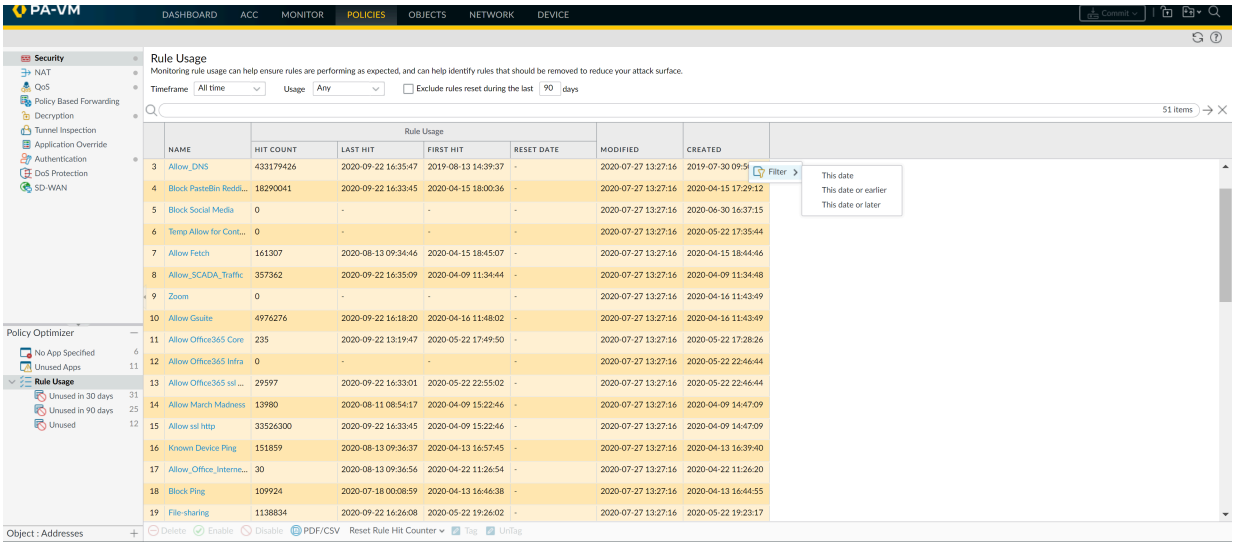
4. (Optional) Specify search filters based on rule data

   1. Hover your cursor over the column header and **Columns**.
   2. Add any additional columns you want to display or use for filter.



   3. Hover your cursor over the column data that you would like to filter on **Filter**. For data that contain dates, select whether to filter using **This date**, **This date or earlier**, or **This date or later**.
   4. **Apply Filter** (→).

https://t.me/learningnets



STEP 7 | Take action on one or more unused policy rules.

1. Select one or more unused policy rules.
2. Perform one of the following actions:
   - **Delete**—Delete one or more selected policy rules.
   - **Enable**—Enable one or more selected policy rules when disabled.
   - **Disable**—Disable one or more selected policy rules.
   - **Tag**—Apply one or more group tags to one or more selected policy rules. The group tag must already exist in order to tag policy rule.
   - **Untag**—Remove one or more group tags from one or more selected policy rules.
3. **Commit** your changes.

https://t.me/learningnets

Use External Services for Monitoring

Using an external service to monitor the firewall enables you to receive alerts for important events, archive monitored information on systems with dedicated long-term storage, and integrate with third-party security monitoring tools. The following are some common scenarios for using external services:

- For immediate notification about important system events or threats, you can Monitor Statistics Using SNMP, Forward Traps to an SNMP Manager, or Configure Email Alerts.
- To send an HTTP-based API request directly to any third-party service that exposes an API to automate a workflow or an action. You can, for example, forward logs that match a defined criteria to create an incidence ticket on ServiceNow instead of relying on an external system to convert syslog messages or SNMP traps to an HTTP request. You can modify the URL, HTTP header, parameters, and the payload in the HTTP request to trigger an action based on the attributes in a firewall log. See Forward Logs to an HTTP(S) Destination.
- For long-term log storage and centralized firewall monitoring, you can Configure Syslog Monitoring to send log data to a syslog server. This enables integration with third-party security monitoring tools such as Splunk or ArcSight.
- For monitoring statistics on the IP traffic that traverses firewall interfaces, you can Configure NetFlow Exports to view the statistics in a NetFlow collector.

You can Configure Log Forwarding from the firewalls directly to external services or from the firewalls to Panorama and then configure Panorama to forwardlogs to the servers. Refer to Log Forwarding Options for the factors to consider when deciding where to forward logs.

You can't aggregate NetFlow records on Panorama; you must send them directly from the firewalls to a NetFlow collector.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Configure Log Forwarding

In an environment where you use multiple firewalls to control and analyze network traffic, any single firewall can display logs and reports only for the traffic it monitors. Because logging in to multiple firewalls can make monitoring a cumbersome task, you can more efficiently achieve global visibility into network activity by forwarding the logs from all firewalls to Panorama or external services. If you Use External Services for Monitoring, the firewall automatically converts the logs to the necessary format: syslog messages, SNMP traps, email notifications, or as an HTTP payload to send the log details to an HTTP(S) server. In cases where some teams in your organization can achieve greater efficiency by monitoring only the logs that are relevant to their operations, you can create forwarding filters based on any log attributes (such as threat type or source user). For example, a security operations analyst who investigates malware attacks might be interested only in Threat logs with the type attribute set to wildfire-virus.

You can forward logs from the firewalls directly to external services or from the firewalls to Panorama and then configure Panorama to forward logs to the servers. Refer to Log Forwarding Options for the factors to consider when deciding where to forward logs.

You can use Secure Copy (SCP) commands from the CLI to export the entire log database to an SCP server and import it to another firewall. Because the log database is too large for an export or import to be practical on the PA-7000 Series firewall, it does not support these options. You can also use the web interface on all platforms to View and Manage Reports, but only on a per log type basis, not for the entire log database.

STEP 1 | Configure a server profile for each external service that will receive log information.

You can use separate profiles to send different sets of logs, filtered by log attributes, to a different server. To increase availability, define multiple servers in a single profile.

Configure one or more of the following server profiles:

- (Required for SMTP over TLS) If you have not already done so, create a certificate profile for the email server.
- 2 To enable the SNMP manager (trap server) to interpret firewall traps, you must load the Palo Alto Networks Supported MIBs into the SNMP manager and, if necessary, compile them. For details, refer to your SNMP management software documentation.
- If the syslog server requires client authentication, you must also 5
- Configure an HTTP server profile (see Forward Logs to an HTTP/S Destination).

STEP 2 | Create a Log Forwarding profile.

The profile defines the destinations for Traffic, Threat, WildFire Submission, URL Filtering, Data Filtering, Tunnel and Authentication logs.

1. Select **Objects** > **Log Forwarding** and **Add** a profile.
2. Enter a **Name** to identify the profile.

   If you want the firewall to automatically assign the profile to new security rules and zones, enter **`default`**. If you don't want a default profile, or you want to override

https://t.me/learningnets

an existing default profile, enter a **Name** that will help you identify the profile when assigning it to security rules and zones.

> *If no log forwarding profile named **default** exists, the profile selection is set to **None** by default in new security rules (**Log Forwarding** field) and new security zones (**Log Setting** field), although you can change the selection.*

3. **Add** one or more *match list profiles*.

   The profiles specify log query filters, forwarding destinations, and automatic actions such as tagging. For each match list profile:

   1. Enter a **Name** to identify the profile.
   2. Select the **Log Type**.
   3. In the **Filter** drop-down, select **Filter Builder**. Specify the following and then **Add** each query:
      - **Connector** logic (and/or)
      - Log **Attribute**
      - **Operator** to define inclusion or exclusion logic
      - Attribute **Value** for the query to match
   4. Select **Panorama** if you want to forward logs to Log Collectors or the Panorama management server.
   5. For each type of external service that you use for monitoring (SNMP, Email, Syslog, and HTTP), **Add** one or more server profiles.
4. (Optional, GlobalProtect Only) If you are using a log forwarding profile with a security policy to automatically quarantine a device using GlobalProtect, select **Quarantine** in the **Built-in Actions** area.
5. Click **OK** to save the Log Forwarding profile.

STEP 3 | Assign the Log Forwarding profile to policy rules and network zones.

Security, Authentication, and DoS Protection rules support log forwarding. In this example, you assign the profile to a Security rule.

Perform the following steps for each rule that you want to trigger log forwarding:

1. Select **Policies** > **Security** and edit the rule.
2. Select **Actions** and select the **Log Forwarding** profile you created.
3. Set the **Profile Type** to **Profiles** or **Group**, and then select the security profiles or **Group Profile** required to trigger log generation and forwarding for:
   - Threat logs—Traffic must match any security profile assigned to the rule.
   - WildFire Submission logs—Traffic must match a WildFire Analysis profile assigned to the rule.
4. For Traffic logs, select **Log At Session Start** and/or **Log At Session End**.
5. Click **OK** to save the rule.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 4 | Configure the destinations for System, Configuration, Correlation, GlobalProtect, HIP Match, and User-ID logs.

Panorama generates Correlation logs based on the firewall logs it receives, rather than aggregating Correlation logs from firewalls.

1. Select **Device** > **Log Settings**.
2. For each log type that the firewall will forward, see Step Add one or more match list profiles.

STEP 5 | (PA-7000 Series firewalls only) Configure a log card interface to perform log forwarding.

1. Select **Network** > **Interfaces** > **Ethernet** and click **Add Interface**.
2. Select the **Slot** and **Interface Name**.
3. Set the **Interface Type** to **Log Card**.
4. Enter the **IP Address**, **Default Gateway**, and (for IPv4 only) **Netmask**.
5. Select **Advanced** and specify the **Link Speed**, **Link Duplex**, and **Link State**.

   *These fields default to **auto**, which specifies that the firewall automatically determines the values based on the connection. However, the minimum recommended **Link Speed** for any connection is **1000** (Mbps).*
6. Click **OK** to save your changes.

STEP 6 | Commit and verify your changes.

1. **Commit** your changes.
2. Verify the log destinations you configured are receiving firewall logs:
   - Panorama—If the firewall forwards logs to a Panorama virtual appliance in Panorama mode or to an M-Series appliance, you must configure a Collector Group before Panorama will receive the logs. You can then verify log forwarding.
   - Email server—Verify that the specified recipients are receiving logs as email notifications.
   - Syslog server—Refer to your syslog server documentation to verify it's receiving logs as syslog messages.
   - SNMP manager—Use an SNMP Manager to Explore MIBs and Objects to verify it's receiving logs as SNMP traps.
   - HTTP server—Forward Logs to an HTTP/S Destination.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Configure Email Alerts

You can configure email alerts for System, Config, HIP Match, Correlation, Threat, WildFire Submission, and Traffic logs. You can use separate profiles to send email notifications for each log type to a different server. To increase availability, define multiple servers (up to four) in a single profile.

As a best practice, configure transport layer security (TLS) to require the firewall to authenticate with the email server before the firewall relays email to the server. This helps prevent malicious activity, such as Simple Mail Transfer Protocol (SMTP) relay, which can be used to send spam or malware, and email spoofing, which can be used for phishing attacks.

STEP 1 | (Required for SMTP over TLS) If you have not already done so, create a certificate profile for the email server.

STEP 2 | Select **Device** > **Server Profiles** > **Email**.

STEP 3 | **Add** an email server profile and enter a **Name**.

STEP 4 | From the read-only window that appears, **Add** the email server and enter a **Name.**

STEP 5 | If the firewall has more than one virtual system (vsys), select the **Location** (vsys or **Shared**) where this profile is available.

STEP 6 | (Optional) Enter an **Email Display Name** to specify the name to display in the From field of the email.

STEP 7 | Enter the email address **From** which the firewall sends emails.

STEP 8 | Enter the email address **To** which the firewall sends emails.

STEP 9 | (Optional) If you want to send emails to a second account, enter the address of the **Additional Recipient**. You can add only one additional recipient. For multiple recipients, add the email address of a distribution list.

STEP 10 | Enter the IP address or hostname of the **Email Gateway** to use for sending emails.

STEP 11 | Select the **Type** of protocol to use to connect to the email server:

- **Unauthenticated SMTP**—Use SMTP to connect to the email server without authentication. The default **Port** is 25, but you can optionally specify a different port. This protocol does not provide the same security as SMTP over TLS, but if you select this protocol, skip the next step.
- **SMTP over TLS**—(Recommended) Use TLS to require authentication to connect to the email server. Continue to the next step to configure the TLS authentication.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 12 | (SMTP over TLS only) Configure the firewall to use TLS authentication to connect to the email server.

1. (Optional) Specify the **Port** to use to connect to the email server (default is 587).
2. **TLS Version**—Specify the TLS version (**1.1** or **1.2**).

   Palo Alto Networks strongly recommends using the latest TLS version.

3. Select the **Authentication Method** for the firewall and the email server:
   - **Auto**—Allow the firewall and the email server to determine the authentication method.
   - **Login**—Use Base64 encoding for the username and password and transmit them separately.
   - **Plain**—Use Base64 encoding for the username and password and transmit them together.
4. Select a **Certificate Profile** to authenticate with the email server.
5. Enter the **Username** and **Password** of the account that sends the emails, then **Confirm Password**.
6. (Optional) To confirm that the firewall can successfully authenticate with the email server, you can **Test Connection**.

STEP 13 | Click **OK** to save the Email server profile.

STEP 14 | (Optional) Select the **Custom Log Format** tab and customize the format of the email messages. For details on how to create custom formats for the various log types, refer to the Common Event Format Configuration Guide.

STEP 15 | Configure email alerts for Traffic, Threat, and WildFire Submission logs.

1. See Create a Log Forwarding profile.
   1. Select **Objects** > **Log Forwarding**, click **Add**, and enter a **Name** to identify the profile.
   2. For each log type and each severity level or WildFire verdict, select the Email server profile and click **OK**.
2. See Assign the Log Forwarding profile to policy rules and network zones.

STEP 16 | Configure email alerts for System, Config, HIP Match, and Correlation logs.

1. Select **Device** > **Log Settings**.
2. For System and Correlation logs, click each Severity level, select the **Email** server profile, and click **OK**.
3. For Config and HIP Match logs, edit the section, select the **Email** server profile, and click **OK**.
4. Click **Commit**.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Use Syslog for Monitoring

Syslog is a standard log transport mechanism that enables the aggregation of log data from different network devices—such as routers, firewalls, printers—from different vendors into a central repository for archiving, analysis, and reporting. Palo Alto Networks firewalls can forward every type of log they generate to an external syslog server. You can use TCP or TLS (TLSv1.2 only) for reliable and secure log forwarding, or UDP for non-secure forwarding.

- Configure Syslog Monitoring
- Syslog Field Descriptions

Configure Syslog Monitoring

To Use Syslog for Monitoring a Palo Alto Networks firewall, create a Syslog server profile and assign it to the log settings for each log type. Optionally, you can configure the header format used in syslog messages and enable client authentication for syslog over TLSv1.2.



For CEF-formated syslog events collection, you must edit the default syslog configuration. The default syslog monitoring configuration is not supported for CEF syslog events collection.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 1 | Configure a Syslog server profile.

You can use separate profiles to send syslogs for each log type to a different server. To increase availability, define multiple servers (up to four) in a single profile.

1. Select **Device** > **Server Profiles** > **Syslog**.
2. Click **Add** and enter a **Name** for the profile.
3. If the firewall has more than one virtual system (vsys), select the **Location** (vsys or **Shared**) where this profile is available.
4. For each syslog server, click **Add** and enter the information that the firewall requires to connect to it:
   - **Name**—Unique name for the server profile.
   - **Syslog Server**—IP address or fully qualified domain name (FQDN) of the syslog server.

     *If you configure an FQDN and use **UDP** transport, if the firewall cannot resolve the FQDN, the firewall uses the existing IP address resolution for the FQDN as the **Syslog Server** address.*
   - **Transport**—Select **TCP**, **UDP**, or **SSL** (TLS) as the protocol for communicating with the syslog server. For **SSL**, the firewall supports only TLSv1.2.
   - **Port**—The port number on which to send syslog messages (default is UDP on port 514); you must use the same port number on the firewall and the syslog server.
   - **Format**—Select the syslog message format to use: **BSD** (the default) or **IETF**. Traditionally, **BSD** format is over UDP and **IETF** format is over TCP or SSL/TLS.
   - **Facility**—Select a syslog standard value (default is **LOG_USER**) to calculate the priority (PRI) field in your syslog server implementation. Select the value that maps to how you use the PRI field to manage your syslog messages.
5. (**Optional**) To customize the format of the syslog messages that the firewall sends, select the **Custom Log Format** tab. For details on how to create custom formats for the various log types, refer to the Common Event Format Configuration Guide.
6. Click **OK** to save the server profile.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 2 | Configure syslog forwarding for Traffic, Threat, and WildFire Submission logs.

1. Configure the firewall to forward logs. For more information, see Step Create a Log Forwarding profile.
   1. Select **Objects** > **Log Forwarding**, click **Add**, and enter a **Name** to identify the profile.
   2. For each log type and each severity level or WildFire verdict, select the **Syslog** server profile and click **OK**.
2. Assign the log forwarding profile to a security policy to trigger log generation and forwarding. For more information, See Step Assign the Log Forwarding profile to policy rules and network zones.
   1. Select **Policies** > **Security** and select a policy rule.
   2. Select the **Actions** tab and select the **Log Forwarding** profile you created.
   3. In the **Profile Type** drop-down, select **Profiles** or **Groups**, and then select the security profiles or **Group Profiles** required to trigger log generation and forwarding.
   4. For Traffic logs, select one or both of the **Log at Session Start** and **Log At Session End** check boxes, and click **OK**.

   For detailed information about configuring a log forwarding profile and assigning the profile to a policy rule, see Configure Log Forwarding.

STEP 3 | Configure syslog forwarding for System, Config, HIP Match, and Correlation logs.

1. Select **Device** > **Log Settings**.
2. For System and Correlation logs, click each Severity level, select the **Syslog** server profile, and click **OK**.
3. For Config, HIP Match, and Correlation logs, edit the section, select the **Syslog** server profile, and click **OK**.

STEP 4 | (Optional) Configure the header format of syslog messages.

The log data includes the unique identifier of the firewall that generated the log. Choosing the header format provides more flexibility in filtering and reporting on the log data for some Security Information and Event Management (SIEM) servers.

This is a global setting and applies to all Syslog server profiles configured on the firewall.

1. Select **Device** > **Setup** > **Management** and edit the Logging and Reporting Settings.
2. Select the **Log Export and Reporting** tab and select the Syslog HOSTNAME Format:
   - **FQDN** (default)—Concatenates the hostname and domain name defined on the sending firewall.
   - **hostname**—Uses the hostname defined on the sending firewall.
   - **ipv4-address**—Uses the IPv4 address of the firewall interface used to send logs. By default, this is the MGT interface.
   - **ipv6-address**—Uses the IPv6 address of the firewall interface used to send logs. By default, this is the MGT interface.
   - **none**—Leaves the hostname field unconfigured on the firewall. There is no identifier for the firewall that sent the logs.
3. Click **OK** to save your changes.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 5 | Create a certificate to secure syslog communication over TLSv1.2.

Required only if the syslog server uses client authentication. The syslog server uses the certificate to verify that the firewall is authorized to communicate with the syslog server.

Ensure the following conditions are met:

- The private key must be available on the sending firewall; the keys can't reside on a Hardware Security Module (HSM).
- The subject and the issuer for the certificate must not be identical.
- The syslog server and the sending firewall must have certificates that the same trusted certificate authority (CA) signed. Alternatively, you can generate a self-signed certificate on the firewall, export the certificate from the firewall, and import it in to the syslog server.
- The connection to a Syslog server over TLS is validated using the Online Certificate Status Protocol (OCSP) or using Certificate Revocation Lists (CRL) so long as each certificate in the trust chain specifies one or both of these extensions. However, you cannot bypass OCSP or CRL failures so you must ensure that the certificate chain is valid and that you can verify each certificate using OCSP or CRL.

1. Select **Device** > **Certificate Management** > **Certificates** > **Device Certificates** and click **Generate**.
2. Enter a **Name** for the certificate.
3. In the **Common Name** field, enter the IP address of the firewall sending logs to the syslog server.
4. In **Signed by**, select the trusted CA or the self-signed CA that the syslog server and the sending firewall both trust.

   The certificate can't be a **Certificate Authority** nor an **External Authority** (certificate signing request [CSR]).
5. Click **Generate**. The firewall generates the certificate and key pair.
6. Click the certificate Name to edit it, select the **Certificate for Secure Syslog** check box, and click **OK**.

STEP 6 | Commit your changes and review the logs on the syslog server.

1. Click **Commit**.
2. To review the logs, refer to the documentation of your syslog management software. You can also review the Syslog Field Descriptions.

STEP 7 | (Optional) Configure the firewall to terminate the connection to the syslog server upon FQDN refresh.

When you configure a syslog server profile using a FQDN, the firewall maintains its connection to the syslog server by default in the event of an FQDN name change.

For example, you have replaced an existing syslog server with a new syslog server that uses a different FQDN name. If you want the firewall to connect to the new syslog server using a new FQDN name, you can configure the firewall to automatically terminate its connection to

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

the old syslog server and establish a connection to the new syslog server using the new FQDN name.

1. Log in to the firewall CLI.
2. Configure the firewall to terminate the connection to the syslog server upon FQDN refresh.

```
admin> set syslogng fqdn-refresh yes
```

Syslog Field Descriptions

The following topics list the standard fields of each log type that Palo Alto Networks firewalls can forward to an external server, as well as the severity levels, custom formats, and escape sequences. To facilitate parsing, the delimiter is a comma: each field is a comma-separated value (CSV) string. The FUTURE_USE tag applies to fields that the firewalls do not currently implement.

WildFire Submissions logs are a subtype of Threat log and use the same syslog format.

- Traffic Log Fields
- Threat Log Fields
- HIP Match Log Fields
- GlobalProtect Log Fields
- IP-Tag Log Fields
- User-ID Log Fields
- Decryption Log Fields
- Tunnel Inspection Log Fields
- SCTP Log Fields
- Config Log Fields
- Authentication Log Fields
- System Log Fields
- Correlated Events Log Fields
- GTP Log Fields
- Custom Log/Event Format
- Escape Sequences

Traffic Log Fields

Format: FUTURE_USE, Receive Time, Serial Number, Type, Threat/Content Type, FUTURE_USE, Generated Time, Source Address, Destination Address, NAT Source IP, NAT Destination IP, Rule Name, Source User, Destination User, Application, Virtual System, Source Zone, Destination Zone, Inbound Interface, Outbound Interface, Log Action, FUTURE_USE, Session ID, Repeat Count, Source Port, Destination Port, NAT Source Port, NAT Destination Port, Flags, Protocol, Action, Bytes, Bytes Sent, Bytes Received, Packets, Start Time, Elapsed Time, Category, FUTURE_USE, Sequence Number, Action Flags, Source Country, Destination Country, FUTURE_USE, Packets

https://t.me/learningnets

Sent, Packets Received, Session End Reason, Device Group Hierarchy Level 1, Device Group Hierarchy Level 2, Device Group Hierarchy Level 3, Device Group Hierarchy Level 4, Virtual System Name, Device Name, Action Source, Source VM UUID, Destination VM UUID, Tunnel ID/IMSI, Monitor Tag/IMEI, Parent Session ID, Parent Start Time, Tunnel Type, SCTP Association ID, SCTP Chunks, SCTP Chunks Sent, SCTP Chunks Received, Rule UUID, HTTP/2 Connection, App Flap Count, Policy ID, Link Switches, SD-WAN Cluster, SD-WAN Device Type, SD-WAN Cluster Type, SD-WAN Site, Dynamic User Group Name, XFF Address, Source Device Category, Source Device Profile, Source Device Model, Source Device Vendor, Source Device OS Family, Source Device OS Version, Source Hostname, Source Mac Address, Destination Device Category, Destination Device Profile, Destination Device Model, Destination Device Vendor, Destination Device OS Family, Destination Device OS Version, Destination Hostname, Destination Mac Address, Container ID, POD Namespace, POD Name, Source External Dynamic List, Destination External Dynamic List, Host ID, Serial Number, Source Dynamic Address Group, Destination Dynamic Address Group, Session Owner, High Resolution Timestamp, A Slice Service Type, A Slice Differentiator, Application Subcategory, Application Category, Application Technology, Application Risk, Application Characteristic, Application Container, Application SaaS, Application Sanctioned State, Offloaded

Field Name	Description
Receive Time (receive_time or cef-formatted-receive_time)	Time the log was received at the management plane.
Serial Number (serial)	Serial number of the firewall that generated the log.
Type (type)	Specifies the type of log; value is TRAFFIC.
Threat/Content Type (subtype)	Subtype of traffic log; values are start, end, drop, and deny • Start—session started • End—session ended • Drop—session dropped before the application is identified and there is no rule that allows the session. • Deny—session dropped after the application is identified and there is a rule to block or no rule that allows the session.
Generated Time (time_generated or cef-formatted-time_generated)	Time the log was generated on the dataplane.
Source Address (src)	Original session source IP address.
Destination Address (dst)	Original session destination IP address.
NAT Source IP (natsrc)	If Source NAT performed, the post-NAT Source IP address.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Field Name	Description
NAT Destination IP (natdst)	If Destination NAT performed, the post-NAT Destination IP address.
Rule Name (rule)	Name of the rule that the session matched.
Source User (srcuser)	Username of the user who initiated the session.
Destination User (dstuser)	Username of the user to which the session was destined.
Application (app)	Application associated with the session.
Virtual System (vsys)	Virtual System associated with the session.
Source Zone (from)	Zone the session was sourced from.
Destination Zone (to)	Zone the session was destined to.
Inbound Interface (inbound_if)	Interface that the session was sourced from.
Outbound Interface (outbound_if)	Interface that the session was destined to.
Log Action (logset)	Log Forwarding Profile that was applied to the session.
Session ID (sessionid)	An internal numerical identifier applied to each session.
Repeat Count (repeatcnt)	Number of sessions with same Source IP, Destination IP, Application, and Subtype seen within 5 seconds.
Source Port (sport)	Source port utilized by the session.
Destination Port (dport)	Destination port utilized by the session.
NAT Source Port (natsport)	Post-NAT source port.
NAT Destination Port (natdport)	Post-NAT destination port.
Flags (flags)	32-bit field that provides details on session; this field can be decoded by AND-ing the values with the logged value: • 0x80000000—session has a packet capture (PCAP) • 0x40000000—option is enabled to allow a client to use multiple paths to connect to a destination host • 0x20000000—file is submitted to WildFire for a verdict

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Field Name	Description
	• 0x10000000—enterprise credential submission by end user detected • 0x08000000— source for the flow is on the allow list and not subject to recon protection • 0x02000000—IPv6 session • 0x01000000—SSL session is decrypted (SSL Proxy) • 0x00800000—session is denied via URL filtering • 0x00400000—session has a NAT translation performed • 0x00200000—user information for the session was captured through Authentication Portal • 0x00100000—application traffic is on a non-standard destination port • 0x00080000 —X-Forwarded-For value from a proxy is in the source user field • 0x00040000—log corresponds to a transaction within a http proxy session (Proxy Transaction) • 0x00020000—Client to Server flow is subject to policy based forwarding • 0x00010000—Server to Client flow is subject to policy based forwarding • 0x00008000—session is a container page access (Container Page) • 0x00002000—session has a temporary match on a rule for implicit application dependency handling. Available in PAN-OS 5.0.0 and above. • 0x00000800—symmetric return is used to forward traffic for this session • 0x00000400—decrypted traffic is being sent out clear text through a mirror port • 0x00000100—payload of the outer tunnel is being inspected
IP Protocol (proto)	IP protocol associated with the session.
Action (action)	Action taken for the session; possible values are: • allow—session was allowed by policy • deny—session was denied by policy • drop—session was dropped silently • drop ICMP—session was silently dropped with an ICMP unreachable message to the host or application

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Field Name	Description
	• reset both—session was terminated and a TCP reset is sent to both the sides of the connection • reset client—session was terminated and a TCP reset is sent to the client • reset server—session was terminated and a TCP reset is sent to the server
Bytes (bytes)	Number of total bytes (transmit and receive) for the session.
Bytes Sent (bytes_sent)	Number of bytes in the client-to-server direction of the session.
Bytes Received (bytes_received)	Number of bytes in the server-to-client direction of the session.
Packets (packets)	Number of total packets (transmit and receive) for the session.
Start Time (start)	Time of session start.
Elapsed Time (elapsed)	Elapsed time of the session.
Category (category)	URL category associated with the session (if applicable).
Sequence Number (seqno)	A 64-bit log entry identifier incremented sequentially; each log type has a unique number space.
Action Flags (actionflags)	A bit field indicating if the log was forwarded to Panorama.
Source Country (srcloc)	Source country or Internal region for private addresses; maximum length is 32 bytes.
Destination Country (dstloc)	Destination country or Internal region for private addresses. Maximum length is 32 bytes.
Packets Sent (pkts_sent)	Number of client-to-server packets for the session.
Packets Received (pkts_received)	Number of server-to-client packets for the session.
Session End Reason (session_end_reason)	The reason a session terminated. If the termination had multiple causes, this field displays only the highest priority reason. The possible session end reason values are as follows, in order of priority (where the first is highest): • threat—The firewall detected a threat associated with a reset, drop, or block (IP address) action.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Field Name	Description
	• policy-deny—The session matched a security rule with a deny or drop action. • decrypt-cert-validation—The session terminated because you configured the firewall to block SSL forward proxy decryption or SSL inbound inspection when the session uses client authentication or when the session uses a server certificate with any of the following conditions: expired, untrusted issuer, unknown status, or status verification time-out. This session end reason also displays when the server certificate produces a fatal error alert of type bad_certificate, unsupported_certificate, certificate_revoked, access_denied, or no_certificate_RESERVED (SSLv3 only). • decrypt-unsupport-param—The session terminated because you configured the firewall to block SSL forward proxy decryption or SSL inbound inspection when the session uses an unsupported protocol version, cipher, or SSH algorithm. This session end reason is displays when the session produces a fatal error alert of type unsupported_extension, unexpected_message, or handshake_failure. • decrypt-error—The session terminated because you configured the firewall to block SSL forward proxy decryption or SSL inbound inspection when firewall resources or the hardware security module (HSM) were unavailable. This session end reason is also displayed when you configured the firewall to block SSL traffic that has SSL errors or that produced any fatal error alert other than those listed for the decrypt-cert-validation and decrypt-unsupport-param end reasons. • tcp-rst-from-client—The client sent a TCP reset to the server. • tcp-rst-from-server—The server sent a TCP reset to the client. • resources-unavailable—The session dropped because of a system resource limitation. For example, the session could have exceeded the number of out-of-order packets allowed per flow or the global out-of-order packet queue. • tcp-fin—Both hosts in the connection sent a TCP FIN message to close the session.
	• tcp-reuse—A session is reused and the firewall closes the previous session.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Field Name	Description
	• decoder—The decoder detects a new connection within the protocol (such as HTTP-Proxy) and ends the previous connection. • aged-out—The session aged out. • unknown—This value applies in the following situations: • Session terminations that the preceding reasons do not cover (for example, a `clear session all` command). • For logs generated in a PAN-OS release that does not support the session end reason field (releases older than PAN-OS 6.1), the value will be **unknown** after an upgrade to the current PAN-OS release or after the logs are loaded onto the firewall. • In Panorama, logs received from firewalls for which the PAN-OS version does not support session end reasons will have a value of **unknown**. • n/a—This value applies when the traffic log type is not **end**.
Device Group Hierarchy (dg_hier_level_1 to dg_hier_level_4)	A sequence of identification numbers that indicate the device group's location within a device group hierarchy. The firewall (or virtual system) generating the log includes the identification number of each ancestor in its device group hierarchy. The shared device group (level 0) is not included in this structure. If the log values are 12, 34, 45, 0, it means that the log was generated by a firewall (or virtual system) that belongs to device group 45, and its ancestors are 34, and 12. To view the device group names that correspond to the value 12, 34 or 45, use one of the following methods: API query: `/api/?type=op&cmd=<show><dg-hierarchy></dg-hierarchy></show>`
Virtual System Name (vsys_name)	The name of the virtual system associated with the session; only valid on firewalls enabled for multiple virtual systems.
Device Name (device_name)	The hostname of the firewall on which the session was logged.
Action Source (action_source)	Specifies whether the action taken to allow or block an application was defined in the application or in policy. The actions can be allow, deny, drop, reset- server, reset-client or reset-both for the session.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Field Name	Description
Source VM UUID (src_uuid)	Identifies the source universal unique identifier for a guest virtual machine in the VMware NSX environment.
Destination VM UUID (dst_uuid)	Identifies the destination universal unique identifier for a guest virtual machine in the VMware NSX environment.
Tunnel ID/IMSI (tunnelid/ imsi)	International Mobile Subscriber Identity (IMSI) is a unique number allocated to each mobile subscriber in the GSM/ UMTS/EPS system. IMSI shall consist of decimal digits (0 through 9) only and maximum number of digits allowed are 15.
Monitor Tag/IMEI (monitortag/imei)	International Mobile Equipment Identity (IMEI) is a unique 15 or 16 digit number allocated to each mobile station equipment.
Parent Session ID (parent_session_id)	ID of the session in which this session is tunneled. Applies to inner tunnel (if two levels of tunneling) or inside content (if one level of tunneling) only.
Parent Start Time (parent_start_time)	Year/month/day hours:minutes:seconds that the parent tunnel session began.
Tunnel Type (tunnel)	Type of tunnel, such as GRE or IPSec.
SCTP Association ID (assoc_id)	Number that identifies all connections for an association between two SCTP endpoints.
SCTP Chunks (chunks)	Sum of SCTP chunks sent and received for an association.
SCTP Chunks Sent (chunks_sent)	Number of SCTP chunks sent for an association.
SCTP Chunks Received (chunks_received)	Number of SCTP chunks received for an association.
Rule UUID (rule_uuid)	The UUID that permanently identifies the rule.
HTTP/2 Connection (http2_connection)	Identifies if traffic used an HTTP/2 Connection by displaying one of the following values: • Parent session ID—HTTP/2 connection • 0—SSL session
App Flap Count (link_change_count)	Number of link flaps that occurred during the session.
Policy ID (policy_id)	Name of the SD-WAN policy.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Field Name	Description
Link Switches (link_switches)	Contains up to four link flap entries, with each entry containing the link name, link tag, link type, physical interface, timestamp, bytes read, bytes written, link health, and link flap cause.
SD-WAN Cluster (sdwan_cluster)	Name of the SD-WAN cluster.
SD-WAN Device Type (sdwan_device_type)	Type of device (`hub` or `branch`).
SD-WAN Cluster Type (sdwan_cluster_type)	Type of cluster (`mesh` or `hub-spoke`).
SD-WAN Site (sdwan_site)	Name of the SD-WAN site.
Dynamic User Group Name (dynusergroup_name)	Name of the dynamic user group that contains the user who initiated the session.
XFF Address (xff_ip)	The IP address of the user who requested the web page or the IP address of the next to last device that the request traversed. If the request goes through one or more proxies, load balancers, or other upstream devices, the firewall displays the IP address of the most recent device.
Source Device Category (src_category)	The category for the device that Device-ID identifies as the source of the traffic.
Source Device Profile (src_profile)	The device profile for the device that Device-ID identifies as the source of the traffic.
Source Device Model (src_model)	The model of the device that Device-ID identifies as the source of the traffic.
Source Device Vendor (src_vendor)	The vendor of the device that Device-ID identifies as the source of the traffic.
Source Device OS Family (src_osfamily)	The operating system type for the device that Device-ID identifies as the source of the traffic.
Source Device OS Version (src_osversion)	The version of the operating system for the device that Device-ID identifies as the source of the traffic.
Source Hostname (src_host)	The hostname of the device that Device-ID identifies as the source of the traffic.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Field Name	Description
Source MAC Address (src_mac)	The MAC address for the device that Device-ID identifies as the source of the traffic.
Destination Device Category (dst_category)	The category for the device that Device-ID identifies as the destination for the traffic.
Destination Device Profile (dst_profile)	The device profile for the device that Device-ID identifies as the destination for the traffic.
Destination Device Model (dst_model)	The model of the device that Device-ID identifies as the destination for the traffic.
Destination Device Vendor (dst_vendor)	The vendor of the device that Device-ID identifies as the destination for the traffic.
Destination Device OS Family (dst_osfamily)	The operating system type for the device that Device-ID identifies as the destination for the traffic.
Destination Device OS Version (dst_osversion)	The version of the operating system for the device that Device-ID identifies as the destination for the traffic.
Destination Hostname (dst_host)	The hostname of the device that Device-ID identifies as the destination for the traffic.
Destination MAC Address (dst_mac)	The MAC address for the device that Device-ID identifies as the destination for the traffic.
Container ID (container_id)	The container ID of the PAN-NGFW pod on the Kubernetes node where the application POD is deployed.
POD Namespace (pod_namespace)	The namespace of the application POD being secured.
POD Name (pod_name)	The application POD being secured.
Source External Dynamic List (src_edl)	The name of the external dynamic list that contains the source IP address of the traffic.
Destination External Dynamic List (dst_edl)	The name of the external dynamic list that contains the destination IP address of the traffic.
Host ID (hostid)	Unique ID GlobalProtect assigns to identify the host.
User Device Serial Number (serialnumber)	Serial number of the user's machine or device.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Field Name	Description
Source Dynamic Address Group (src_dag)	Original session source dynamic address group.
Destination Dynamic Address Group (dst_dag)	Original destination source dynamic address group.
Session Owner (session_owner)	The original high availability (HA) peer session owner in an HA cluster from which the session table data was synchronized upon HA failover.
High Resolution Timestamp (high_res_timestamp)	Time in milliseconds the log was received at the management plane. The format for this new field is YYYY-MM-DDThh:ss:sssTZD: • **YYYY**—Four digit year • **MM**—Two-digit month • **DD**—Two-digit day of the month (01 through 31) • **T**—Indicator for the beginning of the timestamp • **hh**—Two-digit hour using 24-hour time (00 through 23) • **mm**—Two-digit minute (00 through 59) • **ss**—Two-digit second (00 through 60) • **sss**—One or more digits for millisecond • **TZD**—Time zone designator (+hh:mm or -hh:mm) *The High Resolution Timestamp is supported for logs received from managed firewalls running PAN-OS 10.1 and later releases. Logs received from managed firewalls running PAN-OS 9.1 and earlier releases display a `1969-12-31T16:00:00:000-8:00` timestamp regardless of when the log was received.*
A Slice Service Type (nsdsai_sst)	The A Slice Service Type of the Network Slice ID.
A Slice Differentiator (nsdsai_sd)	The A Slice Differentiator of the Network Slice ID.
Application Subcategory (subcategory_of_app)	The application subcategory specified in the application configuration properties.
Application Category (category_of_app)	The application category specified in the application configuration properties. Values are:

https://t.me/learningnets

Field Name	Description
	• business-systems • collaboration • general-internet • media • networking • saas
Application Technology (technology_of_app)	The application technology specified in the application configuration properties. Values are: • browser-based • client-server • network-protocol • peer-to-peer
Application Risk (risk_of_app)	Risk level associated with the application (1=lowest to 5=highest).
Application Characteristic (characteristic_of_app)	Comma-separated list of applicable characteristic of the application
Application Container (container_of_app)	The parent application for an application.
Application SaaS (is_saas_of_app)	Displays `1` if a SaaS application or `0` if not a SaaS application.
Application Sanctioned State (sanctioned_state_of_app)	Displays `1` if application is sanctioned or `0` if application is not sanctioned.
Offloaded (offloaded)	Displays `1` if traffic flow has been offloaded or `0` if traffic flow was not offloaded.

Threat Log Fields

Format: FUTURE_USE, Receive Time, Serial Number, Type, Threat/Content Type, FUTURE_USE, Generated Time, Source Address, Destination Address, NAT Source IP, NAT Destination IP, Rule Name, Source User, Destination User, Application, Virtual System, Source Zone, Destination Zone, Inbound Interface, Outbound Interface, Log Action, FUTURE_USE, Session ID, Repeat Count, Source Port, Destination Port, NAT Source Port, NAT Destination Port, Flags, IP Protocol, Action, URL/Filename, Threat ID, Category, Severity, Direction, Sequence Number, Action Flags, Source Location, Destination Location, FUTURE_USE, Content Type, PCAP_ID, File Digest, Cloud, URL Index, User Agent, File Type, X-Forwarded-For, Referer, Sender, Subject, Recipient, Report ID, Device Group Hierarchy Level 1, Device Group Hierarchy Level 2, Device Group Hierarchy Level 3, Device Group Hierarchy Level 4, Virtual System Name, Device Name, FUTURE_USE, Source

https://t.me/learningnets

VM UUID, Destination VM UUID, HTTP Method, Tunnel ID/IMSI, Monitor Tag/IMEI, Parent Session ID, Parent Start Time, Tunnel Type, Threat Category, Content Version, FUTURE_USE, SCTP Association ID, Payload Protocol ID, HTTP Headers, URL Category List, Rule UUID, HTTP/2 Connection, Dynamic User Group Name, XFF Address, Source Device Category, Source Device Profile, Source Device Model, Source Device Vendor, Source Device OS Family, Source Device OS Version, Source Hostname, Source MAC Address, Destination Device Category, Destination Device Profile, Destination Device Model, Destination Device Vendor, Destination Device OS Family, Destination Device OS Version, Destination Hostname, Destination MAC Address, Container ID, POD Namespace, POD Name, Source External Dynamic List, Destination External Dynamic List, Host ID, Serial Number, Domain EDL, Source Dynamic Address Group, Destination Dynamic Address Group, Partial Hash, High Resolution Timestamp, Reason, Justification, A Slice Service Type, Application Subcategory, Application Category, Application Technology, Application Risk, Application Characteristic, Application Container, Application SaaS, Application Sanctioned State

Field Name	Description
Receive Time (receive_time or cef-formatted-receive_time)	Time the log was received at the management plane.
Serial Number (serial #)	Serial number of the firewall that generated the log.
Type (type)	Specifies the type of log; value is THREAT.
Threat/Content Type (subtype)	Subtype of threat log. Values include the following: • data—Data pattern matching a Data Filtering profile. • file—File type matching a File Blocking profile. • flood—Flood detected via a Zone Protection profile. • packet—Packet-based attack protection triggered by a Zone Protection profile. • scan—Scan detected via a Zone Protection profile. • spyware —Spyware detected via an Anti-Spyware profile. • url—URL filtering log. • ml-virus—Virus detected by WildFire Inline ML via an Antivirus profile. • virus—Virus detected via an Antivirus profile. • vulnerability —Vulnerability exploit detected via a Vulnerability Protection profile. • wildfire —A WildFire verdict generated when the firewall submits a file to WildFire per a WildFire Analysis profile and a verdict (malware, phishing, grayware, or benign, depending on what you are logging) is logged in the WildFire Submissions log. • wildfire-virus—Virus detected via an Antivirus profile.

https://t.me/learningnets

Field Name	Description
Generate Time (time_generated or cef-formatted-time_generated)	Time the log was generated on the dataplane.
Source address (src)	Original session source IP address.
Destination address (dst)	Original session destination IP address.
NAT Source IP (natsrc)	If source NAT performed, the post-NAT source IP address.
NAT Destination IP (natdst)	If destination NAT performed, the post-NAT destination IP address.
Rule Name (rule)	Name of the rule that the session matched.
Source User (srcuser)	Username of the user who initiated the session.
Destination User (dstuser)	Username of the user to which the session was destined.
Application (app)	Application associated with the session.
Virtual System (vsys)	Virtual System associated with the session.
Source Zone (from)	Zone the session was sourced from.
Destination Zone (to)	Zone the session was destined to.
Inbound Interface (inbound_if)	Interface that the session was sourced from.
Outbound Interface (outbound_if)	Interface that the session was destined to.
Log Action (logset)	Log Forwarding Profile that was applied to the session.
Session ID (sessionid)	An internal numerical identifier applied to each session.
Repeat Count (repeatcnt)	Number of sessions with same Source IP, Destination IP, Application, and Content/Threat Type seen within 5 seconds.
Source Port (sport)	Source port utilized by the session.
Destination Port (dport)	Destination port utilized by the session.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Field Name	Description
NAT Source Port (natsport)	Post-NAT source port.
NAT Destination Port (natdport)	Post-NAT destination port.
Flags (flags)	32-bit field that provides details on session; this field can be decoded by AND-ing the values with the logged value: • 0x80000000—session has a packet capture (PCAP) • 0x40000000—option is enabled to allow a client to use multiple paths to connect to a destination host • 0x20000000—file is submitted to WildFire for a verdict • 0x10000000—enterprise credential submission by end user detected • 0x08000000— source for the flow is on an allow list and not subject to recon protection • 0x02000000—IPv6 session • 0x01000000—SSL session is decrypted (SSL Proxy) • 0x00800000—session is denied via URL filtering • 0x00400000—session has a NAT translation performed • 0x00200000—user information for the session was captured through Authentication Portal • 0x00100000—application traffic is on a non-standard destination port • 0x00080000 —X-Forwarded-For value from a proxy is in the source user field • 0x00040000 —log corresponds to a transaction within a http proxy session (Proxy Transaction) • 0x00020000—Client to Server flow is subject to policy based forwarding • 0x00010000—Server to Client flow is subject to policy based forwarding • 0x00008000 —session is a container page access (Container Page) • 0x00002000 —session has a temporary match on a rule for implicit application dependency handling. Available in PAN-OS 5.0.0 and above. • 0x00000800 —symmetric return is used to forward traffic for this session

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Field Name	Description
	• 0x00000400—decrypted traffic is being sent out clear text through a mirror port • 0x00000010—payload of the outer tunnel is being inspected
IP Protocol (proto)	IP protocol associated with the session.
Action (action)	Action taken for the session; values are alert, allow, deny, drop, drop-all-packets, reset-client, reset-server, reset-both, block-url. • alert—threat or URL detected but not blocked • allow— flood detection alert • deny—flood detection mechanism activated and deny traffic based on configuration • drop— threat detected and associated session was dropped • reset-client —threat detected and a TCP RST is sent to the client • reset-server —threat detected and a TCP RST is sent to the server • reset-both —threat detected and a TCP RST is sent to both the client and the server • block-url —URL request was blocked because it matched a URL category that was set to be blocked • block-ip—threat detected and client IP is blocked • random-drop—flood detected and packet was randomly dropped • sinkhole—DNS sinkhole activated • syncookie-sent—syncookie alert • block-continue (URL subtype only)—a HTTP request is blocked and redirected to a Continue page with a button for confirmation to proceed • continue (URL subtype only)—response to a block-continue URL continue page indicating a block-continue request was allowed to proceed • block-override (URL subtype only)—a HTTP request is blocked and redirected to an Admin override page that requires a pass code from the firewall administrator to continue • override-lockout (URL subtype only)—too many failed admin override pass code attempts from the source IP. IP is now blocked from the block-override redirect page • override (URL subtype only)—response to a block-override page where a correct pass code is provided and the request is allowed • block (Wildfire only)—file was blocked by the firewall and uploaded to Wildfire

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Field Name	Description
URL/Filename (misc)	Field with variable length. A Filename has a maximum of 63 characters. A URL has a maximum of 1023 characters The actual URI when the subtype is url File name or file type when the subtype is file File name when the subtype is virus File name when the subtype is wildfire-virus File name when the subtype is wildfire URL or File name when the subtype is vulnerability if applicable URL when Threat Category is `domain-edl`
Threat/Content Name (threatid)	Palo Alto Networks identifier for known and custom threats. It is a description string followed by a 64-bit numerical identifier in parentheses for some Subtypes: • 8000 – 8099— scan detection • 8500 – 8599— flood detection • 9999— URL filtering log • 10000 – 19999 —spyware phone home detection • 20000 – 29999 —spyware download detection • 30000 – 44999 —vulnerability exploit detection • 52000 – 52999— filetype detection • 60000 – 69999 —data filtering detection If the Domain EDL field is populated, then this field is populated with the same value. *Threat ID ranges for virus detection, WildFire signature feed, and DNS C2 signatures used in previous releases have been replaced with permanent, globally unique IDs. Refer to the Threat/Content Type (subtype) and Threat Category (thr_category) field names to create updated reports, filter threat logs, and ACC activity.*
Category (category)	For URL Subtype, it is the URL Category; For WildFire subtype, it is the verdict on the file and is either 'malware', 'phishing', 'grayware', or 'benign'; For other subtypes, the value is 'any'.
Severity (severity)	Severity associated with the threat; values are informational, low, medium, high, critical.

https://t.me/learningnets

Field Name	Description
Direction (direction)	Indicates the direction of the attack, client-to-server or server-to-client: • 0—direction of the threat is client to server • 1—direction of the threat is server to client
Sequence Number (seqno)	A 64-bit log entry identifier incremented sequentially. Each log type has a unique number space.
Action Flags (actionflags)	A bit field indicating if the log was forwarded to Panorama.
Source Country (srcloc)	Source country or Internal region for private addresses. Maximum length is 32 bytes.
Destination Country (dstloc)	Destination country or Internal region for private addresses. Maximum length is 32 bytes.
Content Type (contenttype)	Applicable only when Subtype is URL. Content type of the HTTP response data. Maximum length 32 bytes.
PCAP ID (pcap_id)	The packet capture (pcap) ID is a 64 bit unsigned integral denoting an ID to correlate threat pcap files with extended pcaps taken as a part of that flow. All threat logs will contain either a pcap_id of 0 (no associated pcap), or an ID referencing the extended pcap file.
File Digest (filedigest)	Only for WildFire subtype; all other types do not use this field The filedigest string shows the binary hash of the file sent to be analyzed by the WildFire service.
Cloud (cloud)	Only for WildFire subtype; all other types do not use this field. The cloud string displays the FQDN of either the WildFire appliance (private) or the WildFire cloud (public) from where the file was uploaded for analysis.
URL Index (url_idx)	Used in URL Filtering and WildFire subtypes. When an application uses TCP keepalives to keep a connection open for a length of time, all the log entries for that session have a single session ID. In such cases, when you have a single threat log (and session ID) that includes multiple URL entries, the url_idx is a counter that allows you to correlate the order of each log entry within the single session. For example, to learn the URL of a file that the firewall forwarded to WildFire for analysis, locate the session ID and the url_idx from the WildFire Submissions log and search for the same session ID

https://t.me/learningnets

Field Name	Description
	and url_idx in your URL filtering logs. The log entry that matches the session ID and url_idx will contain the URL of the file that was forwarded to WildFire.
User Agent (user_agent)	Only for the URL Filtering subtype; all other types do not use this field. The User Agent field specifies the web browser that the user used to access the URL, for example Internet Explorer. This information is sent in the HTTP request to the server.
File Type (filetype)	Only for WildFire subtype; all other types do not use this field. Specifies the type of file that the firewall forwarded for WildFire analysis.
X-Forwarded-For (xff)	Only for the URL Filtering subtype; all other types do not use this field. The X-Forwarded-For field in the HTTP header contains the IP address of the user who requested the web page. It allows you to identify the IP address of the user, which is useful particularly if you have a proxy server on your network that replaces the user IP address with its own address in the source IP address field of the packet header.
Referer (referer)	Only for the URL Filtering subtype; all other types do not use this field. The Referer field in the HTTP header contains the URL of the web page that linked the user to another web page; it is the source that redirected (referred) the user to the web page that is being requested.
Sender (sender)	Specifies the name of the sender of an email.
Subject (subject)	Specifies the subject of an email.
Recipient (recipient)	Specifies the name of the receiver of an email.
Report ID (reportid)	Only for WildFire subtype; all other types do not use this field. Identifies the analysis request on the WildFire cloud or the WildFire appliance.
Device Group Hierarchy (dg_hier_level_1 to dg_hier_level_4)	A sequence of identification numbers that indicate the device group's location within a device group hierarchy. The firewall (or virtual system) generating the log includes the identification number of each ancestor in its device group hierarchy. The shared device group (level 0) is not included in this structure.

https://t.me/learningnets

Field Name	Description
	If the log values are 12, 34, 45, 0, it means that the log was generated by a firewall (or virtual system) that belongs to device group 45, and its ancestors are 34, and 12. To view the device group names that correspond to the value 12, 34 or 45, use one of the following methods: API query: `/api/?type=op&cmd=<show><dg-hierarchy></dg-hierarchy></show>`
Virtual System Name (vsys_name)	The name of the virtual system associated with the session; only valid on firewalls enabled for multiple virtual systems.
Device Name (device_name)	The hostname of the firewall on which the session was logged.
Source VM UUID (src_uuid)	Identifies the source universal unique identifier for a guest virtual machine in the VMware NSX environment.
Destination VM UUID (dst_uuid)	Identifies the destination universal unique identifier for a guest virtual machine in the VMware NSX environment.
HTTP Method (http_method)	Only in URL filtering logs. Describes the HTTP Method used in the web request. Only the following methods are logged: Connect, Delete, Get, Head, Options, Post, Put.
Tunnel ID/IMSI (tunnel_id/imsi)	International Mobile Subscriber Identity (IMSI) is a unique number allocated to each mobile subscriber in the GSM/UMTS/EPS system. IMSI shall consist of decimal digits (0 through 9) only and maximum number of digits allowed are 15.
Monitor Tag/IMEI (monitortag/imei)	International Mobile Equipment Identity (IMEI) is a unique 15 or 16 digit number allocated to each mobile station equipment.
Parent Session ID (parent_session_id)	ID of the session in which this session is tunneled. Applies to inner tunnel (if two levels of tunneling) or inside content (if one level of tunneling) only.
Parent Session Start Time (parent_start_time)	Year/month/day hours:minutes:seconds that the parent tunnel session began.
Tunnel Type (tunnel)	Type of tunnel, such as GRE or IPSec.
Threat Category (thr_category)	Describes threat categories used to classify different types of threat signatures.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Field Name	Description
	If a domain external dynamic list generated the log, `domain-edl` populates this field.
Content Version (contentver)	Applications and Threats version on your firewall when the log was generated.
SCTP Association ID (assoc_id)	Number that identifies all connections for an association between two SCTP endpoints.
Payload Protocol ID (ppid)	ID of the protocol for the payload in the data portion of the data chunk.
HTTP Headers (http_headers)	Indicates the inserted HTTP header in the URL log entries on the firewall.
URL Category List (url_category_list)	Lists the URL Filtering categories that the firewall used to enforce policy.
Rule UUID (rule_uuid)	The UUID that permanently identifies the rule.
HTTP/2 Connection (http2_connection)	Identifies if traffic used an HTTP/2 connection by displaying one of the following values: • TCP connection session ID—session is HTTP/2 • 0—session is not HTTP/2
Dynamic User Group Name (dynusergroup_name)	The name of the dynamic user group that contains the user who initiated the session.
XFF Address (xff_ip)	The IP address of the user who requested the web page or the IP address of the next to last device that the request traversed. If the request goes through one or more proxies, load balancers, or other upstream devices, the firewall displays the IP address of the most recent device.
Source Device Category (src_category)	The category for the device that Device-ID identifies as the source of the traffic.
Source Device Profile (src_profile)	The device profile for the device that Device-ID identifies as the source of the traffic.
Source Device Model (src_model)	The model of the device that Device-ID identifies as the source of the traffic.

https://t.me/learningnets

Field Name	Description
Source Device Vendor (src_vendor)	The vendor of the device that Device-ID identifies as the source of the traffic.
Source Device OS Family (src_osfamily)	The operating system type for the device that Device-ID identifies as the source of the traffic.
Source Device OS Version (src_osversion)	The version of the operating system for the device that Device-ID identifies as the source of the traffic.
Source Hostname (src_host)	The hostname of the device that Device-ID identifies as the source of the traffic.
Source MAC Address (src_mac)	The MAC address for the device that Device-ID identifies as the source of the traffic.
Destination Device Category (dst_category)	The category for the device that Device-ID identifies as the destination for the traffic.
Destination Device Profile (dst_profile)	The device profile for the device that Device-ID identifies as the destination for the traffic.
Destination Device Model (dst_model)	The model of the device that Device-ID identifies as the destination for the traffic.
Destination Device Vendor (dst_vendor)	The vendor of the device that Device-ID identifies as the destination for the traffic.
Destination Device OS Family (dst_osfamily)	The operating system type for the device that Device-ID identifies as the destination for the traffic.
Destination Device OS Version (dst_osversion)	The version of the operating system for the device that Device-ID identifies as the destination for the traffic.
Destination Hostname (dst_host)	The hostname of the device that Device-ID identifies as the destination for the traffic.
Destination MAC Address (dst_mac)	The MAC address for the device that Device-ID identifies as the destination for the traffic.
Container ID (container_id)	The container ID of the PAN-NGFW pod on the Kubernetes node where the application POD is deployed.
POD Namespace (pod_namespace)	The namespace of the application POD being secured.
POD Name (pod_name)	The application POD being secured.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Field Name	Description
Source External Dynamic List (src_edl)	The name of the external dynamic list that contains the source IP address of the traffic.
Destination External Dynamic List (dst_edl)	The name of the external dynamic list that contains the destination IP address of the traffic.
Host ID (hostid)	Unique ID GlobalProtect assigns to identify the host.
User Device Serial Number (serialnumber)	Serial number of the user's machine or device.
Domain EDL (domain_edl)	The name of the external dynamic list that contains the domain name of the traffic.
Source Dynamic Address Group (src_dag)	Original session source dynamic address group.
Destination Dynamic Address Group (dst_dag)	Original destination source dynamic address group.
Partial Hash (partial_hash)	Machine Learning partial hash.
High Resolution Timestamp (high_res timestamp)	Time in milliseconds the log was received at the management plane. The format for this new field is YYYY-MM-DDThh:ss:sssTZD: • **YYYY**—Four digit year • **MM**—Two-digit month • **DD**—Two-digit day of the month (01 through 31) • **T**—Indicator for the beginning of the timestamp • **hh**—Two-digit hour using 24-hour time (00 through 23) • **mm**—Two-digit minute (00 through 59) • **ss**—Two-digit second (00 through 60) • **sss**—One or more digits for millisecond • **TZD**—Time zone designator (+hh:mm or -hh:mm) *The High Resolution Timestamp is supported for logs received from managed firewalls running PAN-OS 10.1 and later releases. Logs received from managed firewalls running PAN-OS 9.1 and earlier releases display a `1969-12-31T16:00:00:000-8:00` timestamp regardless of when the log was received.*

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Field Name	Description
Reason (reason)	Reason for Data Filtering action.
Justification (justification)	Justification for Data Filtering action.
A Slice Service Type (nssai_sst)	The A Slice Service Type of the Network Slice ID.
Application Subcategory (subcategory_of_app)	The application subcategory specified in the application configuration properties.
Application Category (category_of_app)	The application category specified in the application configuration properties. Values are: • business-systems • collaboration • general-internet • media • networking • saas
Application Technology (technology_of_app)	The application technology specified in the application configuration properties. Values are: • browser-based • client-server • network-protocol • peer-to-peer
Application Risk (risk_of_app)	Risk level associated with the application (1=lowest to 5=highest).
Application Characteristic (characteristic_of_app)	Comma-separated list of applicable characteristic of the application
Application Container (container_of_app)	The parent application for an application.
Application SaaS (is_saas_of_app)	Displays `1` if a SaaS application or `0` if not a SaaS application.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Field Name	Description
Application Sanctioned State (sanctioned_state_of_app)	Displays `1` if application is sanctioned or `0` if application is not sanctioned.

HIP Match Log Fields

Format: FUTURE_USE, Receive Time, Serial Number, Type, Threat/Content Type, FUTURE_USE, Generated Time, Source User, Virtual System, Machine Name, Operating System, Source Address, HIP, Repeat Count, HIP Type, FUTURE_USE, FUTURE_USE, Sequence Number, Action Flags, Device Group Hierarchy Level 1, Device Group Hierarchy Level 2, Device Group Hierarchy Level 3, Device Group Hierarchy Level 4, Virtual System Name, Device Name, Virtual System ID, IPv6 Source Address, Host ID, User Device Serial Number, Device MAC Address, High Resolution Timestamp

Field Name	Description
Receive Time (receive_time or cef-formatted-receive_time)	Time the log was received at the management plane.
Serial Number (serial)	Serial number of the firewall that generated the log.
Type (type)	Specifies the type of log; value is HIP-MATCH.
Threat/Content Type (subtype)	Subtype of HIP match log; unused.
Generated Time (time_generated or cef-formatted-time_generated)	Time the log was generated on the dataplane.
Source User (srcuser)	Username of the user who initiated the session.
Virtual System (vsys)	Virtual System associated with the HIP match log.
Machine Name (machinename)	Name of the user's machine.
Operating System (os)	The operating system installed on the user's machine or device (or on the client system).

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Field Name	Description
Source Address (src)	IP address of the source user.
HIP (matchname)	Name of the HIP object or profile.
Repeat Count (repeatcnt)	Number of times the HIP profile matched.
HIP Type (matchtype)	Whether the hip field represents a HIP object or a HIP profile.
Sequence Number (seqno)	A 64-bit log entry identifier incremented sequentially; each log type has a unique number space.
Action Flags (actionflags)	A bit field indicating if the log was forwarded to Panorama.
Device Group Hierarchy (dg_hier_level_1 to dg_hier_level_4)	A sequence of identification numbers that indicate the device group's location within a device group hierarchy. The firewall (or virtual system) generating the log includes the identification number of each ancestor in its device group hierarchy. The shared device group (level 0) is not included in this structure. If the log values are 12, 34, 45, 0, it means that the log was generated by a firewall (or virtual system) that belongs to device group 45, and its ancestors are 34, and 12. To view the device group names that correspond to the value 12, 34 or 45, use one of the following methods: API query: `/api/?type=op&cmd=<show><dg-hierarchy></dg-hierarchy></show>`
Virtual System Name (vsys_name)	The name of the virtual system associated with the session; only valid on firewalls enabled for multiple virtual systems.
Device Name (device_name)	The hostname of the firewall on which the session was logged.
Virtual System ID (vsys_id)	A unique identifier for a virtual system on a Palo Alto Networks firewall.
IPv6 System Address (srcipv6)	IPv6 address of the user's machine or device.
Host ID (hostid)	Unique ID GlobalProtect assigns to identify the host.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Field Name	Description
User Device Serial Number (serialnumber)	Serial number of the user's machine or device.
Device MAC Address (mac)	The MAC address of the user's machine or device.
High Resolution Timestamp (high_res_timestamp)	Time in milliseconds the log was received at the management plane. The format for this new field is YYYY-MM-DDThh:ss:sssTZD: • **YYYY**—Four digit year • **MM**—Two-digit month • **DD**—Two-digit day of the month (01 through 31) • **T**—Indicator for the beginning of the timestamp • **hh**—Two-digit hour using 24-hour time (00 through 23) • **mm**—Two-digit minute (00 through 59) • **ss**—Two-digit second (00 through 60) • **sss**—One or more digits for millisecond • **TZD**—Time zone designator (+hh:mm or -hh:mm) *The High Resolution Timestamp is supported for logs received from managed firewalls running PAN-OS 10.1 and later releases. Logs received from managed firewalls running PAN-OS 9.1 and earlier releases display a `1969-12-31T16:00:00:000-8:00` timestamp regardless of when the log was received.*

GlobalProtect Log Fields

Format: FUTURE_USE, Receive Time, Serial Number, Type, Threat/Content Type, FUTURE_USE, Generated Time, Virtual System, Event ID, Stage, Authentication Method, Tunnel Type, Source User, Source Region, Machine Name, Public IP, Public IPv6, Private IP, Private IPv6, Host ID, Serial Number, Client Version, Client OS, Client OS Version, Repeat Count, Reason, Error, Description, Status, Location, Login Duration, Connect Method, Error Code, Portal, Sequence Number, Action Flags, High Res Timestamp, Selection Type, Response Time, Priority, Attempted Gateways, Gateway, Device Group Hierarchy Level 1, Device Group Hierarchy Level 2, Device Group Hierarchy Level 3, Device Group Hierarchy Level 4, Virtual System Name, Device Name, Virtual System ID

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Field Name	Description
Receive Time (receive_time)	The time that the log was received at the management plane.
Serial # (serial)	The serial number of the firewall that generated the log.
Type (type)	Specifies the type of log; value is GLOBALPROTECT.
Threat/Content Type (subtype)	Subtype of threat log. Values include the following: • data—Data pattern matching a Data Filtering profile. • file—File type matching a File Blocking profile. • flood—Flood detected via a Zone Protection profile. • packet—Packet-based attack protection triggered by a Zone Protection profile. • scan—Scan detected via a Zone Protection profile. • spyware —Spyware detected via an Anti-Spyware profile. • url—URL filtering log. • virus—Virus detected via an Antivirus profile. • vulnerability —Vulnerability exploit detected via a Vulnerability Protection profile. • wildfire —A WildFire verdict generated when the firewall submits a file to WildFire per a WildFire Analysis profile and a verdict (malicious, phishing, grayware, or benign, depending on what you are logging) is logged in the WildFire Submissions log. • wildfire-virus—Virus detected via an Antivirus profile.
Generate Time (time_generated)	The time that the log was generated on the dataplane.
Virtual System (vsys)	The Virtual System associated with the session.
Event ID (eventid)	A string showing the name of the event.
Stage (stage)	A string showing the stage of the connection (for example, `before-login`, `login`, or `tunnel`).
Authentication Method (auth_method)	A string showing the authentication type, such as `LDAP`, `RADIUS`, or `SAML`.
Tunnel Type (tunnel_type)	The type of tunnel (either SSLVPN or IPSec).

https://t.me/learningnets

Field Name	Description
Source User (srcuser)	The username of the user who initiated the session.
Source Region (srcregion)	The region for the user who initiated the session.
Machine Name (machinename)	The name of the user's machine.
Public IP (public_ip)	The public IP address for the user who initiated the session.
Public IPv6 (public_ipv6)	The public IPv6 address for the user who initiated the session.
Private IP (private_ip)	The private IP address for the user who initiated the session.
Private IPv6 (private_ipv6)	The private IPv6 address for the user who initiated the session.
Host ID (hostid)	The unique ID that GlobalProtect assigns to identify the host.
Serial Number (serialnumber)	The serial number of the user's machine or device.
Client Version (client_ver)	The client's GlobalProtect app version.
Client OS (client_os)	The client device's OS type (for example, Windows or Linux).
Client OS Version (client_os_ver)	The client device's OS version.
Repeat Count (repeatcnt)	The number of sessions with the same source IP address, destination IP address, application, and subtype that GlobalProtect has detected within the last five seconds.
Reason (reason)	A string that shows the reason for the quarantine.
Error (error)	A string showing that error that has occurred in any event.
Description (opaque)	Additional information for any event that has occurred.
Status (status)	The status (success or failure) of the event.
Location (location)	A string showing the administrator-defined location of the GlobalProtect portal or gateway.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Field Name	Description
Login Duration (login_duration)	The length of time, in seconds, the user is connected to the GlobalProtect gateway from logging in to logging out.
Connect Method (connect_method)	A string showing the how the GlobalProtect app connects to Gateway, (for example, `on-demand` or `user-logon`.
Error Code (error_code)	An integer associated with any errors that occurred.
Portal (portal)	The name of the GlobalProtect portal or gateway.
Sequence Number (seqno)	A 64-bit log entry identifier incremented sequentially; each log type has a unique number space.
Action Flags (actionflags)	A bit field indicating if the log was forwarded to Panorama.
Gateway Selection Method (selection_type)	The connection method that is selected to connect to the gateway. • manual—The gateway to which you want the GlobalProtect app to manually connect. • preferred—The preferred gateway to which you want the GlobalProtect app to connect. • auto—Automatically connect to the **Best Available** gateway based on the priority assigned to the gateway and the response time.
SSL Response Time (response_time)	The SSL response time of the selected gateway that is measured in milliseconds on the endpoint during tunnel setup.
Gateway Priority (priority)	The priority order of the gateway that is based on highest (1), high (2), medium (3), low (4), or lowest (5) to which the GlobalProtect app can connect.
Attempted Gateways (attempted_gateways)	The fields that are collected for each gateway connection attempt with the gateway name, SSL response time, and priority (see Gateway Priority in a Multiple Gateway Configuration. Each field entry is separated by commas such as `g82-gateway,12,3`. Each gateway entry is separated by semicolons such as `g83-gateway,10,2;g84-gateway,-1,1`.
Gateway Name (gateway)	The name of the gateway that is specified on the portal configuration.
Device Group Hierarchy	A sequence of identification numbers that indicate the device group's location within a device group hierarchy. The firewall (or virtual system) generating the log includes the identification number of each ancestor

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Field Name	Description
(dg_hier_level_1 to dg_hier_level_4)	in its device group hierarchy. The shared device group (level 0) is not included in this structure. If the log values are 12, 34, 45, 0, it means that the log was generated by a firewall (or virtual system) that belongs to device group 45, and its ancestors are 34, and 12. To view the device group names that correspond to the value 12, 34 or 45, use one of the following methods: API query: `/api/?type=op&cmd=<show><dg-hierarchy></dg-hierarchy></show>`
Virtual System Name (vsys_name)	The name of the virtual system associated with the session; only valid on firewalls enabled for multiple virtual systems.
Device Name (device_name)	The hostname of the firewall on which the session was logged.
Virtual System ID (vsys_id)	A unique identifier for a virtual system on a Palo Alto Networks firewall.

IP-Tag Log Fields

Format: FUTURE_USE , Receive Time, Serial, Type, Threat/Content Type, FUTURE_USE, Generate Time, Virtual System, Source IP, Tag Name , Event ID, Repeat Count , Timeout, Data Source Name, Data Source Type, Data Source Subtype, Sequence Number, Action Flags, DG Hierarchy Level 1 , DG Hierarchy Level 2, DG Hierarchy Level 3, DG Hierarchy Level 4, Virtual System Name, Device Name, Virtual System ID, High Resolution Timestamp

Field Name	Description
Receive Time (receive_time or cef-formatted-receive_time)	The time the log was received at the management plane.
Serial Number (serial)	The serial number of the firewall that generated the log.
Type (type)	Specifies the type of log; value is IPTAG.
Threat/Content Type (subtype)	The subtype of the HIP match log; unused.
Generated Time (time_generated	The time the log was generated on the dataplane.

https://t.me/learningnets

Field Name	Description
or cef-formatted-time_generated)	
Virtual System (vsys)	The virtual system associated with the HIP match log.
Source IP (src)	The IP address of the source user.
Tag Name (tag_name)	The tag mapped to the source IP address.
Event ID (event_id)	A string showing the name of the event.
Repeat Count (repeatcnt)	The number of sessions with the same Source IP, Destination IP, Application, and Subtype seen within 5 seconds.
Timeout (timeout)	The amount of time before the IP address-to-tag mapping expires for the source IP address.
Data Source Name (datasourcename)	The name of the source from which mapping information is collected.
Data Source Type (datasource_type)	The source from which mapping information is collected.
Data Source Subtype (datasource_subtype)	The mechanism used to identify the IP address-to-username mappings within a data source.
Sequence Number (seqno)	A 64-bit log entry identifier incremented sequentially. Each log type has a unique number space.
Action Flags (actionflags)	A bit field indicating whether the log was forwarded to Panorama.
Device Group Hierarchy (dg_hier_level_1 to dg_hier_level_4)	A sequence of identification numbers that indicates the location of the device group within a device group hierarchy. The firewall (or virtual system) generating the log includes the identification number of each ancestor in its device group hierarchy except the shared device group (level 0), which is not included in this structure. If the log values are 12, 34, 45, and 0, it means that the log was generated by a firewall (or virtual system) that belongs to device group 45 and its ancestors are 34 and 12. To view the device group names that correspond to the value 12, 34, or 45, use one of the following methods: API query: `/api/?type=op&cmd=<show><dg-hierarchy></dg-hierarchy></show>`

https://t.me/learningnets

Field Name	Description
Virtual System Name (vsys_name)	The name of the virtual system associated with the session; only valid on firewalls enabled for multiple virtual systems.
Device Name (device_name)	The hostname of the firewall on which the session was logged.
Virtual System ID (vsys_id)	A unique identifier for a virtual system on a Palo Alto Networks firewall.
High Resolution Timestamp (high_res timestamp)	Time in milliseconds the log was received at the management plane. The format for this new field is YYYY-MM-DDThh:ss:sssTZD: • **YYYY**—Four digit year • **MM**—Two-digit month • **DD**—Two-digit day of the month (01 through 31) • **T**—Indicator for the beginning of the timestamp • **hh**—Two-digit hour using 24-hour time (00 through 23) • **mm**—Two-digit minute (00 through 59) • **ss**—Two-digit second (00 through 60) • **sss**—One or more digits for millisecond • **TZD**—Time zone designator (+hh:mm or -hh:mm) *The High Resolution Timestamp is supported for logs received from managed firewalls running PAN-OS 10.1 and later releases. Logs received from managed firewalls running PAN-OS 9.1 and earlier releases display a `1969-12-31T16:00:00:000-8:00` timestamp regardless of when the log was received.*

User-ID Log Fields

Format: FUTURE_USER, Receive Time, Serial Number, Type, Threat/Content Type, FUTURE_USE, Generated Time, Virtual System, Source IP, User, Data Source Name, Event ID, Repeat Count, Time Out Threshold, Source Port, Destination Port, Data Source, Data Source Type, Sequence Number, Action Flags, Device Group Hierarchy Level 1, Device Group Hierarchy Level 2, Device Group Hierarchy Level 3, Device Group Hierarchy Level 4, Virtual System Name, Device Name, Virtual System ID, Factor Type, Factor Completion Time, Factor Number, FUTURE_USE, FUTURE_USE, User Group Flags, User by Source, High Resolution Timestamp

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Field Name	Description
Receive Time (receive_time or cef-formatted-receive_time)	Time the log was received at the management plane.
Serial Number (serial)	Serial number of the firewall that generated the log.
Type (type)	Specifies the type of log; value is USERID.
Threat/Content Type (subtype)	Subtype of User-ID log; values are login, logout, register-tag, and unregister-tag. • login—User logged in. • logout—User logged out. • register-tag—Indicates a tag or tags were registered for the user. • unregister-tag—Indicates a tag or tags were unregistered for the user.
Generated Time (time_generated or cef-formatted-time_generated)	The time the log was generated on the dataplane.
Virtual System (vsys)	Virtual System associated with the configuration log.
Source IP (ip)	Original session source IP address.
User (user)	Identifies the end user.
Data Source Name (datasourcename)	User-ID source that sends the IP (Port)-User Mapping.
Event ID (eventid)	String showing the name of the event.
Repeat Count (repeatcnt)	Number of sessions with same Source IP, Destination IP, Application, and Subtype seen within 5 seconds.
Time Out Threshold (timeout)	Timeout after which the IP/User Mappings are cleared.
Source Port (beginport)	Source port utilized by the session.
Destination Port (endport)	Destination port utilized by the session.
Data Source (datasource)	Source from which mapping information is collected.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Field Name	Description
Data Source Type (datasourcetype)	Mechanism used to identify the IP/User mappings within a data source.
Sequence Number (seqno)	Serial number of the firewall that generated the log.
Action Flags (actionflags)	A bit field indicating if the log was forwarded to Panorama.
Device Group Hierarchy (dg_hier_level_1 to dg_hier_level_4)	A sequence of identification numbers that indicate the device group's location within a device group hierarchy. The firewall (or virtual system) generating the log includes the identification number of each ancestor in its device group hierarchy. The shared device group (level 0) is not included in this structure. If the log values are 12, 34, 45, 0, it means that the log was generated by a firewall (or virtual system) that belongs to device group 45, and its ancestors are 34, and 12. To view the device group names that correspond to the value 12, 34 or 45, use one of the following methods: API query: `/api/?type=op&cmd=<show><dg-hierarchy></dg-hierarchy></show>`
Virtual System Name (vsys_name)	The name of the virtual system associated with the session; only valid on firewalls enabled for multiple virtual systems.
Device Name (device_name)	The hostname of the firewall on which the session was logged.
Virtual System ID (vsys_id)	A unique identifier for a virtual system on a Palo Alto Networks firewall.
Factor Type (factortype)	Vendor used to authenticate a user when Multi Factor authentication is present.
Factor Completion Time (factorcompletiontime)	Time the authentication was completed.
Factor Number (factorno)	Indicates the use of primary authentication (1) or additional factors (2, 3).
User Group Flags (ugflags)	Displays whether the user group that was found during user group mapping. Supported values are: • User Group Found—Indicates whether the user could be mapped to a group.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Field Name	Description
	• Duplicate User—Indicates whether duplicate users were found in a user group. Displays N/A if no user group is found.
User by Source (userbysource)	Indicates the username received from the source through IP address-to-username mapping.
High Resolution Timestamp (high_res timestamp)	Time in milliseconds the log was received at the management plane. The format for this new field is YYYY-MM-DDThh:ss:sssTZD: • **YYYY**—Four digit year • **MM**—Two-digit month • **DD**—Two-digit day of the month (01 through 31) • **T**—Indicator for the beginning of the timestamp • **hh**—Two-digit hour using 24-hour time (00 through 23) • **mm**—Two-digit minute (00 through 59) • **ss**—Two-digit second (00 through 60) • **sss**—One or more digits for millisecond • **TZD**—Time zone designator (+hh:mm or -hh:mm) *The High Resolution Timestamp is supported for logs received from managed firewalls running PAN-OS 10.1 and later releases. Logs received from managed firewalls running PAN-OS 9.1 and earlier releases display a `1969-12-31T16:00:00:000-8:00` timestamp regardless of when the log was received.*

Decryption Log Fields

Format: FUTURE_USE, Receive Time, Serial Number, Type, Threat/Content Type, Config Version, Generate Time, Source Address, Destination Address, NAT Source IP, NAT Destination IP, Rule, Source User, Destination User, Application, Virtual System, Source Zone, Destination Zone, Inbound Interface, Outbound Interface, Log Action, Time Logged, Session ID, Repeat Count, Source Port, Destination Port, NAT Source Port, NAT Destination Port, Flags, IP Protocol, Action, Tunnel, FUTURE_USE, FUTURE_USE, Source VM UUID, Destination VM UUID, UUID for rule, Stage for Client to Firewall, Stage for Firewall to Server, TLS Version, Key Exchange Algorithm, Encryption Algorithm, Hash Algorithm, Policy Name, Elliptic Curve, Error Index, Root Status, Chain Status, Proxy Type, Certificate Serial Number, Fingerprint, Certificate Start Date, Certificate End Date, Certificate Version, Certificate Size, Common Name Length, Issuer Common Name Length, Root Common Name Length, SNI Length, Certificate Flags, Subject Common Name, Issuer Subject Common Name, Root Subject Common Name, Server Name Indication, Error, Container ID, POD Namespace, POD Name, Source External Dynamic List, Destination External Dynamic List, Source Dynamic Address Group, Destination Dynamic Address Group, High Res Timestamp, Source Device Category, Source Device Profile, Source Device Model, Source Device Vendor, Source Device OS Family, Source Device OS Version, Source Hostname, Source Mac Address, Destination

https://t.me/learningnets

Device Category, Destination Device Profile, Destination Device Model, Destination Device Vendor, Destination Device OS Family, Destination Device OS Version, Destination Hostname, Destination Mac Address, Sequence Number, Action Flags, Device Group Hierarchy Level 1, Device Group Hierarchy Level 2, Device Group Hierarchy Level 3, Device Group Hierarchy Level 4, Virtual System Name, Device Name, Virtual System ID, Application Subcategory, Application Category, Application Technology, Application Risk, Application Characteristic, Application Container, Application SaaS, Application Sanctioned State

Field Name	Description
Receive Time (receive_time or cef-formatted-receive_time)	Time the log was received at the management plane.
Serial Number (serial)	Serial number of the firewall that generated the log.
Type (type)	Specifies the type of log; value is DECRYPTION.
Threat/ContentType (subtype)	Not used in the Decryption log.
Config Version (config_ver)	The software version.
Generate Time (time_generated)	Time the log was generated on the dataplane.
Source Address (src)	Original session source IP address.
Destination Address (dst)	Original session destination IP address.
NAT Source IP (natsrc)	If Source NAT performed, the post-NAT Source IP address.
NAT Destination IP (natdst)	If Destination NAT performed, the post-NAT Destination IP address.
Rule (rule)	Security policy rule that controls the session traffic.
Source User (srcuser)	Username of the user who initiated the session.
Destination User (dstuser)	Username of the user to which the session was destined.
Application (app)	Application associated with the session.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Field Name	Description
Virtual System (vsys)	Virtual System associated with the session.
Source Zone (from)	Zone the session was sourced from.
Destination Zone (to)	Zone the session was destined to.
Inbound Interface (inbound_if)	Interface that the session was sourced from.
Outbound Interface (outbound_if)	Interface that the session was destined to.
Log Action (logset)	Log Forwarding profile applied to the session.
Time Logged (time_received)	The time the log was received.
Session ID (sessionid)	An internal numerical identifier applied to each session.
Repeat Count (repeatcnt)	Number of sessions with the same Source IP, Destination IP, Application, and Content/Threat Type seen within 5 seconds.
Source Port (sport)	Source port utilized by the session.
Destination Port (dport)	Destination port utilized by the session.
NAT Source Port (natsport)	Post-NAT source port.
NAT Destination Port (natdport)	Post-NAT destination port.
Flags (flags)	32-bit field that provides details on session; this field can be decoded by AND-ing the values with the logged value: • 0x80000000—session has a packet capture (PCAP) • 0x40000000—option is enabled to allow a client to use multiple paths to connect to a destination host • 0x20000000—file is submitted to WildFire for a verdict • 0x10000000—enterprise credential submission by end user detected • 0x08000000— source for the flow is on the allow list and not subject to recon protection • 0x02000000—IPv6 session

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Field Name	Description
	• 0x01000000—SSL session is decrypted (SSL Proxy) • 0x00800000—session is denied via URL filtering • 0x00400000—session has a NAT translation performed • 0x00200000—user information for the session was captured through Authentication Portal • 0x00100000—application traffic is on a non-standard destination port • 0x00080000 —X-Forwarded-For value from a proxy is in the source user field • 0x00040000—log corresponds to a transaction within a http proxy session (Proxy Transaction) • 0x00020000—Client to Server flow is subject to policy based forwarding • 0x00010000—Server to Client flow is subject to policy based forwarding • 0x00008000—session is a container page access (Container Page) • 0x00002000—session has a temporary match on a rule for implicit application dependency handling. Available in PAN-OS 5.0.0 and above. • 0x00000800—symmetric return is used to forward traffic for this session • 0x00000400—decrypted traffic is being sent out clear text through a mirror port • 0x00000100—payload of the outer tunnel is being inspected
IP Protocol (proto)	IP protocol associated with the session.
Action (action)	Action taken for the session; possible values are: • allow—session was allowed by policy • deny—session was denied by policy • drop—session was dropped silently • drop ICMP—session was silently dropped with an ICMP unreachable message to the host or application • reset both—session was terminated and a TCP reset is sent to both the sides of the connection • reset client—session was terminated and a TCP reset is sent to the client • reset server—session was terminated and a TCP reset is sent to the server

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Field Name	Description
Tunnel (tunnel)	Type of tunnel.
Source VM UUID (src_uuid)	The source universal unique identifier for a guest virtual machine in the VMware NSX environment.
Destination VM UUID (dst_uuid)	The destination universal unique identifier for a guest virtual machine in the VMware NSX environment.
UUID for rule (rule_uuid)	The UUID that permanently identifies the rule.
Stage for Client to Firewall (hs_stage_c2f)	The stage of the TLS handshake from the client to the firewall, for example, Client Hello, Server Hello, Certificate, Client/Server key exchange, etc.
Stage for Firewall to Server (hs_stage_f2s)	The stage of the TLS handshake from the firewall to the server.
TLS Version (tls_version)	The version of TLS protocol used for the session.
Key Exchange Algorithm (tls_keyxchg)	The key exchange algorithm used for the session.
Encryption Algorithm (tls_enc)	The algorithm used to encrypt the session data, such as AES-128-CBC, AES-256-GCM, etc.
Hash Algorithm (tls_auth)	The authentication algorithm used for the session, for example, SHA, SHA256, SHA384, etc.
Policy Name (policy_name)	The name of the Decryption policy associated with the session.
Elliptic Curve (ec_curve)	The elliptic cryptography curve that the client and server negotiate and use for connections that use ECDHE cipher suites.
Error Index (err_index)	The type of error that occurred: Cipher, Resource, Resume, Version, Protocol, Certificate, Feature, or HSM.
Root Status (root_status)	The status of the root certificate, for example, trusted, untrusted, or uninspected.
Chain Status (chain_status)	Whether the chain is trusted. Values are: • Uninspected • Untrusted

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Field Name	Description
	• Trusted • Incomplete
Proxy Type (proxy_type)	The Decryption proxy type, such as Forward for Forward Proxy, Inbound for Inbound Inspection, No Decrypt for undecrypted traffic, GlobalProtect, etc.
Certificate Serial Number (cert_serial)	The unique identifier of the certificate (generated by the certificate issuer).
Certificate Fingerprint (fingerprint)	A hash of the certificate in x509 binary format.
Certificate Start Date (notbefore)	The time the certificate became valid (certificate in invalid before this time).
Certificate End Date (notafter)	The time the certificate expires (certificate becomes invalid after this time).
Certificate Version (cert_ver)	The certificate version (V1, V2, or V3).
Certificate Size (cert_size)	The certificate key size.
Common Name Length (cn_len)	The length of the subject common name.
Issuer Common Name Length (issuer_len)	The length of the issuer common name.
Root Common Name Length (rootcn_len)	The length of the root common name.
SNI Length (sni_len)	The length of the Server Name Indication (hostname).
Certificate Flags (cert_flags)	The certificate flags can return seven values: • Session is resumed (b_resume_session) • Certificate (subject) common name is truncated (b_cert_cn_truncated) • Issuer common name is truncated (b_issuer_cn_truncated) • Root common name is truncated (b_root_cn_truncated) • Server Name Indication (SNI) is truncated (b_sni_truncated)

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Field Name	Description
	• Certificate type, RSA or ECDSA (b_cert_type) • Unused (padding3)
Subject Common Name (cn)	The domain name (the name of the server that the certificate protects).
Issuer Common Name (issuer_cn)	The name of the organization that verified the certificate's contents.
Root Common Name (root_cn)	The name of the root certificate authority.
Server Name Indication (sni)	The hostname of the server that the client is trying to contact. Using SNIs enables a server to host multiple websites and present multiple certificates on the same IP address and TCP port because each website has a unique SNI.
Error (error)	A string showing the error that has occurred in the event.
Container ID (container_id)	A unique alphanumeric string that identifies the container if the firewall runs in a cloud container.
POD Namespace (pod_namespace)	The name of the Kubernetes pod namespace.
POD Name (pod_name)	The name of the kubernetes pod.
Source External Dynamic List (src_edl)	The name of the external dynamic list that contains the source IP address of the traffic.
Destination External Dynamic List (dst_edl)	The name of the external dynamic list that contains the destination IP address of the traffic.
Source Dynamic Address Group (src_dag)	The dynamic address group that Device-ID identifies as the source of the traffic.
Destination Dynamic Address Group (dst_dag)	The dynamic address group that Device-ID identifies as the destination for the traffic.
High Resolution Timestamp (high_res_timestamp)	Time in milliseconds the log was received at the management plane. The format for this field is YYYY-MM-DDThh:ss:sssTZD: • **YYYY**—Four digit year • **MM**—Two-digit month

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Field Name	Description
	• **DD**—Two-digit day of the month (01 through 31) • **T**—Indicator for the beginning of the timestamp • **hh**—Two-digit hour using 24-hour time (00 through 23) • **mm**—Two-digit minute (00 through 59) • **ss**—Two-digit second (00 through 60) • **sss**—One or more digits for millisecond • **TZD**—Time zone designator (+hh:mm or -hh:mm) *The High Resolution Timestamp is supported for logs received from managed firewalls running PAN-OS 10.1 and later releases. Logs received from managed firewalls running PAN-OS 9.1 and earlier releases display a `1969-12-31T16:00:00:000-8:00` timestamp regardless of when the log was received.*
Source Device Category (src_category)	The category for the device that Device-ID identifies as the source of the traffic.
Source Device Profile (src_profile)	The device profile for the device that Device-ID identifies as the source of the traffic.
Source Device Model (src_model)	The model of the device that Device-ID identifies as the source of the traffic.
Source Device Vendor (src_vendor)	The vendor of the device that Device-ID identifies as the source of the traffic.
Source Device OS Family (src_osfamily)	The operating system type for the device that Device-ID identifies as the source of the traffic.
Source Device OS Version (src_osversion)	The version of the operating system for the device that Device-ID identifies as the source of the traffic.
Source Hostname (src_host)	The hostname of the device that Device-ID identifies as the source of the traffic.
Source MAC Address (src_mac)	The MAC address for the device that Device-ID identifies as the source of the traffic.
Destination Device Category (dst_category)	The category for the device that Device-ID identifies as the destination for the traffic.

https://t.me/learningnets

Field Name	Description
Destination Device Profile (dst_profile)	The device profile for the device that Device-ID identifies as the destination for the traffic.
Destination Device Model (dst_model)	The model of the device that Device-ID identifies as the destination for the traffic.
Destination Device Vendor (dst_vendor)	The vendor of the device that Device-ID identifies as the destination for the traffic.
Destination Device OS Family (dst_osfamily)	The operating system type for the device that Device-ID identifies as the destination for the traffic.
Destination Device OS Version (dst_osversion)	The version of the operating system for the device that Device-ID identifies as the destination for the traffic.
Destination Hostname (dst_host)	The hostname of the device that Device-ID identifies as the destination for the traffic.
Destination MAC Address (dst_mac)	The MAC address for the device that Device-ID identifies as the destination for the traffic.
Sequence Number (seqno)	A 64-bit log entry identifier incremented sequentially; each log type has unique number space.
Action Flags (actionflags)	A bit field indicating if the log was forwarded to Panorama.
Device Group Hierarchy (dg_hier_level_1 to dg_hier_level_4)	A sequence of identification numbers that indicate the device group's location within a device group hierarchy. The firewall (or virtual system) generating the log includes the identification number of each ancestor in its device group hierarchy. The shared device group (level 0) is not included in this structure. If the log values are 12, 34, 45, 0, it means that the log was generated by a firewall (or virtual system) that belongs to device group 45, and its ancestors are 34, and 12. To view the device group names that correspond to the value 12, 34 or 45, use one of the following methods: API query: `/api/?type=op&cmd=<show><dg-hierarchy></dg-hierarchy></show>`

https://t.me/learningnets

Field Name	Description
Virtual System Name (vsys_name)	The name of the virtual system associated with the session; only valid on firewalls enabled for multiple virtual systems.
Device Name (device_name)	The hostname of the firewall on which the session was logged.
Virtual System ID (vsys_id)	A unique identifier for a virtual system on a Palo Alto Networks firewall.
Application Subcategory (subcategory_of_app)	The application subcategory specified in the application configuration properties.
Application Category (category_of_app)	The application category specified in the application configuration properties. Values are: • business-systems • collaboration • general-internet • media • networking • saas
Application Technology (technology_of_app)	The application technology specified in the application configuration properties. Values are: • browser-based • client-server • network-protocol • peer-to-peer
Application Risk (risk_of_app)	Risk level associated with the application (1=lowest to 5=highest).
Application Characteristic (characteristic_of_app)	Comma-separated list of applicable characteristic of the application
Application Container (container_of_app)	The parent application for an application.
Application SaaS (is_saas_of_app)	Displays `1` if a SaaS application or `0` if not a SaaS application.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Field Name	Description
Application Sanctioned State (sanctioned_state_of_app)	Displays 1 if application is sanctioned or 0 if application is not sanctioned.

Tunnel Inspection Log Fields

Format: FUTURE_USE, Receive Time, Serial Number, Type, Subtype, FUTURE_USE, Generated Time, Source Address, Destination Address, NAT Source IP, NAT Destination IP, Rule Name, Source User, Destination User, Application, Virtual System, Source Zone, Destination Zone, Inbound Interface, Outbound Interface, Log Action, FUTURE_USE, Session ID, Repeat Count, Source Port, Destination Port, NAT Source Port, NAT Destination Port, Flags, Protocol, Action, Severity, Sequence Number, Action Flags, Source Location, Destination Location, Device Group Hierarchy Level 1, Device Group Hierarchy Level 2, Device Group Hierarchy Level 3, Device Group Hierarchy Level 4, Virtual System Name, Device Name, Tunnel ID/IMSI, Monitor Tag/IMEI, Parent Session ID, Parent Start Time, Tunnel, Bytes, Bytes Sent, Bytes Received, Packets, Packets Sent, Packets Received, Maximum Encapsulation, Unknown Protocol, Strict Check, Tunnel Fragment, Sessions Created, Sessions Closed, Session End Reason, Action Source, Start Time, Elapsed Time, Tunnel Inspection Rule, Remote User IP, Remote User ID, Rule UUID, PCAP ID, Dynamic User Group, Source External Dynamic List, Destination External Dynamic List, High Resolution Timestamp, A Slice Differentiator, A Slice Service Type, PDU Session ID, Application Subcategory, Application Category, Application Technology, Application Risk, Application Characteristic, Application Container, Application SaaS, Application Sanctioned State

Field Name	Description
Receive Time (receive_time or cef-formatted-receive_time)	Month, day, and time the log was received at the management plane.
Serial Number (serial)	Serial number of the firewall that generated the log.
Type (type)	Type of log as it pertains to the session: START or END.
Threat/Content Type (subtype)	Subtype of traffic log; values are start, end, drop, and deny • Start—session started • End—session ended • Drop—session dropped before the application is identified and there is no rule that allows the session. • Deny—session dropped after the application is identified and there is a rule to block or no rule that allows the session.
Generated Time (time_generated	Time the log was generated on the dataplane.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Field Name	Description
or cef-formatted-time_generated)	
Source Address (src)	Source IP address of packets in the session.
Destination Address (dst)	Destination IP address of packets in the session.
NAT Source IP (natsrc)	If Source NAT performed, the post-NAT Source IP address.
NAT Destination IP (natdst)	If Destination NAT performed, the post-NAT Destination IP address.
Rule Name (rule)	Name of the Security policy rule in effect on the session.
Source User (srcuser)	Source User ID of packets in the session.
Destination User (dstuser)	Destination User ID of packets in the session.
Application (app)	Tunneling protocol used in the session.
Virtual System (vsys)	Virtual System associated with the session.
Source Zone (from)	Source zone of packets in the session.
Destination Zone (to)	Destination zone of packets in the session.
Inbound Interface (inbound_if)	Interface that the session was sourced from.
Outbound Interface (outbound_if)	Interface that the session was destined to.
Log Action (logset)	Log Forwarding Profile that was applied to the session.
Session ID (sessionid)	Session ID of the session being logged.
Repeat Count (repeatcnt)	Number of sessions with same Source IP, Destination IP, Application, and Subtype seen within 5 seconds.
Source Port (sport)	Source port utilized by the session.
Destination Port (dport)	Destination port utilized by the session.

https://t.me/learningnets

Field Name	Description
NAT Source Port (natsport)	Post-NAT source port.
NAT Destination Port (natdport)	Post-NAT destination port.
Flags (flags)	32-bit field that provides details on session; this field can be decoded by AND-ing the values with the logged value: • 0x80000000 —session has a packet capture (PCAP) • 0x02000000 —IPv6 session • 0x01000000 —SSL session was decrypted (SSL Proxy) • 0x00800000 —session was denied via URL filtering • 0x00400000 —session has a NAT translation performed (NAT) • 0x00200000 —user information for the session was captured through Authentication Portal • 0x00080000 —X-Forwarded-For value from a proxy is in the source user field • 0x00040000 —log corresponds to a transaction within a http proxy session (Proxy Transaction) • 0x00008000 —session is a container page access (Container Page) • 0x00002000 —session has a temporary match on a rule for implicit application dependency handling. Available in PAN-OS 5.0.0 and above. • 0x00000800 —symmetric return was used to forward traffic for this session
IP Protocol (proto)	IP protocol associated with the session.
Action (action)	Action taken for the session; possible values are: • Allow—session was allowed by policy • Deny—session was denied by policy • Drop—session was dropped silently • Drop ICMP—session was silently dropped with an ICMP unreachable message to the host or application • Reset both—session was terminated and a TCP reset is sent to both the sides of the connection • Reset client—session was terminated and a TCP reset is sent to the client • Reset server—session was terminated and a TCP reset is sent to the server

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Field Name	Description
Severity (severity)	Severity associated with the event; values are informational, low, medium, high, critical.
Sequence Number (seqno)	A 64-bit log entry identifier incremented sequentially; each log type has a unique number space. This field is not supported on PA-7000 Series firewalls.
Action Flags (actionflags)	A bit field indicating if the log was forwarded to Panorama.
Source Location (srcloc)	Source country or Internal region for private addresses; maximum length is 32 bytes.
Destination Location (dstloc)	Destination country or Internal region for private addresses. Maximum length is 32 bytes.
Device Group Hierarchy (dg_hier_level_1 to dg_hier_level_4)	A sequence of identification numbers that indicate the device group's location within a device group hierarchy. The firewall (or virtual system) generating the log includes the identification number of each ancestor in its device group hierarchy. The shared device group (level 0) is not included in this structure. If the log values are 12, 34, 45, 0, it means that the log was generated by a firewall (or virtual system) that belongs to device group 45, and its ancestors are 34, and 12. To view the device group names that correspond to the value 12, 34 or 45, use one of the following methods: API query: `/api/?type=op&cmd=<show><dg-hierarchy></dg-hierarchy></show>`
Virtual System Name (vsys_name)	The name of the virtual system associated with the session; only valid on firewalls enabled for multiple virtual systems.
Device Name (device_name)	The hostname of the firewall on which the session was logged.
Tunnel ID (tunnelid)	ID of the tunnel being inspected or the International Mobile Subscriber Identity (IMSI) ID of the mobile user.
Monitor Tag (monitortag)	Monitor name you configured for the Tunnel Inspection policy rule or the International Mobile Equipment Identity (IMEI) ID of the mobile device.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Field Name	Description
Parent Session ID (parent_session_id)	ID of the session in which this session is tunneled. Applies to inner tunnel (if two levels of tunneling) or inside content (if one level of tunneling) only.
Parent Start Time (parent_start_time)	Year/month/day hours:minutes:seconds that the parent tunnel session began.
Tunnel Type (tunnel)	Type of tunnel, such as GRE or IPSec.
Bytes (bytes)	Number of bytes in the session.
Bytes Sent (bytes_sent)	Number of bytes in the client-to-server direction of the session.
Bytes Received (bytes_received)	Number of bytes in the server-to-client direction of the session.
Packets (packets)	Number of total packets (transmit and receive) for the session.
Packets Sent (pkts_sent)	Number of client-to-server packets for the session.
Packets Received (pkts_received)	Number of server-to-client packets for the session.
Maximum Encapsulation (max_encap)	Number of packets the firewall dropped because the packet exceeded the maximum number of encapsulation levels configured in the Tunnel Inspection policy rule (Drop packet if over maximum tunnel inspection level).
Unknown Protocol (unknown_proto)	Number of packets the firewall dropped because the packet contains an unknown protocol, as enabled in the Tunnel Inspection policy rule (Drop packet if unknown protocol inside tunnel).
Strict Checking (strict_check)	Number of packets the firewall dropped because the tunnel protocol header in the packet failed to comply with the RFC for the tunnel protocol, as enabled in the Tunnel Inspection policy rule (**Drop packet if tunnel protocol fails strict header check**).
Tunnel Fragment (tunnel_fragment)	Number of packets the firewall dropped because of fragmentation errors.
Sessions Created (sessions_created)	Number of inner sessions created.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Field Name	Description
Sessions Closed (sessions_closed)	Number of completed/closed sessions created.
Session End Reason (session_end_reason)	The reason a session terminated. If the termination had multiple causes, this field displays only the highest priority reason. The possible session end reason values are as follows, in order of priority (where the first is highest): • threat—The firewall detected a threat associated with a reset, drop, or block (IP address) action. • policy-deny—The session matched a security rule with a deny or drop action. • decrypt-cert-validation—The session terminated because you configured the firewall to block SSL forward proxy decryption or SSL inbound inspection when the session uses client authentication or when the session uses a server certificate with any of the following conditions: expired, untrusted issuer, unknown status, or status verification time-out. This session end reason also displays when the server certificate produces a fatal error alert of type bad_certificate, unsupported_certificate, certificate_revoked, access_denied, or no_certificate_RESERVED (SSLv3 only). • decrypt-unsupport-param—The session terminated because you configured the firewall to block SSL forward proxy decryption or SSL inbound inspection when the session uses an unsupported protocol version, cipher, or SSH algorithm. This session end reason is displays when the session produces a fatal error alert of type unsupported_extension, unexpected_message, or handshake_failure. • decrypt-error—The session terminated because you configured the firewall to block SSL forward proxy decryption or SSL inbound inspection when firewall resources or the hardware security module (HSM) were unavailable. This session end reason is also displayed when you configured the firewall to block SSL traffic that has SSH errors or that produced any fatal error alert other than those listed for the decrypt-cert-validation and decrypt-unsupport-param end reasons. • tcp-rst-from-client—The client sent a TCP reset to the server. • tcp-rst-from-server—The server sent a TCP reset to the client. • resources-unavailable—The session dropped because of a system resource limitation. For example, the session could have exceeded the number of out-of-order packets allowed per flow or the global out-of-order packet queue. • tcp-fin—One host or both hosts in the connection sent a TCP FIN message to close the session.

https://t.me/learningnets

Field Name	Description
	• tcp-reuse—A session is reused and the firewall closes the previous session. • decoder—The decoder detects a new connection within the protocol (such as HTTP-Proxy) and ends the previous connection. • aged-out—The session aged out. • unknown—This value applies in the following situations: • Session terminations that the preceding reasons do not cover (for example, a `clear session all` command). • For logs generated in a PAN-OS release that does not support the session end reason field (releases older than PAN-OS 6.1), the value will be `unknown` after an upgrade to the current PAN-OS release or after the logs are loaded onto the firewall. • In Panorama, logs received from firewalls for which the PAN-OS version does not support session end reasons will have a value of `unknown`. • n/a—This value applies when the traffic log type is not **end**.
Action Source (action_source)	Specifies whether the action taken to allow or block an application was defined in the application or in policy. The actions can be allow, deny, drop, reset- server, reset-client or reset-both for the session.
Start Time (start)	Year/month/day hours:minutes:seconds that the session began.
Elapsed Time (elapsed)	Elapsed time of the session.
Tunnel Inspection Rule (tunnel_insp_rule)	Name of the tunnel inspection rule matching the cleartext tunnel traffic.
Remote User IP (remote_user_ip)	IPv4 or IPv6 address of a remote user.
Remote User ID (remote_user_id)	IMSI identity of a remote user, and if available, one IMEI identity or one MSISDN identity.
Security Rule UUID (rule_uuid)	The UUID that permanently identifies the rule.
PCAP ID (pcap_id)	Unique packet capture ID that defines the location of the pcap file on the firewall.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Field Name	Description
Dynamic User Group Name (dynusergroup_name)	The name of the dynamic user group that contains the user who initiated the session.
Source External Dynamic List (src_edl)	The name of the external dynamic list that contains the source IP address of the traffic.
Destination External Dynamic List (dst_edl)	The name of the external dynamic list that contains the destination IP address of the traffic.
High Resolution Timestamp (high_res timestamp)	Time in milliseconds the log was received at the management plane. The format for this new field is YYYY-MM-DDThh:ss:sssTZD: • **YYYY**—Four digit year • **MM**—Two-digit month • **DD**—Two-digit day of the month (01 through 31) • **T**—Indicator for the beginning of the timestamp • **hh**—Two-digit hour using 24-hour time (00 through 23) • **mm**—Two-digit minute (00 through 59) • **ss**—Two-digit second (00 through 60) • **sss**—One or more digits for millisecond • **TZD**—Time zone designator (+hh:mm or -hh:mm) *The High Resolution Timestamp is supported for logs received from managed firewalls running PAN-OS 10.1 and later releases. Logs received from managed firewalls running PAN-OS 9.1 and earlier releases display a `1969-12-31T16:00:00:000-8:00` timestamp regardless of when the log was received.*
A Slice Differentiator (nssai_sd)	The A Slice Differentiator of the Network Slice ID.
A Slice Service Type (nssai_sd)	The A Slice Service Type of the Network Slice ID.
PDU Session ID (pdu_session_id)	Session ID for the collection of L4 segments inside a tunnel.
Application Subcategory (subcategory_of_app)	The application subcategory specified in the application configuration properties.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Field Name	Description
Application Category (category_of_app)	The application category specified in the application configuration properties. Values are: • business-systems • collaboration • general-internet • media • networking • saas
Application Technology (technology_of_app)	The application technology specified in the application configuration properties. Values are: • browser-based • client-server • network-protocol • peer-to-peer
Application Risk (risk_of_app)	Risk level associated with the application (1=lowest to 5=highest).
Application Characteristic (characteristic_of_app)	Comma-separated list of applicable characteristic of the application
Application Container (container_of_app)	The parent application for an application.
Application SaaS (is_saas_of_app)	Displays `1` if a SaaS application or `0` if not a SaaS application.
Application Sanctioned State (sanctioned_state_of_app)	Displays `1` if application is sanctioned or `0` if application is not sanctioned.

SCTP Log Fields

Format: FUTURE_USE, Receive Time, Serial Number, Type, FUTURE_USE, FUTURE_USE, Generated Time, Source Address, Destination Address, FUTURE_USE, FUTURE_USE, Rule Name, FUTURE_USE, FUTURE_USE, FUTURE_USE, Virtual System, Source Zone, Destination Zone, Inbound Interface, Outbound Interface, Log Action, FUTURE_USE, Session ID, Repeat Count, Source Port, Destination Port, FUTURE_USE, FUTURE_USE, FUTURE_USE, FUTURE_USE, IP Protocol, Action, Device Group Hierarchy Level 1, Device Group Hierarchy Level 2, Device Group Hierarchy Level 3, Device Group Hierarchy Level 4, Virtual System Name, Device Name, Sequence Number, FUTURE_USE, SCTP Association ID, Payload Protocol ID, Severity, SCTP Chunk Type,

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

FUTURE_USE, SCTP Verification Tag 1, SCTP Verification Tag 2, SCTP Cause Code, Diameter App ID, Diameter Command Code, Diameter AVP Code, SCTP Stream ID, SCTP Association End Reason, Op Code, SCCP Calling Party SSN, SCCP Calling Party Global Title, SCTP Filter, SCTP Chunks, SCTP Chunks Sent, SCTP Chunks Received, Packets, Packets Sent, Packets Received, UUID for rule, High Resolution Timestamp

Field Name	Description
Receive Time (receive_time or cef-formatted-receive_time)	Time the log was received at the management plane.
Serial Number (serial)	Serial number of the firewall that generated the log.
Type (type)	Specifies the type of log; value is SCTP.
Generated Time (time_generated or cef-formatted-time_generated)	Time the log was generated on the dataplane.
Source Address (src)	Original session source IP address.
Destination Address (dst)	Original session destination IP address.
Rule Name (rule)	Name of the Security policy rule in effect on the session.
Virtual System (vsys)	Virtual System associated with the session.
Source Zone (from)	Zone the session was sourced from.
Destination Zone (to)	Zone the session was destined to.
Inbound Interface (inbound_if)	Interface that the session was sourced from.
Outbound Interface (outbound_if)	Interface that the session was destined to.
Log Action (logset)	Log Forwarding Profile that was applied to the session.
Session ID (sessionid)	An internal numerical identifier applied to each session.
Repeat Count (repeatcnt)	Number of sessions with same Source IP, Destination IP, Application, and Subtype seen within 5 seconds.
Source Port (sport)	Source port utilized by the session.
Destination Port (dport)	Destination port utilized by the session.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Field Name	Description
IP Protocol (proto)	IP protocol associated with the session.
Action (action)	Action taken for the session; possible values are: • allow—session was allowed by the policy • deny—session was denied by the policy
Device Group Hierarchy (dg_hier_level_1 to dg_hier_level_4)	A sequence of identification numbers that indicate the device group's location within a device group hierarchy. The firewall (or virtual system) generating the log includes the identification number of each ancestor in its device group hierarchy. The shared device group (level 0) is not included in this structure. If the log values are 12, 34, 45, 0, it means that the log was generated by a firewall (or virtual system) that belongs to device group 45, and its ancestors are 34, and 12. To view the device group names that correspond to the value 12, 34 or 45, use one of the following methods: API query: `/api/?type=op&cmd=<show><dg-hierarchy></dg-hierarchy></show>`
Virtual System Name (vsys_name)	The name of the virtual system associated with the session; only valid on firewalls enabled for multiple virtual systems.
Device Name (device_name)	The hostname of the firewall on which the session was logged.
Sequence Number (seqno)	A 64-bit log entry identifier incremented sequentially; each log type has a unique number space.
SCTP Association ID (assoc_id)	An internal 56-bit numerical logical identifier applied to each SCTP association.
Payload Protocol ID (ppid)	Identifies the Payload Protocol ID (PPID) in the data chunk which triggered this event. PPID is assigned by Internet Assigned Numbers Authority (IANA).
Severity (severity)	Severity associated with the event; values are informational, low, medium, high, critical.
SCTP Chunk Type (sctp_chunk_type)	Describes the type of information contained in a chunk, such as control or data.

https://t.me/learningnets

Field Name	Description
SCTP Event Type (sctp_event_type)	Defines the event triggered per SCTP chunk or packet when SCTP protection profile is applied to the SCTP traffic. It is also triggered by start or end of a SCTP association.
SCTP Verification Tag 1 (verif_tag_1)	Used by endpoint1 which initiates the association to verify if the SCTP packet received belongs to current SCTP association and validate the endpoint2.
SCTP Verification Tag 2 (verif_tag_2)	Used by endpoint2 to verify if the SCTP packet received belongs to current SCTP association and validate the endpoint1.
SCTP Cause Code (sctp_cause_code)	Sent by an endpoint to specify reason for an error condition to other endpoint of same SCTP association.
Diameter App ID (diam_app_id)	The diameter application in the data chunk which triggered the event. Diameter Application ID is assigned by Internet Assigned Numbers Authority (IANA).
Diameter Command Code (diam_cmd_code)	The diameter command code in the data chunk which triggered the event. Diameter Command Code is assigned by Internet Assigned Numbers Authority (IANA)
Diameter AVP Code (diam_avp_code)	The diameter AVP code in the data chunk which triggered the event.
SCTP Stream ID (stream_id)	ID of the stream which carries the data chunk which triggered the event.
SCTP Association End Reason (assoc_end_reason)	Reason an association was terminated. If the termination had multiple causes, the highest priority reason is displayed. The possible session end reasons in descending priority are: • shutdown-from-endpoint (highest)—endpoint sends out SHUTDOWN • abort-from-endpoint—endpoint sends out ABORT • unknown (lowest)—the association aged out, or association termination reason is not covered by one of the previous reasons (for example, a clear session all command).
Op Code (op_code)	Identifies the operation code of application layer SS7 protocols, like MAP or CAP, in the data chunk which triggered the event.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Field Name	Description
SCCP Calling Party SSN (sccp_calling_ssn)	The Signaling Connection Control Part (SCCP) calling party subsystem number (SSN) in the data chunk which triggered the event.
SCCP Calling Party Global Title (sccp_calling_gt)	The Signaling Connection Control Part (SCCP) calling party global title (GT) in the data chunk which triggered the event.
SCTP Filter (sctp_filter)	Name of the filter that the SCTP chunk matched.
SCTP Chunks (chunks)	Number of total chunks (transmit and receive) for the association.
SCTP Chunks Sent (chunks_sent)	Number of endpoint1(which initiates association)-to-endpoint2 chunks for the association.
SCTP Chunks Received (chunks_received)	Number of endpoint2-to-endpoint1(which initiates association) chunks for the association.
Packets (packets)	Number of total packets (transmit and receive) for the session.
Packets Sent (pkts_sent)	Number of client-to-server packets for the session.
Packets Received (pkts_received)	Number of server-to-client packets for the session.
UUID for rule (rule_uuid)	The UUID that permanently identifies the rule.
High Resolution Timestamp (high_res_timestamp)	Time in milliseconds the log was received at the management plane. The format for this new field is YYYY-MM-DDThh:ss:sssTZD: • **YYYY**—Four digit year • **MM**—Two-digit month • **DD**—Two-digit day of the month (01 through 31) • **T**—Indicator for the beginning of the timestamp • **hh**—Two-digit hour using 24-hour time (00 through 23) • **mm**—Two-digit minute (00 through 59) • **ss**—Two-digit second (00 through 60) • **sss**—One or more digits for millisecond • **TZD**—Time zone designator (+hh:mm or -hh:mm)

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Field Name	Description
	The High Resolution Timestamp is supported for logs received from managed firewalls running PAN-OS 10.1 and later releases. Logs received from managed firewalls running PAN-OS 9.1 and earlier releases display a `1969-12-31T16:00:00:000-8:00` *timestamp regardless of when the log was received.*

Authentication Log Fields

Format: FUTURE_USE, Receive Time, Serial Number, Type, Threat/Content Type, FUTURE_USE, Generated Time, Virtual System, Source IP, User, Normalize User, Object, Authentication Policy, Repeat Count, Authentication ID, Vendor, Log Action, Server Profile, Description, Client Type, Event Type, Factor Number, Sequence Number, Action Flags, Device Group Hierarchy 1, Device Group Hierarchy 2, Device Group Hierarchy 3, Device Group Hierarchy 4, Virtual System Name, Device Name, Virtual System ID, Authentication Protocol, UUID for rule, High Resolution Timestamp, Source Device Category, Source Device Profile, Source Device Model, Source Device Vendor, Source Device OS Family, Source Device OS Version, Source Hostname, Source Mac Address, Region, FUTURE_USE, User Agent, Session ID

Field Name	Description
Receive Time (receive_time or cef-formatted-receive_time)	Time the log was received at the management plane.
Serial Number (serial)	Serial number of the device that generated the log.
Type (type)	Specifies the type of log; value is AUTHENTICATION.
Threat/Content Type (subtype)	Subtype of the system log; refers to the system daemon generating the log; values are crypto, dhcp, dnsproxy, dos, general, global-protect, ha, hw, nat, ntpd, pbf, port, pppoe, ras, routing, satd, sslmgr, sslvpn, userid, url-filtering, vpn.
Generated Time (time_generated or cef-formatted-time_generated)	Time the log was generated on the dataplane.
Virtual System (vsys)	Virtual System associated with the session.
Source IP (ip)	Original session source IP address.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Field Name	Description
User (user)	End user being authenticated.
Normalize User (normalize_user)	Normalized version of username being authenticated (such as appending a domain name to the username).
Object (object)	Name of the object associated with the system event.
Authentication Policy (authpolicy)	Policy invoked for authentication before allowing access to a protected resource.
Repeat Count (repeatcnt)	Number of sessions with same Source IP, Destination IP, Application, and Subtype seen within 5 seconds.
Authentication ID (authid)	Unique ID given across primary authentication and additional (multi factor) authentication.
Vendor (vendor)	Vendor providing additional factor authentication.
Log Action (logset)	Log Forwarding Profile that was applied to the session.
Server Profile (serverprofile)	Authentication server used for authentication.
Description (desc)	Additional authentication information.
Client Type (clienttype)	Type of client used to complete authentication (such as authentication portal).
Event Type (event)	Result of the authentication attempt.
Factor Number (factorno)	Indicates the use of primary authentication (1) or additional factors (2, 3).
Sequence Number (seqno)	A 64-bit log entry identifier incremented sequentially. Each log type has a unique number space.
Action Flags (actionflags)	A bit field indicating if the log was forwarded to Panorama.
Device Group Hierarchy (dg_hier_level_1 to dg_hier_level_4)	A sequence of identification numbers that indicate the device group's location within a device group hierarchy. The firewall (or virtual system) generating the log includes the identification number of each ancestor in its device group hierarchy. The shared device group (level 0) is not included in this structure. If the log values are 12, 34, 45, 0, it means that the log was generated by a firewall (or virtual system) that belongs to device group 45,

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Field Name	Description
	and its ancestors are 34, and 12. To view the device group names that correspond to the value 12, 34 or 45, use one of the following methods: API query: `/api/?type=op&cmd=<show><dg-hierarchy></dg-hierarchy></show>`
Virtual System Name (vsys_name)	The name of the virtual system associated with the session; only valid on firewalls enabled for multiple virtual systems.
Device Name (device_name)	The hostname of the firewall on which the session was logged.
Virtual System ID (vsys_id)	A unique identifier for a virtual system on a Palo Alto Networks firewall.
Authentication Protocol (authproto)	Indicates the authentication protocol used by the server. For example, PEAP with GTC.
UUID for rule (rule_uuid)	The UUID that permanently identifies the rule.
High Resolution Timestamp (high_res_timestamp)	Time in milliseconds the log was received at the management plane. The format for this new field is YYYY-MM-DDThh:ss:sssTZD: • **YYYY**—Four digit year • **MM**—Two-digit month • **DD**—Two-digit day of the month (01 through 31) • **T**—Indicator for the beginning of the timestamp • **hh**—Two-digit hour using 24-hour time (00 through 23) • **mm**—Two-digit minute (00 through 59) • **ss**—Two-digit second (00 through 60) • **sss**—One or more digits for millisecond • **TZD**—Time zone designator (+hh:mm or -hh:mm) *The High Resolution Timestamp is supported for logs received from managed firewalls running PAN-OS 10.1 and later releases. Logs received from managed firewalls running PAN-OS 9.1 and earlier releases display a `1969-12-31T16:00:00:000-8:00` timestamp regardless of when the log was received.*

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Field Name	Description
Source Device Category (src_category)	The category for the device that Device-ID identifies as the source of the traffic.
Source Device Profile (src_profile)	The device profile for the device that Device-ID identifies as the source of the traffic.
Source Device Model (src_model)	The model of the device that Device-ID identifies as the source of the traffic.
Source Device Vendor (src_vendor)	The vendor of the device that Device-ID identifies as the source of the traffic.
Source Device OS Family (src_osfamily)	The operating system type for the device that Device-ID identifies as the source of the traffic.
Source Device OS Version (src_osversion)	The version of the operating system for the device that Device-ID identifies as the source of the traffic.
Source Hostname (src_host)	The hostname of the device that Device-ID identifies as the source of the traffic.
Source MAC Address (src_mac)	The MAC address for the device that Device-ID identifies as the source of the traffic.
Region (region)	The geographical region where the traffic originates.
User Agent (user_agent)	The string from the HTTP request header `User-Agent`.
Session ID	A string that uniquely identifies the traffic session.

Config Log Fields

Format: FUTURE_USE, Receive Time, Serial Number, Type, Subtype, FUTURE_USE, Generated Time, Host, Virtual System, Command, Admin, Client, Result, Configuration Path, Before Change Detail, After Change Detail, Sequence Number, Action Flags, Device Group Hierarchy Level 1, Device Group Hierarchy Level 2, Device Group Hierarchy Level 3, Device Group Hierarchy Level 4, Virtual System Name, Device Name, Device Group, Audit Comment

Field Name	Description
Receive Time (receive_time or	Time the log was received at the management plane.

https://t.me/learningnets

Field Name	Description
cef-formatted-receive_time)	
Serial Number (serial)	Serial number of the device that generated the log.
Type (type)	Specifies the type of log; value is CONFIG.
Threat/Content Type (subtype)	Subtype of the configuration log; unused.
Generated Time (time_generated or cef-formatted-time_generated)	Time the log was generated on the dataplane.
Host (host)	Hostname or IP address of the client machine
Virtual System (vsys)	Virtual System associated with the configuration log
Command (cmd)	Command performed by the Admin; values are add, clone, commit, delete, edit, move, rename, set.
Admin (admin)	Username of the Administrator performing the configuration
Client (client)	Client used by the Administrator; values are Web and CLI
Result (result)	Result of the configuration action; values are Submitted, Succeeded, Failed, and Unauthorized
Configuration Path (path)	The path of the configuration command issued; up to 512 bytes in length
Before Change Detail (before_change_detail)	This field is in custom logs only; it is not in the default format. It contains the full xpath before the configuration change.
After Change Detail (after_change_detail)	This field is in custom logs only; it is not in the default format. It contains the full xpath after the configuration change.
Sequence Number (seqno)	A 64bit log entry identifier incremented sequentially; each log type has a unique number space.
Action Flags (actionflags)	A bit field indicating if the log was forwarded to Panorama.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Field Name	Description
Device Group Hierarchy (dg_hier_level_1 to dg_hier_level_4)	A sequence of identification numbers that indicate the device group's location within a device group hierarchy. The firewall (or virtual system) generating the log includes the identification number of each ancestor in its device group hierarchy. The shared device group (level 0) is not included in this structure. If the log values are 12, 34, 45, 0, it means that the log was generated by a firewall (or virtual system) that belongs to device group 45, and its ancestors are 34, and 12. To view the device group names that correspond to the value 12, 34 or 45, use one of the following methods: API query: `/api/?type=op&cmd=<show><dg-hierarchy></dg-hierarchy></show>`
Virtual System Name (vsys_name)	The name of the virtual system associated with the session; only valid on firewalls enabled for multiple virtual systems.
Device Name (device_name)	The hostname of the firewall on which the session was logged.
Device Group (dg_id)	The device group the firewall belongs to if managed by a Panorama™ management server.
Audit Comment (comment)	The audit comment entered in a policy rule configuration change.

System Log Fields

Format: FUTURE_USE, Receive Time, Serial Number, Type, Content/Threat Type, FUTURE_USE, Generated Time, Virtual System, Event ID, Object, FUTURE_USE, FUTURE_USE, Module, Severity, Description, Sequence Number, Action Flags, Device Group Hierarchy Level 1, Device Group Hierarchy Level 2, Device Group Hierarchy Level 3, Device Group Hierarchy Level 4, Virtual System Name, Device Name, FUTURE_USE, FUTURE_USE, High Resolution Timestamp

Field Name	Description
Receive Time (receive_time or cef-formatted-receive_time)	Time the log was received at the management plane.
Serial Number (serial)	Serial number of the firewall that generated the log.
Type (type)	Specifies the type of log; value is SYSTEM.

https://t.me/learningnets

Field Name	Description
Content/Threat Type (subtype)	Subtype of the system log; refers to the system daemon generating the log; values are crypto, dhcp, dnsproxy, dos, general, global-protect, ha, hw, nat, ntpd, pbf, port, pppoe, ras, routing, satd, sslmgr, sslvpn, userid, url-filtering, vpn.
Generated Time (time_generated or cef-formatted-time_generated)	Time the log was generated on the dataplane.
Virtual System (vsys)	Virtual System associated with the configuration log.
Event ID (eventid)	String showing the name of the event.
Object (object)	Name of the object associated with the system event.
Module (module)	This field is valid only when the value of the Subtype field is general. It provides additional information about the sub-system generating the log; values are general, management, auth, ha, upgrade, chassis.
Severity (severity)	Severity associated with the event; values are informational, low, medium, high, critical.
Description (opaque)	Detailed description of the event, up to a maximum of 512 bytes.
Sequence Number (seqno)	A 64-bit log entry identifier incremented sequentially; each log type has a unique number space.
Action Flags (actionflags)	A bit field indicating if the log was forwarded to Panorama.
Device Group Hierarchy (dg_hier_level_1 to dg_hier_level_4)	A sequence of identification numbers that indicate the device group's location within a device group hierarchy. The firewall (or virtual system) generating the log includes the identification number of each ancestor in its device group hierarchy. The shared device group (level 0) is not included in this structure. If the log values are 12, 34, 45, 0, it means that the log was generated by a firewall (or virtual system) that belongs to device group 45, and its ancestors are 34, and 12. To view the device group names that correspond to the value 12, 34 or 45, use one of the following methods: API query: `/api/?type=op&cmd=<show><dg-hierarchy></dg-hierarchy></show>`

https://t.me/learningnets

Field Name	Description
Virtual System Name (vsys_name)	The name of the virtual system associated with the session; only valid on firewalls enabled for multiple virtual systems.
Device Name (device_name)	The hostname of the firewall on which the session was logged.
High Resolution Timestamp (high_res_timestamp)	Time in milliseconds the log was received at the management plane. The format for this new field is YYYY-MM-DDThh:ss:sssTZD: • **YYYY**—Four digit year • **MM**—Two-digit month • **DD**—Two-digit day of the month (01 through 31) • **T**—Indicator for the beginning of the timestamp • **hh**—Two-digit hour using 24-hour time (00 through 23) • **mm**—Two-digit minute (00 through 59) • **ss**—Two-digit second (00 through 60) • **sss**—One or more digits for millisecond • **TZD**—Time zone designator (+hh:mm or -hh:mm) *The High Resolution Timestamp is supported for logs received from managed firewalls running PAN-OS 10.1 and later releases. Logs received from managed firewalls running PAN-OS 9.1 and earlier releases display a `1969-12-31T16:00:00:000-8:00` timestamp regardless of when the log was received.*

Correlated Events Log Fields

Format: FUTURE_USE, Receive Time, Serial Number, Type, Content/Threat Type, FUTURE_USE, Generated Time, Source Address. Source User, Virtual System, Category, Severity, Device Group Hierarchy Level 1, Device Group Hierarchy Level 2, Device Group Hierarchy Level 3, Device Group Hierarchy Level 4, Virtual System Name, Device Name, Virtual System ID, Object Name, Object ID, Evidence

Field Name	Description
Receive Time (receive_time or cef-formatted-receive_time)	Time the log was received at the management plane.
Serial Number (serial)	Serial number of the device that generated the log.

https://t.me/learningnets

Field Name	Description
Type (type)	Specifies the type of log; value is CORRELATION.
Content/Threat Type (subtype)	Subtype of the system log; refers to the system daemon generating the log; values are crypto, dhcp, dnsproxy, dos, general, global-protect, ha, hw, nat, ntpd, pbf, port, pppoe, ras, routing, satd, sslmgr, sslvpn, userid, url-filtering, vpn.
Generated Time (time_generated or cef-formatted-time_generated)	Time the log was generated on the dataplane.
Source Address (src)	IP address of the user who initiated the event.
Source User (srcuser)	Username of the user who initiated the event.
Virtual System (vsys)	Virtual System associated with the configuration log.
Category (category)	A summary of the kind of threat or harm posed to the network, user, or host.
Severity (severity)	Severity associated with the event; values are informational, low, medium, high, critical.
Device Group Hierarchy (dg_hier_level_1 to dg_hier_level_4)	A sequence of identification numbers that indicate the device group's location within a device group hierarchy. The firewall (or virtual system) generating the log includes the identification number of each ancestor in its device group hierarchy. The shared device group (level 0) is not included in this structure. If the log values are 12, 34, 45, 0, it means that the log was generated by a firewall (or virtual system) that belongs to device group 45, and its ancestors are 34, and 12. To view the device group names that correspond to the value 12, 34 or 45, use one of the following methods: API query: `/api/?type=op&cmd=<show><dg-hierarchy></dg-hierarchy></show>`
Virtual System Name (vsys_name)	The name of the virtual system associated with the session; only valid on firewalls enabled for multiple virtual systems.
Device Name (device_name)	The hostname of the firewall on which the session was logged.

https://t.me/learningnets

Field Name	Description
Virtual System ID (vsys_id)	A unique identifier for a virtual system on a Palo Alto Networks firewall.
Object Name (objectname)	Name of the correlation object that was matched on.
Object ID (object_id)	Name of the object associated with the system event.
Evidence (evidence)	A summary statement that indicates how many times the host has matched against the conditions defined in the correlation object. For example, Host visited known malware URl (19 times).

GTP Log Fields

Format: FUTURE_USE, Receive Time, Serial Number, Type, Threat/Content Type, FUTURE_USE, Generated Time, Source Address, Destination Address, FUTURE_USE, FUTURE_USE, Rule Name, FUTURE_USE, FUTURE_USE, Application, Virtual System, Source Zone, Destination Zone, Inbound Interface, Outbound Interface, Log Action, FUTURE_USE, Session ID, FUTURE_USE, Source Port, Destination Port, FUTURE_USE, FUTURE_USE, FUTURE_USE, Protocol, Action, GTP Event Type, MSISDN, Access Point Name, Radio Access Technology, GTP Message Type, End User IP Address, Tunnel Endpoint Identifier1, Tunnel Endpoint Identifier2, GTP Interface, GTP Cause, Severity, Serving Country MCC, Serving Network MNC, Area Code, Cell ID, GTP Event Code, FUTURE_USE, FUTURE_USE, Source Location, Destination Location, FUTURE_USE, FUTURE_USE, FUTURE_USE, FUTURE_USE, FUTURE_USE, FUTURE_USE, FUTURE_USE, Tunnel ID/IMSI, Monitor Tag/IMEI, FUTURE_USE, FUTURE_USE, FUTURE_USE, FUTURE_USE, FUTURE_USE, FUTURE_USE, FUTURE_USE, FUTURE_USE, FUTURE_USE, FUTURE_USE, FUTURE_USE, FUTURE_USE, FUTURE_USE, FUTURE_USE, FUTURE_USE, FUTURE_USE, FUTURE_USE, FUTURE_USE, Start Time, Elapsed Time, Tunnel Inspection Rule, Remote User IP, Remote User ID, UUID for rule, PCAP ID, High Resolution Timestamp, A Slice Service Type, A Slice Differentiator, Application Subcategory, Application Category, Application Technology, Application Risk, Application Characteristic, Application Container, Application SaaS, Application Sanctioned State

Field Name	Description
Receive Time (receive_time or cef-formatted-receive_time)	Month, Day and time the log was received at the management plane.
Serial Number (serial)	Serial number of the firewall that generated the log.
Type (type)	Specifies the type of log; value is GTP.
Threat/Content Type (subtype)	Subtype of traffic log; values are start, end, drop, and deny • Start—session started • End—session ended

https://t.me/learningnets

Field Name	Description
	• Drop—session dropped before the application is identified and there is no rule that allows the session. • Deny—session dropped after the application is identified and there is a rule to block or no rule that allows the session.
Generated Time (time_generated or cef-formatted-time_generated)	Time the log was generated on the dataplane.
Source Address (src)	Source IP address of packets in the session.
Destination Address (dst)	Destination IP address of packets in the session.
Rule Name (rule)	Name of the Security policy rule in effect on the session.
Application (app)	Tunneling protocol used in the session.
Virtual System (vsys)	Virtual System associated with the session.
Source Zone (from)	Source zone of packets in the session.
Destination Zone (to)	Destination zone of packets in the session.
Inbound Interface (inbound_if)	Interface that the session was sourced from.
Outbound Interface (outbound_if)	Interface that the session was destined to.
Log Action (logset)	Log Forwarding Profile that was applied to the session.
Session ID (sessionid)	Session ID of the session being logged.
Source Port (sport)	Source port utilized by the session.
Destination Port (dport)	Destination port utilized by the session.
IP Protocol (proto)	IP protocol associated with the session.
Action (action)	Action taken for the session; possible values are: • allow—session was allowed by policy • deny—session was denied by policy

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Field Name	Description
GTP Event Type (event_type)	Defines event triggered by a GTP message when checks in GTP protection profile are applied to the GTP traffic. Also triggered by the start or end of a GTP session.
MSISDN (msisdn)	Service identity associated with the mobile subscriber composed of a Country Code, National Destination Code and a Subscriber. Consists of decimal digits (0-9) only with a maximum of 15 digits.
Access Point Name (apn)	Reference to a Packet Data Network Data Gateway (PGW)/ Gateway GPRS Support Node in a mobile network. Composed of a mandatory APN Network Identifier and an optional APN Operator Identifier.
Radio Access Technology (rat)	Type of technology used for radio access. For example, EUTRAN, WLAN, Virtual, HSPA Evolution, GAN and GERAN.
GTP Message Type (msg_type)	Indicates the GTP message type.
End IP Address (end_ip_adr)	IP address of a mobile subscriber allocated by a PGW/GGSN.
Tunnel Endpoint Identifier1 (teid1)	Identifies the GTP tunnel in the network node. TEID1 is the first TEID in the GTP message.
Tunnel Endpoint Identifier2 (teid2)	Identifies the GTP tunnel in the network node. TEID2 is the second TEID in the GTP message.
GTP Interface (gtp_interface)	3GPP interface from which a GTP message is received.
GTP Cause (cause_code)	GTP cause value in logs responses which contain an Information Element that provides information about acceptance or rejection of GTP requests by a network node.
Severity (severity)	Severity associated with the event; values are informational, low, medium, high, critical.
Serving Network MCC (mcc)	Mobile country code of serving core network operator.
Serving Network MNC (mnc)	Mobile network code of serving core network operator.
Area Code (area_code)	Area within a Public Land Mobile Network (PLMN).
Cell ID (cell_id)	Base station within an area code.
GTP Event Code (event_code)	Event code describing the GTP event.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Field Name	Description
Source Location (srcloc)	Source country or Internal region for private addresses; maximum length is 32 bytes.
Destination Location (dstloc)	Destination country or Internal region for private addresses; maximum length is 32 bytes.
Tunnel ID/IMSI (imsi)	International Mobile Subscriber Identity (IMSI) is a unique number allocated to each mobile subscriber in the GSM/UMTS/EPS system. IMSI shall consist of decimal digits (0 through 9) only and maximum number of digits allowed are 15.
Monitor Tag/IMEI (imei)	International Mobile Equipment Identity (IMEI) is a unique 15 or 16 digit number allocated to each mobile station equipment.
Start Time (start)	Time of session start.
Elapsed Time (elapsed)	Elapsed time of the session.
Tunnel Inspection Rule (tunnel_insp_rule)	Name of the tunnel inspection rule matching the cleartext tunnel traffic
Remote User IP (remote_user_ip)	IPv4 or IPv6 address used by a remote user.
Remote User ID (remote_user_id)	IMSI identity of a remote user, and if available, one IMEI identity and/or one MSISDN identity.
UUID for rule (rule_uuid)	Universally Unique ID for rule.
PCAP ID (pcap_id)	Unique packet capture ID that is used to locate the pcap file saved on the firewall.
High Resolution Timestamp (high_res_timestamp)	Time in milliseconds the log was received at the management plane. The format for this new field is YYYY-MM-DDThh:ss:sssTZD: • **YYYY**—Four digit year • **MM**—Two-digit month • **DD**—Two-digit day of the month (01 through 31) • **T**—Indicator for the beginning of the timestamp • **hh**—Two-digit hour using 24-hour time (00 through 23) • **mm**—Two-digit minute (00 through 59) • **ss**—Two-digit second (00 through 60) • **sss**—One or more digits for millisecond

https://t.me/learningnets

Field Name	Description
	• **TZD**—Time zone designator (+hh:mm or -hh:mm)  *The High Resolution Timestamp is supported for logs received from managed firewalls running PAN-OS 10.1 and later releases. Logs received from managed firewalls running PAN-OS 9.1 and earlier releases display a `1969-12-31T16:00:00:000-8:00` timestamp regardless of when the log was received.*
A Slice Service Type (nsdsai_sst)	The A Slice Service Type of the Network Slice ID.
A Slice Differentiator (nsdsai_sd)	The A Slice Differentiator of the Network Slice ID.
Application Subcategory (subcategory_of_app)	The application subcategory specified in the application configuration properties.
Application Category (category_of_app)	The application category specified in the application configuration properties. Values are: • business-systems • collaboration • general-internet • media • networking • saas
Application Technology (technology_of_app)	The application technology specified in the application configuration properties. Values are: • browser-based • client-server • network-protocol • peer-to-peer
Application Risk (risk_of_app)	Risk level associated with the application (1=lowest to 5=highest).
Application Characteristic (characteristic_of_app)	Comma-separated list of applicable characteristic of the application

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Field Name	Description
Application Container (container_of_app)	The parent application for an application.
Application SaaS (is_saas_of_app)	Displays `1` if a SaaS application or `0` if not a SaaS application.
Application Sanctioned State (sanctioned_state_of_app)	Displays `1` if application is sanctioned or `0` if application is not sanctioned.
Application Subcategory (subcategory_of_app)	The application subcategory specified in the application configuration properties.

Syslog Severity

The syslog severity is set based on the log type and contents.

Log Type/Severity	Syslog Severity
Traffic	Info
Config	Info
Threat/System—Informational	Info
Threat/System—Low	Notice
Threat/System—Medium	Warning
Threat/System—High	Warning
Threat/System—Critical	Critical

Custom Log/Event Format

To facilitate the integration with external log parsing systems, the firewall allows you to customize the log format; it also allows you to add custom *Key*: *Value* attribute pairs. Custom message formats can be configured under **Device** > **Server Profiles** > **Syslog** > **Syslog Server Profile** > **Custom Log Format**.

To achieve ArcSight Common Event Format (CEF) compliant log formatting, refer to the CEF Configuration Guide.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Escape Sequences

Any field that contains a comma or a double-quote is enclosed in double quotes. Furthermore, if a double-quote appears inside a field it is escaped by preceding it with another double-quote. To maintain backward compatibility, the Misc field in threat log is always enclosed in double-quotes.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

SNMP Monitoring and Traps

The following topics describe how Palo Alto Networks firewalls, Panorama, and WF-500 appliances implement SNMP, and the procedures to configure SNMP monitoring and trap delivery.

- SNMP Support
- Use an SNMP Manager to Explore MIBs and Objects
- Enable SNMP Services for Firewall-Secured Network Elements
- Monitor Statistics Using SNMP
- Forward Traps to an SNMP Manager
- Supported MIBs

SNMP Support

You can use an SNMP manager to monitor event-driven alerts and operational statistics for the firewall, Panorama, or WF-500 appliance and for the traffic they process. The statistics and traps can help you identify resource limitations, system changes or failures, and malware attacks. You configure alerts by forwarding log data as traps, and enable the delivery of statistics in response to GET messages (requests) from your SNMP manager. Each trap and statistic has an object identifier (OID). Related OIDs are organized hierarchically within the Management Information Bases (MIBs) that you load into the SNMP manager to enable monitoring.

When an event triggers SNMP trap generation (for example, an interface goes down), the firewall, Panorama virtual appliance, M-Series appliance, and WF-500 appliance respond by updating the corresponding SNMP object (for example, the interfaces MIB) instead of waiting for the periodic update of all objects that occurs every ten seconds. This ensures that your SNMP manager displays the latest information when polling an object to confirm an event.

The firewall, Panorama, and WF-500 appliance support SNMP Version 2c and Version 3. Decide which to use based on the version that other devices in your network support and on your network security requirements. SNMPv3 is more secure and enables more granular access control for system statistics than SNMPv2c. The following table summarizes the security features of each version. You select the version and configure the security features when you Monitor Statistics Using SNMP and Forward Traps to an SNMP Manager.

SNMPVe	Authentication	Message Privacy	Messag	MIB Access Granularity
SNMPv2c	Community string	No (cleartext)	No	SNMP community access for all MIBs on a device
SNMPv3	EngineID, username, and authentication password (SHA hashing for the password)	Privacy password for AES (128, 192, or 256) encryption	Yes	User access based on views that include or exclude specific OIDs

https://t.me/learningnets

SNMPVe	Authentication	Message Privacy	Messag	MIB Access Granularity
		of SNMP messages		

SNMP Implementation illustrates a deployment in which firewalls forward traps to an SNMP manager while also forwarding logs to Log Collectors. Alternatively, you could configure the Log Collectors to forward the firewall traps to the SNMP manager. For details on these deployments, refer to Log Forwarding Options in Centralized Logging and Reporting. In all deployments, the SNMP manager gets statistics directly from the firewall, Panorama, or WF-500 appliance. In this example, a single SNMP manager collects both traps and statistics, though you can use separate managers for these functions if that better suits your network.

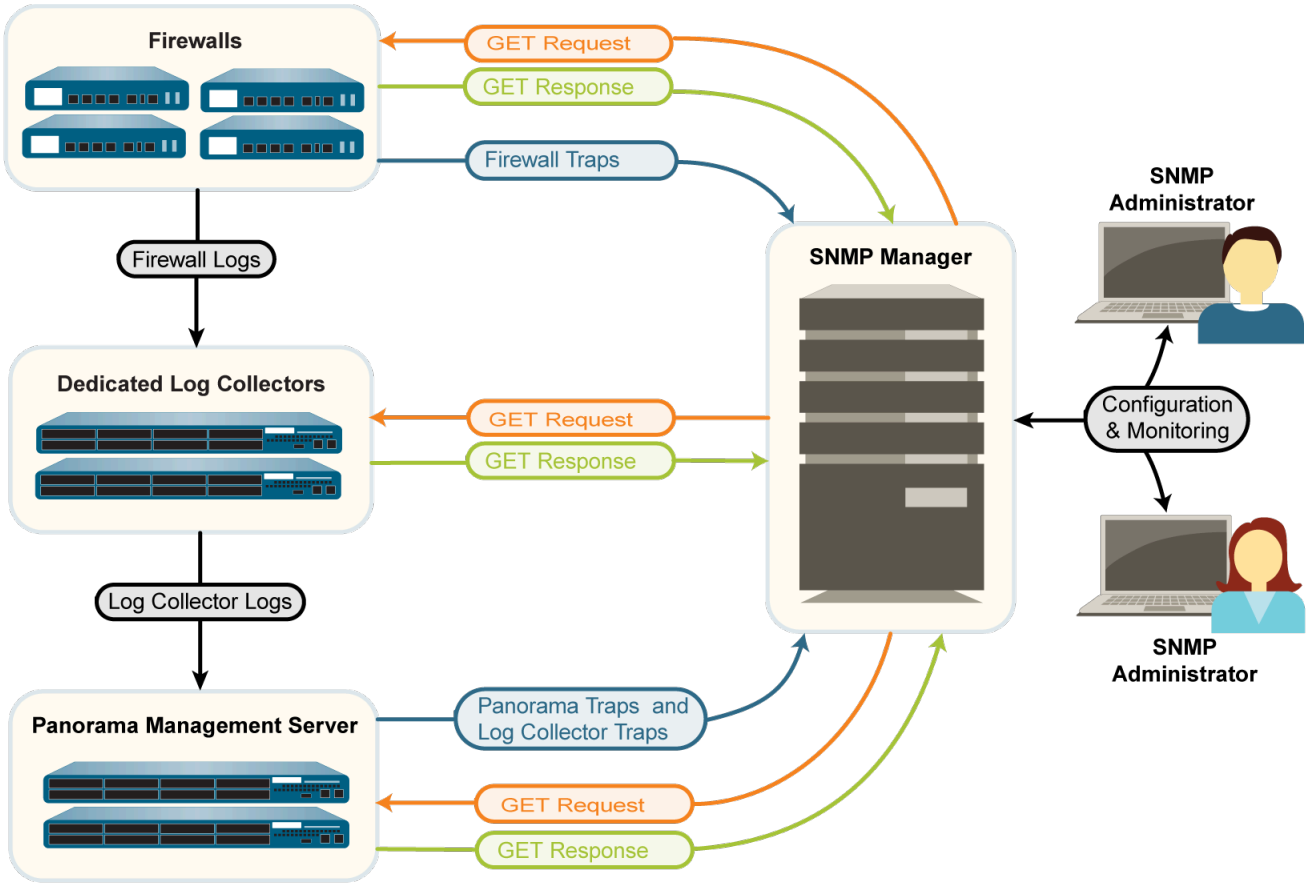


Figure 2: SNMP Implementation

Use an SNMP Manager to Explore MIBs and Objects

To use SNMP for monitoring Palo Alto Networks firewalls, Panorama, or WF-500 appliances, you must first load the Supported MIBs into your SNMP manager and determine which object identifiers (OIDs) correspond to the system statistics and traps you want to monitor. The following topics provide an overview of how to find OIDs and MIBs in an SNMP manager. For the specific steps to perform these tasks, refer to your SNMP management software.

- Identify a MIB Containing a Known OID

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

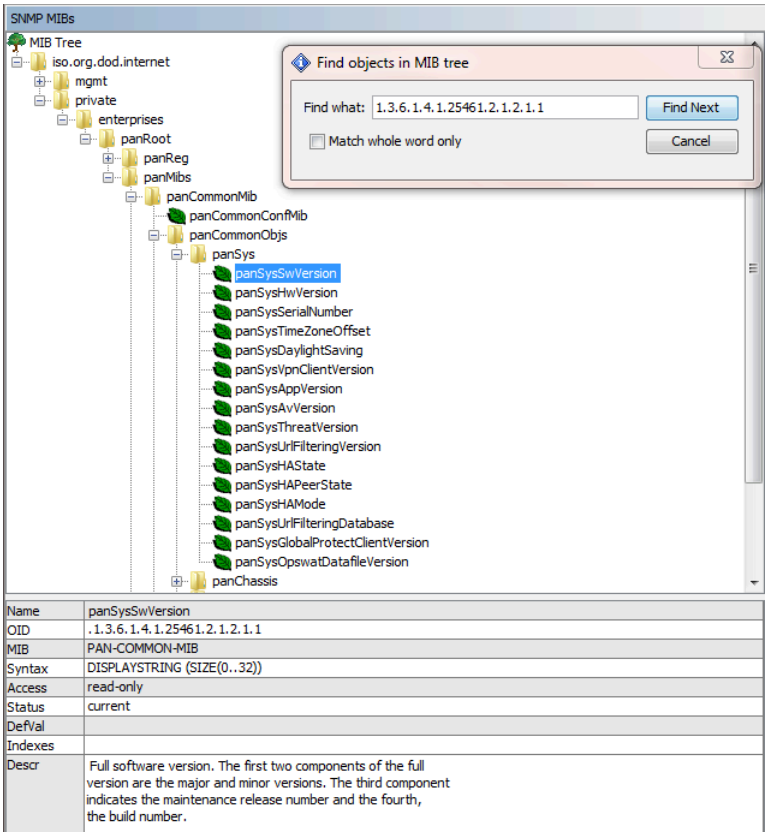
- Walk a MIB
- Identify the OID for a System Statistic or Trap

Identify a MIB Containing a Known OID

If you already know the OID for a particular SNMP object (statistic or trap) and want to know the OIDs of similar objects so you can monitor them, you can explore the MIB that contains the known OID.

STEP 1 | Load all the Supported MIBs into your SNMP manager.

STEP 2 | Search the entire MIB tree for the known OID. The search result displays the MIB path for the OID, as well as information about the OID (for example, name, status, and description). You can then select other OIDs in the same MIB to see information about them.

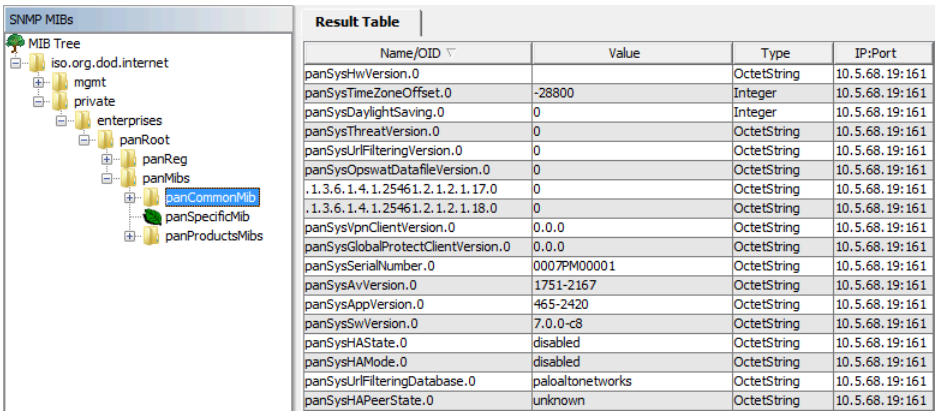


STEP 3 | (Optional) Walk a MIB to display all its objects.

Walk a MIB

If you want to see which SNMP objects (system statistics and traps) are available for monitoring, displaying all the objects of a particular MIB can be useful. To do this, load the Supported MIBs into your SNMP manager and perform a *walk* on the desired MIB. To list the traps that Palo Alto Networks firewalls, Panorama, and WF-500 appliance support, walk the panCommonEventEventsV2 MIB. In the following example, walking the PAN-COMMON-MIB.my displays the following list of OIDs and their values for certain statistics:

https://t.me/learningnets



Identify the OID for a System Statistic or Trap

To use an SNMP manager for monitoring Palo Alto Networks firewalls, Panorama, or WF-500 appliances, you must know the OIDs of the system statistics and traps you want to monitor.

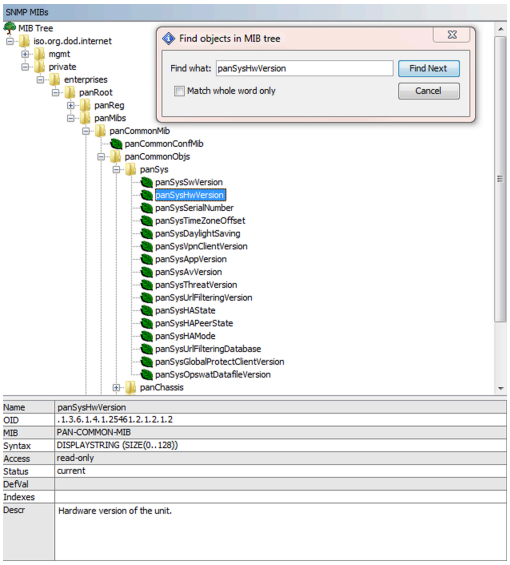
STEP 1 | Review the Supported MIBs to determine which one contains the type of statistic you want. For example, the PAN-COMMON-MIB.my contains hardware version information. The panCommonEventEventsV2 MIB contains all the traps that Palo Alto Networks firewalls, Panorama, and WF-500 appliances support.

STEP 2 | Open the MIB in a text editor and perform a keyword search. For example, using **Hardware version** as a search string in PAN-COMMON-MIB identifies the panSysHwVersion object:

```
panSysHwVersion OBJECT-TYPE
 SYNTAX DisplayString (SIZE(0..128))
 MAX-ACCESS read-only
 STATUS current
 DESCRIPTION
 "Hardware version of the unit."
 ::= {panSys 2}
```


https://t.me/learningnets

STEP 3 | In a MIB browser, search the MIB tree for the identified object name to display its OID. For example, the panSysHwVersion object has an OID of 1.3.6.1.4.1.25461.2.1.2.1.2.



Enable SNMP Services for Firewall-Secured Network Elements

If you will use Simple Network Management Protocol (SNMP) to monitor or manage network elements (for example, switches and routers) that are within the security zones of Palo Alto Networks firewalls, you must create a security rule that allows SNMP services for those elements.

You don't need a security rule to enable SNMP monitoring of Palo Alto Networks firewalls, Panorama, or WF-500 appliances. For details, see Monitor Statistics Using SNMP.

STEP 1 | Create an application group.

1. Select **Objects** > **Application Group** and click **Add**.
2. Enter a **Name** to identify the application group.
3. Click **Add**, type **snmp**, and select **snmp** and **snmp-trap** from the drop-down.
4. Click **OK** to save the application group.

STEP 2 | Create a security rule to allow SNMP services.

1. Select **Policies** > **Security** and click **Add**.
2. In the **General** tab, enter a **Name** for the rule.
3. In the **Source** and **Destination** tabs, click **Add** and enter a **Source Zone** and a **Destination Zone** for the traffic.
4. In the **Applications** tab, click **Add**, type the name of the applications group you just created, and select it from the drop-down.
5. In the **Actions** tab, verify that the **Action** is set to **Allow**, and then click **OK** and **Commit**.

Monitor Statistics Using SNMP

The statistics that a Simple Network Management Protocol (SNMP) manager collects from Palo Alto Networks firewalls can help you gauge the health of your network (systems and connections),

https://t.me/learningnets

identify resource limitations, and monitor traffic or processing loads. The statistics include information such as interface states (up or down), active user sessions, concurrent sessions, session utilization, temperature, and system uptime.

You can't configure an SNMP manager to control Palo Alto Networks firewalls (using SET messages), only to collect statistics from them (using GET messages). For details on how SNMP is implemented for Palo Alto Networks firewalls, see SNMP Support.

STEP 1 | Configure the SNMP Manager to get statistics from firewalls.

The following steps provide an overview of the tasks you perform on the SNMP manager. For the specific steps, refer to the documentation of your SNMP manager.

1. To enable the SNMP manager to interpret firewall statistics, load the Supported MIBs for Palo Alto Networks firewalls and, if necessary, compile them.
2. For each firewall that the SNMP manager will monitor, define the connection settings (IP address and port) and authentication settings (SNMPv2c community string or SNMPv3 EngineID/username/password) for the firewall.

   All Palo Alto Networks firewalls use port 161.

   The SNMP manager can use the same or different connection and authentication settings for multiple firewalls. The settings must match those you define when you configure SNMP on the firewall (see Step 3). For example, if you use SNMPv2c, the community string you define when configuring the firewall must match the community string you define in the SNMP manager for that firewall.
3. Determine the object identifiers (OIDs) of the statistics you want to monitor. For example, to monitor the session utilization percentage of a firewall, a MIB browser shows that this statistic corresponds to OID 1.3.6.1.4.1.25461.2.1.2.3.1.0 in PAN-COMMON-MIB.my. For details, see Use an SNMP Manager to Explore MIBs and Objects.
4. Configure the SNMP manager to monitor the desired OIDs.

STEP 2 | Enable SNMP traffic on a firewall interface.

This is the interface that will receive statistics requests from the SNMP manager.

PAN-OS doesn't synchronize management (MGT) interface settings for firewalls in a high availability (HA) configuration. You must configure the interface for each HA peer.

Perform this step in the firewall web interface.

- To enable SNMP traffic on the MGT interface, select **Device** > **Setup** > **Interfaces**, edit the **Management** interface, select **SNMP**, and then click **OK** and **Commit**.
- To enable SNMP traffic on any other interface, create an interface management profile for SNMP services and assign the profile to the interface that will receive the SNMP requests. The interface type must be Layer 3 Ethernet.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 3 | Configure the firewall to respond to statistics requests from an SNMP manager.

PAN-OS doesn't synchronize SNMP response settings for firewalls in a high availability (HA) configuration. You must configure these settings for each HA peer.

1. Select **Device** > **Setup** > **Operations** and, in the Miscellaneous section, click **SNMP Setup**.
2. Select the SNMP **Version** and configure the authentication values as follows. For version details, see SNMP Support.
   - **V2c**—Enter the **SNMP Community String**, which identifies a community of SNMP managers and monitored devices, and serves as a password to authenticate the community members to each other.

     As a best practice, don't use the default community string ***public****; it's well known and therefore not secure.*
   - **V3**—Create at least one SNMP view group and one user. User accounts and views provide authentication, privacy, and access control when firewalls forward traps and SNMP managers get firewall statistics.
     - **Views**—Each view is a paired OID and bitwise mask: the OID specifies a MIB and the mask (in hexadecimal format) specifies which objects are accessible within (include matching) or outside (exclude matching) that MIB. Click **Add** in the first list and enter a **Name** for the group of views. For each view in the group, click **Add** and configure the view **Name**, **OID**, matching **Option** (**include** or **exclude**), and **Mask**.
     - **Users**—Click **Add** in the second list, enter a username under **Users**, select the **View** group from the drop-down, enter the authentication password (**Auth Password**) used to authenticate to the SNMP manager, and enter the privacy password (**Priv Password**) used to encrypt SNMP messages to the SNMP manager.
3. Click **OK** and **Commit**.

STEP 4 | Monitor the firewall statistics in an SNMP manager.

Refer to the documentation of your SNMP manager for details.

When monitoring statistics related to firewall interfaces, you must match the interface indexes in the SNMP manager with interface names in the firewall web interface. For details, see Firewall Interface Identifiers in SNMP Managers and NetFlow Collectors.

Forward Traps to an SNMP Manager

Simple Network Management Protocol (SNMP) traps can alert you to system events (failures or changes in hardware or software of Palo Alto Networks firewalls) or to threats (traffic that matches a firewall security rule) that require immediate attention.

To see the list of traps that Palo Alto Networks firewalls support, use your SNMP Manager to access the panCommonEventEventsV2 MIB. For details, see Use an SNMP Manager to Explore MIBs and Objects.

For details on how for Palo Alto Networks firewalls implement SNMP, see SNMP Support.

https://t.me/learningnets

STEP 1 | Enable the SNMP manager to interpret the traps it receives.

Load the Supported MIBs for Palo Alto Networks firewalls and, if necessary, compile them. For the specific steps, refer to the documentation of your SNMP manager.

STEP 2 | Configure an SNMP Trap server profile.

The profile defines how the firewall accesses the SNMP managers (trap servers). You can define up to four SNMP managers for each profile.



Optionally, configure separate SNMP Trap server profiles for different log types, severity levels, and WildFire verdicts.

1. Log in to the firewall web interface.
2. Select **Device** > **Server Profiles** > **SNMP Trap**.
3. Click **Add** and enter a **Name** for the profile.
4. If the firewall has more than one virtual system (vsys), select the **Location** (vsys or **Shared**) where this profile is available.
5. Select the SNMP **Version** and configure the authentication values as follows. For version details, see SNMP Support.
   - **V2c**—For each server, click **Add** and enter the server **Name**, IP address (**SNMP Manager**), and **Community String**. The community string identifies a community of SNMP managers and monitored devices, and serves as a password to authenticate the community members to each other.

     As a best practice, don't use the default community string ***public****; it's well known and therefore not secure.*
   - **V3**—For each server, click **Add** and enter the server **Name**, IP address (**SNMP Manager**), SNMP **User** account (this must match a username defined in the SNMP manager), **EngineID** used to uniquely identify the firewall (you can leave the field blank to use the firewall serial number), authentication password (**Auth Password**) used to authenticate to the server, and privacy password (**Priv Password**) used to encrypt SNMP messages to the server.
6. Click **OK** to save the server profile.

STEP 3 | Configure log forwarding.

1. Configure the destinations of Traffic, Threat, and WildFire traps:
   1. Create a Log Forwarding profile. For each log type and each severity level or WildFire verdict, select the **SNMP Trap** server profile.
   2. Assign the Log Forwarding profile to policy rules and network zones. The rules and zones will trigger trap generation and forwarding.
2. Configure the destinations for System, Configuration, User-ID, HIP Match, and Correlation logs. For each log (trap) type and severity level, select the **SNMP Trap** server profile.
3. Click **Commit**.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 4 | Monitor the traps in an SNMP manager.

Refer to the documentation of your SNMP manager.

When monitoring traps related to firewall interfaces, you must match the interface indexes in the SNMP manager with interface names in the firewall web interface. For details, see Firewall Interface Identifiers in SNMP Managers and NetFlow Collectors.

Supported MIBs

The following table lists the Simple Network Management Protocol (SNMP) management information bases (MIBs) that Palo Alto Networks firewalls, Panorama, and WF-500 appliances support. You must load these MIBs into your SNMP manager to monitor the objects (system statistics and traps) that are defined in the MIBs. For details, see Use an SNMP Manager to Explore MIBs and Objects.

MIB Type	Supported MIBs
Standard—The Internet Engineering Task Force (IETF) maintains most standard MIBs. You can download the MIBs from the IETF website. *Palo Alto Networks firewalls, Panorama, and WF-500 appliances don't support every object (OID) in every one of these MIBs. See the Supported MIBs links for an overview of the supported OIDs.*	MIB-II IF-MIB HOST-RESOURCES-MIB ENTITY-MIB ENTITY-SENSOR-MIB ENTITY-STATE-MIB IEEE 802.3 LAG MIB LLDP-V2-MIB.my BFD-STD-MIB
Enterprise—You can download the enterprise MIBs from the Palo Alto Networks Technical Documentation portal.	PAN-COMMON-MIB.my PAN-GLOBAL-REG-MIB.my PAN-GLOBAL-TC-MIB.my PAN-LC-MIB.my

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

MIB Type	Supported MIBs
	PAN-PRODUCT-MIB.my PAN-ENTITY-EXT-MIB.my PAN-TRAPS.my

MIB-II

MIB-II provides object identifiers (OIDs) for network management protocols in TCP/IP-based networks. Use this MIB to monitor general information about systems and interfaces. For example, you can analyze trends in bandwidth usage by interface type (ifType object) to determine if the firewall needs more interfaces of that type to accommodate spikes in traffic volume.

Palo Alto Networks firewalls, Panorama, and WF-500 appliances support only the following object groups:

Object Group	Description
system	Provides system information such as the hardware model, system uptime, FQDN, and physical location.
interfaces	Provides statistics for physical and logical interfaces such as type, current bandwidth (speed), operational status (for example, up or down), and discarded packets. Logical interface support includes VPN tunnels, aggregate groups, Layer 2 subinterfaces, Layer 3 subinterfaces, loopback interfaces, and VLAN interfaces.

RFC 1213 defines this MIB.

IF-MIB

IF-MIB supports interface types (physical and logical) and larger counters (64K) beyond those defined in MIB-II. Use this MIB to monitor interface statistics in addition to those that MIB-II provides. For example, to monitor the current bandwidth of high-speed interfaces (greater than 2.2Gps) such as the 10G interfaces of the PA-5200 Series firewalls, you must check the ifHighSpeed object in IF-MIB instead of the ifSpeed object in MIB-II. IF-MIB statistics can be useful when evaluating the capacity of your network.

Palo Alto Networks firewalls, Panorama, and WF-500 appliances support only the ifXTable in IF-MIB, which provides interface information such as the number of multicast and broadcast packets transmitted and received, whether an interface is in promiscuous mode, and whether an interface has a physical connector.

RFC 2863 defines this MIB.

HOST-RESOURCES-MIB

HOST-RESOURCES-MIB provides information for host computer resources. Use this MIB to monitor CPU and memory usage statistics. For example, checking the current CPU load (hrProcessorLoad object) can help you troubleshoot performance issues on the firewall.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Palo Alto Networks firewalls, Panorama, and WF-500 appliances support portions of the following object groups:

Object Group	Description
hrDevice	Provides information such as CPU load, storage capacity, and partition size. The hrProcessorLoad OIDs provide an average of the cores that process packets. For the PA-7000 and PA-5200 Series firewalls, which have multiple dataplanes (DPs), you can monitor individual dataplane processor utilization. Set alerts when utilization reaches a specific threshold for each DP processor to avoid service availability issues.
hrSystem	Provides information such as system uptime, number of current user sessions, and number of current processes.
hrStorage	Provides information such as the amount of used storage.

RFC 2790 defines this MIB.

ENTITY-MIB

ENTITY-MIB provides OIDs for multiple logical and physical components. Use this MIB to determine what physical components are loaded on a system (for example, fans and temperature sensors) and see related information such as models and serial numbers. You can also use the index numbers for these components to determine their operational status in the ENTITY-SENSOR-MIB and ENTITY-STATE-MIB.

Palo Alto Networks firewalls, Panorama, and WF-500 appliances support only portions of the entPhysicalTable group:

Object	Description
entPhysicalIndex	A single namespace that includes disk slots and disk drives.
entPhysicalDescr	The component description.
entPhysicalVendorType	The sysObjectID (see PAN-PRODUCT-MIB.my) when it is available (chassis and module objects).
entPhysicalContainedIn	The value of entPhysicalIndex for the component that contains this component.
entPhysicalClass	Chassis (3), container (5) for a slot, power supply (6), fan (7), sensor (8) for each temperature or other environmental, and module (9) for each line card.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

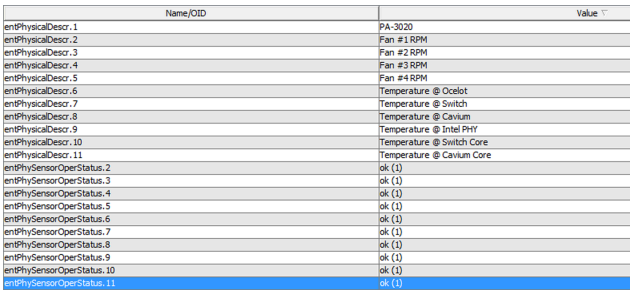
Object	Description
entPhysicalParentRelPos	The relative position of this *child* component among its *sibling* components. Sibling components are defined as entPhysicalEntry components that share the same instance values of each of the entPhysicalContainedIn and entPhysicalClass objects.
entPhysicalName	Supported only if the management (MGT) interface allows for naming the line card.
entPhysicalHardwareRev	The vendor-specific hardware revision of the component.
entPhysicalFirwareRev	The vendor-specific firmware revision of the component.
entPhysicalSoftwareRev	The vendor-specific software revision of the component.
entPhysicalSerialNum	The vendor-specific serial number of the component.
entPhysicalMfgName	The name of the manufacturer of the component.
entPhysicalMfgDate	The date when the component was manufactured.
entPhysicalModelName	The disk model number.
entPhysicalAlias	An alias that the network manager specified for the component.
entPhysicalAssetID	A user-assigned asset tracking identifier that the network manager specified for the component.
entPhysicalIsFRU	Indicates whether the component is a field replaceable unit (FRU).
entPhysicalUris	The Common Language Equipment Identifier (CLEI) number of the component (for example, URN:CLEI:CNME120ARA).

RFC 4133 defines this MIB.

ENTITY-SENSOR-MIB

ENTITY-SENSOR-MIB adds support for physical sensors of networking equipment beyond what ENTITY-MIB defines. Use this MIB in tandem with the ENTITY-MIB to monitor the operational status of the physical components of a system (for example, fans and temperature sensors). For example, to troubleshoot issues that might result from environmental conditions, you can map the entity indexes from the ENTITY-MIB (entPhysicalDescr object) to operational status values (entPhysSensorOperStatus object) in the ENTITY-SENSOR-MIB. In the following example, all the fans and temperature sensors for a PA-3020 firewall are working:

https://t.me/learningnets



Name/OID	Value
entPhysicalDescr.1	PA-3020
entPhysicalDescr.2	Fan #1 RPM
entPhysicalDescr.3	Fan #2 RPM
entPhysicalDescr.4	Fan #3 RPM
entPhysicalDescr.5	Fan #4 RPM
entPhysicalDescr.6	Temperature @ Ocelot
entPhysicalDescr.7	Temperature @ Switch
entPhysicalDescr.8	Temperature @ Cavium
entPhysicalDescr.9	Temperature @ Intel PHY
entPhysicalDescr.10	Temperature @ Switch Core
entPhysicalDescr.11	Temperature @ Cavium Core
entPhySensorOperStatus.2	ok (1)
entPhySensorOperStatus.3	ok (1)
entPhySensorOperStatus.4	ok (1)
entPhySensorOperStatus.5	ok (1)
entPhySensorOperStatus.6	ok (1)
entPhySensorOperStatus.7	ok (1)
entPhySensorOperStatus.8	ok (1)
entPhySensorOperStatus.9	ok (1)
entPhySensorOperStatus.10	ok (1)
entPhySensorOperStatus.11	ok (1)

The same OID might refer to different sensors on different platforms. Use the ENTITY-MIB for the targeted platform to match the value to the description.

Palo Alto Networks firewalls, Panorama, and WF-500 appliances support only portions of the entPhySensorTable group. The supported portions vary by platform and include only thermal (temperature in Celsius) and fan (in RPM) sensors.

RFC 3433 defines the ENTITY-SENSOR-MIB.

ENTITY-STATE-MIB

ENTITY-STATE-MIB provides information about the state of physical components beyond what ENTITY-MIB defines, including the administrative and operational state of components in chassis-based platforms. Use this MIB in tandem with the ENTITY-MIB to monitor the operational state of the components of a PA-7000 Series firewall (for example, line cards, fan trays, and power supplies). For example, to troubleshoot log forwarding issues for Threat logs, you can map the log processing card (LPC) indexes from the ENTITY-MIB (entPhysicalDescr object) to operational state values (entStateOper object) in the ENTITY-STATE-MIB. The operational state values use numbers to indicate state: 1 for unknown, 2 for disabled, 3 for enabled, and 4 for testing. The PA-7000 Series firewall is the only Palo Alto Networks firewall that supports this MIB.

RFC 4268 defines the ENTITY-STATE-MIB.

IEEE 802.3 LAG MIB

Use the IEEE 802.3 LAG MIB to monitor the status of aggregate groups that have Link Aggregation Control Protocol (LACP in an Aggregate Interface Group) enabled. When the firewall logs LACP events, it also generates traps that are useful for troubleshooting. For example, the traps can tell you whether traffic interruptions between the firewall and an LACP peer resulted from lost connectivity or from mismatched interface speed and duplex values.

PAN-OS implements the following SNMP tables for LACP.

The dot3adTablesLastChanged object indicates the time of the most recent change to dot3adAggTable, dot3adAggPortListTable, and dot3adAggPortTable.

Table	Description
Aggregator Configuration Table (dot3adAggTable)	This table contains information about every aggregate group that is associated with a firewall. Each aggregate group has one entry. Some table objects have restrictions, which the dot3adAggIndex object describes. This index is the unique identifier that the local

https://t.me/learningnets

Table	Description
	system assigns to the aggregate group. It identifies an aggregate group instance among the subordinate managed objects of the containing object. The identifier is read-only. *The ifTable MIB (a list of interface entries) does not support logical interfaces and therefore does not have an entry for the aggregate group.*
Aggregation Port List Table (dot3adAggPortListTable)	This table lists the ports associated with each aggregate group in a firewall. Each aggregate group has one entry. The dot3adAggPortListPorts attribute lists the complete set of ports associated with an aggregate group. Each bit set in the list represents a port member. For non-chassis platforms, this is a 64-bit value. For chassis platforms, the value is an array of eight 64-bit entries.
Aggregation Port Table (dot3adAggPortTable)	This table contains LACP configuration information about every port associated with an aggregate group in a firewall. Each port has one entry. The table has no entries for ports that are not associated with an aggregate group.
LACP Statistics Table (dot3adAggPortStatsTable)	This table contains link aggregation information about every port associated with an aggregate group in a firewall. Each port has one row. The table has no entries for ports that are not associated with an aggregate group.

The IEEE 802.3 LAG MIB includes the following LACP-related traps:

Trap Name	Description
panLACPLostConnectivityTrap	The peer lost connectivity to the firewall.
panLACPUnresponsiveTrap	The peer does not respond to the firewall.
panLACPNegoFailTrap	LACP negotiation with the peer failed.
panLACPSpeedDuplexTrap	The link speed and duplex settings on the firewall and peer do not match.
panLACPLinkDownTrap	An interface in the aggregate group is down.
panLACPLacpDownTrap	An interface was removed from the aggregate group.
panLACPLacpUpTrap	An interface was added to the aggregate group.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

For the MIB definitions, refer to IEEE 802.3 LAG MIB.

LLDP-V2-MIB.my

Use the LLDP-V2-MIB to monitor Link Layer Discovery Protocol (LLDP) events. For example, you can check the lldpV2StatsRxPortFramesDiscardedTotal object to see the number of LLDP frames that were discarded for any reason. The Palo Alto Networks firewall uses LLDP to discover neighboring devices and their capabilities. LLDP makes troubleshooting easier, especially for virtual wire deployments where the ping or traceroute utilities won't detect the firewall.

Palo Alto Networks firewalls support all the LLDP-V2-MIB objects except:

- The following lldpV2Statistics objects:
  - lldpV2StatsRemTablesLastChangeTime
  - lldpV2StatsRemTablesInserts
  - lldpV2StatsRemTablesDeletes
  - lldpV2StatsRemTablesDrops
  - lldpV2StatsRemTablesAgeouts
- The following lldpV2RemoteSystemsData objects:
  - The lldpV2RemOrgDefInfoTable table
  - In the lldpV2RemTable table: lldpV2RemTimeMark

RFC 4957 defines this MIB.

BFD-STD-MIB

Use the Bidirectional Forwarding Detection (BFD) MIB to monitor and receive failure alerts for the bidirectional path between two forwarding engines, such as interfaces, data links, or the actual engines. For example, you can check the bfdSessState object to see the state of a BFD session between forwarding engines. In the Palo Alto Networks implementation, one of the forwarding engines is a firewall interface and the other is an adjacent configured BFD peer.

RFC 7331 defines this MIB.

PAN-COMMON-MIB.my

Use the PAN-COMMON-MIB to monitor the following information for Palo Alto Networks firewalls, Panorama, and WF-500 appliances:

Object Group	Description
panSys	Contains such objects as system software/hardware versions, dynamic content versions, serial number, HA mode/state, and global counters. The global counters include those related to Denial of Service (DoS), IP fragmentation, TCP state, and dropped packets. Tracking these counters enables you to monitor traffic irregularities that result from DoS attacks, system or connection faults, or resource limitations. PAN-COMMON-MIB supports global counters for firewalls but not for Panorama.

https://t.me/learningnets

Object Group	Description
panChassis	Chassis type and M-Series appliance mode (Panorama or Log Collector).
panSession	Session utilization information. For example, the total number of active sessions on the firewall or a specific virtual system.
panMgmt	Status of the connection from the firewall to the Panorama management server.
panGlobalProtect	GlobalProtect gateway utilization as a percentage, maximum tunnels allowed, and number of active tunnels.
panLogCollector	Logging statistics for each Log Collector, including logging rate, log quotas, disk usage, retention periods, log redundancy (enabled or disabled), the forwarding status from firewalls to Log Collectors, the forwarding status from Log Collectors to external services, and the status of firewall-to-Log Collector connections.
panDeviceLogging	Logging statistics for each firewall, including logging rate, disk usage, retention periods, the forwarding status from individual firewalls to Panorama and external servers, and the status of firewall-to-Log Collector connections.

PAN-GLOBAL-REG-MIB.my

PAN-GLOBAL-REG-MIB.my contains global, top-level OID definitions for various sub-trees of Palo Alto Networks enterprise MIB modules. This MIB doesn't contain objects for you to monitor; it is required only for referencing by other MIBs.

PAN-GLOBAL-TC-MIB.my

PAN-GLOBAL-TC-MIB.my defines conventions (for example, character length and allowed characters) for the text values of objects in Palo Alto Networks enterprise MIB modules. All Palo Alto Networks products use these conventions. This MIB doesn't contain objects for you to monitor; it is required only for referencing by other MIBs.

PAN-LC-MIB.my

PAN-LC-MIB.my contains definitions of managed objects that Log Collectors (M-Series appliances in Log Collector mode) implement. Use this MIB to monitor the logging rate, log database storage duration (in days), and disk usage (in MB) of each logical disk (up to four) on a Log Collector. For example, you can use this information to determine whether you should add more Log Collectors or forward logs to an external server (for example, a syslog server) for archiving.

PAN-PRODUCT-MIB.my

PAN-PRODUCT-MIB.my defines sysObjectID OIDs for all Palo Alto Networks products. This MIB doesn't contain objects for you to monitor; it is required only for referencing by other MIBs.

https://t.me/learningnets

PAN-ENTITY-EXT-MIB.my

Use PAN-ENTITY-EXT-MIB.my in tandem with the ENTITY-MIB to monitor power usage for the physical components of a PA-7000 Series firewall (for example, fan trays, and power supplies), which is the only Palo Alto Networks firewall that supports this MIB. For example, when troubleshooting log forwarding issues, you might want to check the power usage of the log processing cards (LPCs): you can map the LPC indexes from the ENTITY-MIB (entPhysicalDescr object) to values in the PAN-ENTITY-EXT-MIB (panEntryFRUModelPowerUsed object).

PAN-TRAPS.my

Use PAN-TRAPS.my to see a complete listing of all the generated traps and information about them (for example, a description). For a list of traps that Palo Alto Networks firewalls, Panorama, and WF-500 appliances support, refer to the PAN-COMMON-MIB.my **panCommonEvents** > **panCommonEventsEvents** > **panCommonEventEventsV2** object.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

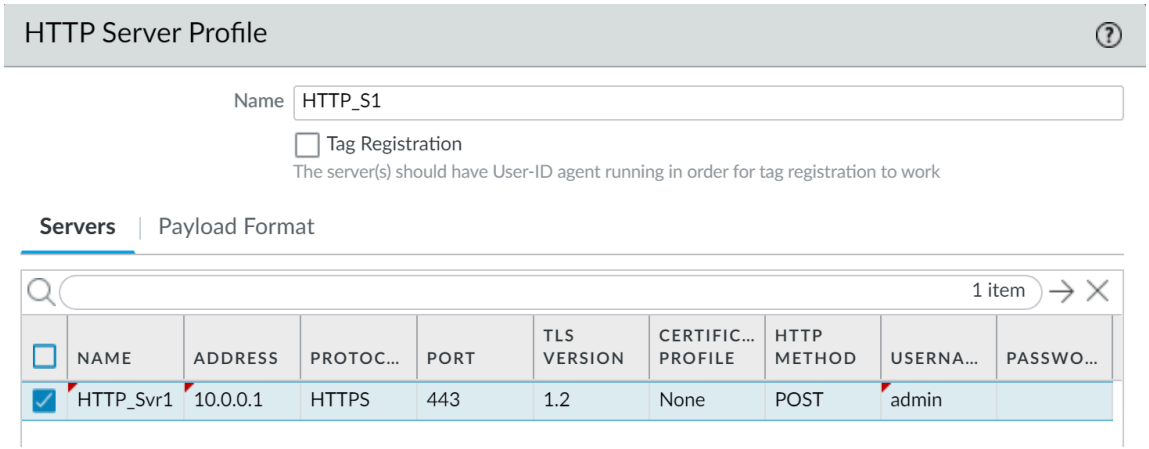
Forward Logs to an HTTP/S Destination

The firewall and Panorama™ can forward logs to an HTTP/S server. You can choose to forward all logs or specific logs to trigger an action on an external HTTP-based service when an event occurs. When forwarding logs to an HTTP server, configure the firewall to send an HTTP-based API request directly to a third-party service to trigger an action that is based on the attributes in a firewall log. You can configure the firewall to work with any HTTP-based service that exposes an API and you can modify the URL, HTTP header, parameters, and the payload in the HTTP request to meet your integration needs.

STEP 1 | Create an HTTP server profile to forward logs to an HTTP/S destination.

The HTTP server profile allows you to specify how to access the server and define the format in which to forward logs to the HTTP/S destination. By default, the firewall uses the management port to forward these logs. However, you can assign a different source interface and IP address in **Device** > **Setup** > **Services** > **Service Route Configuration**.

1. Select **Device** > **Server Profiles** > **HTTP** and **Add** a new profile.
2. Specify a **Name** for the server profile, and select the **Location**. The profile can be **Shared** across all virtual systems or can belong to a specific virtual system.
3. **Add** the details for each server. Each profile can have a maximum of four servers.
4. Enter a **Name** and IP **Address**.
5. Select the **Protocol** (**HTTP** or **HTTPS**). The default **Port** is 80 or 443 respectively but you can modify the port number to match the port on which your HTTP server listens.
6. Select the **TLS Version** supported on the server—**1.0**, **1.1**, or **1.2** (default).
7. Select the **Certificate Profile** to use for the TLS connection with the server.
8. Select the **HTTP Method** that the third-party service supports—**DELETE**,**GET**, **POST** (default), or **PUT**.
9. (Optional) Enter the **Username** and **Password** for authenticating to the server, if needed.
10. (Optional) Select **Test Server Connection** to verify network connectivity between the firewall and the HTTP/S server.



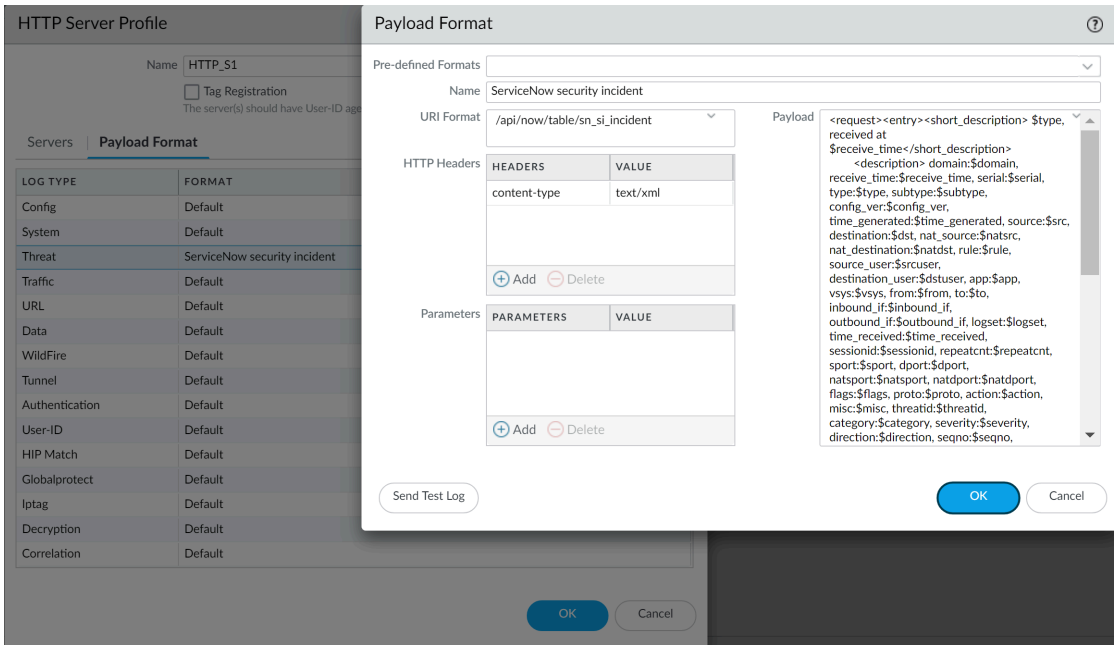
 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 2 | Select the **Payload Format** for the HTTP request.

1. Select the **Log Type** link for each log type for which you want to define the HTTP request format.
2. Select the **Pre-defined Formats** (available through content updates) or create a custom format.

   If you create a custom format, the **URI** is the resource endpoint on the HTTP service. The firewall appends the URI to the IP address you defined earlier to construct the URL for the HTTP request. Ensure that the URI and payload format matches the syntax that your third-party vendor requires. You can use any attribute supported on the selected log type within the HTTP Header, the Parameter and Value pairs, and in the request payload.

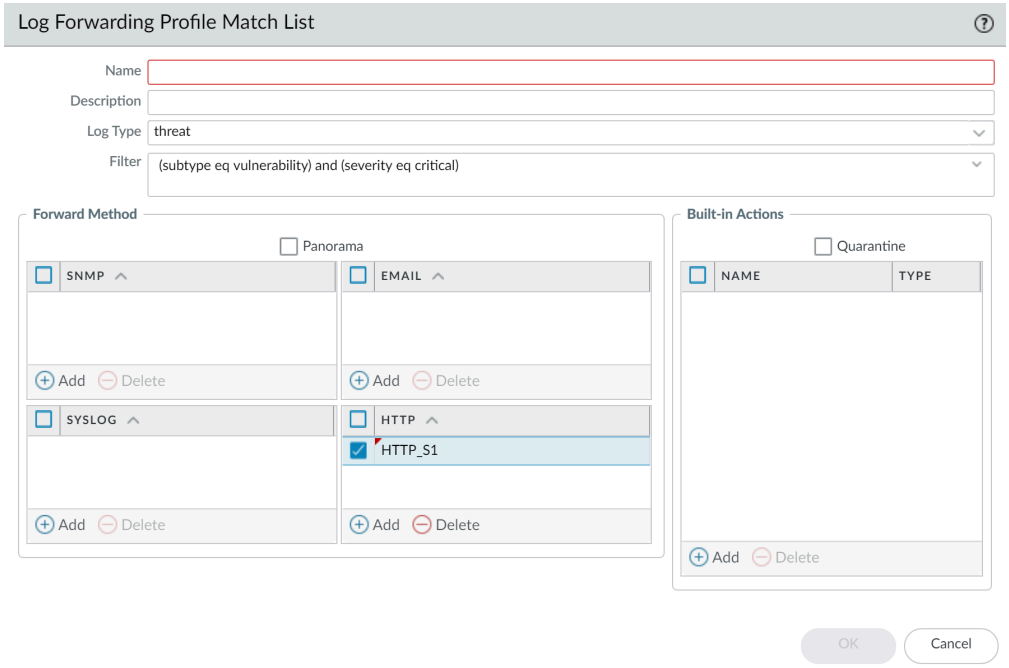


3. **Send Test Log** to verify that the HTTP server receives the request. When you interactively send a test log, the firewall uses the format as is and does not replace the variable with a value from a firewall log. If your HTTP server sends a 404 response, provide values for the parameters so that the server can process the request successfully.

https://t.me/learningnets

STEP 3 | Define the match criteria for when the firewall will forward logs to the HTTP server and attach the HTTP server profile you will use.

1. Select the log types for which you want to trigger a workflow:
   - Add a Log Forwarding Profile (**Objects** > **Log Forwarding**) for logs that pertain to user activity (for example, Traffic, Threat, or Authentication logs).
   - Select **Device** > **Log Settings** for logs that pertain to system events, such as Configuration or System logs.
2. Select the Log Type and use the new **Filter Builder** to define the match criteria.
3. **Add** the HTTP server profile for forwarding logs to the HTTP destination.



https://t.me/learningnets

NetFlow Monitoring

NetFlow is an industry-standard protocol that the firewall can use to export statistics about the IP traffic ingressing its interfaces. The firewall exports the statistics as NetFlow fields to a NetFlow collector. The NetFlow collector is a server you use to analyze network traffic for security, administration, accounting and troubleshooting. All Palo Alto Networks firewalls support NetFlow Version 9. The firewalls support only unidirectional NetFlow, not bidirectional. The firewalls perform NetFlow processing on all IP packets on the interfaces and do not support sampled NetFlow. You can export NetFlow records for Layer 3, Layer 2, virtual wire, tap, VLAN, loopback, and tunnel interfaces. For aggregate Ethernet sub-interfaces, you can export records for the individual sub-interfaces that data flows through within the group. To identify firewall interfaces in a NetFlow collector, see Firewall Interface Identifiers in SNMP Managers and NetFlow Collectors. The firewalls support standard and enterprise (PAN-OS specific) NetFlow Templates, which NetFlow collectors use to decipher the NetFlow fields.

- Configure NetFlow Exports
- NetFlow Templates

Configure NetFlow Exports

To use a NetFlow collector for analyzing the network traffic ingressing firewall interfaces, perform the following steps to configure NetFlow record exports.

STEP 1 | Create a NetFlow server profile.

The profile defines which NetFlow collectors will receive the exported records and specifies export parameters.

1. Select **Device** > **Server Profiles** > **NetFlow** and **Add** a profile.
2. Enter a **Name** to identify the profile.
3. Specify the rate at which the firewall refreshes NetFlow Templates in **Minutes** (default is 30) and **Packets** (exported records—default is 20), according to the requirements of your NetFlow collector. The firewall refreshes the templates after either threshold is passed.
4. Specify the **Active Timeout**, which is the frequency in minutes at which the firewall exports records (default is 5).
5. Select **PAN-OS Field Types** if you want the firewall to export App-ID and User-ID fields.
6. **Add** each NetFlow collector (up to two per profile) that will receive records. For each collector, specify the following:
   - **Name** to identify the collector.
   - **NetFlow Server** hostname or IP address.
   - Access **Port** (default 2055).
7. Click **OK** to save the profile.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 2 | Assign the NetFlow server profile to the firewall interfaces where traffic you want to analyze is ingressing.

In this example, you assign the profile to an existing Ethernet interface.

1. Select **Network** > **Interfaces** > **Ethernet** and click an interface name to edit it.

   You can export NetFlow records for Layer 3, Layer 2, virtual wire, tap, VLAN, loopback, and tunnel interfaces. For aggregate Ethernet interfaces, you can export records for the individual sub-interfaces that data flows through within the group.

2. Select the NetFlow server profile (**NetFlow Profile**) you configured and click **OK**.

STEP 3 | **(Required for PA-7000 Series, PA-5400 Series, and PA-5200 Series firewalls)** Configure a service route for the interface that the firewall will use to send NetFlow records.

You cannot use the management (MGT) interface to send NetFlow records from the PA-7000 Series, PA-5400 Series, and PA-5200 Series firewalls. For other firewall models, a service route is optional. For all firewalls, the interface that sends NetFlow records does not have to be the same as the interface for which the firewall collects the records.

1. Select **Device** > **Setup** > **Services**.
2. **(Firewall with multiple virtual systems)** Select one of the following:
   - **Global**—Select this option if the service route applies to all virtual systems on the firewall.
   - **Virtual Systems**—Select this option if the service route applies to a specific virtual system. Set the **Location** to the virtual system.
3. Select **Service Route Configuration** and Customize.
4. Select the protocol (**IPv4** or **IPv6**) that the interface uses. You can configure the service route for both protocols if necessary.
5. Click **Netflow** in the Service column.
6. Select the **Source Interface**.

   Any, ***Use default***, *and* ***MGT*** *are not valid interface options for sending NetFlow records from PA-7000 Series, PA-5400 Series, or PA-5200 Series firewalls.*

7. Select a **Source Address** (IP address).
8. Click **OK** twice to save your changes.

STEP 4 | **Commit** your changes.

STEP 5 | Monitor the firewall traffic in a NetFlow collector.

Refer to your NetFlow collector documentation.

When monitoring statistics, you must match the interface indexes in the NetFlow collector with interface names in the firewall web interface. For details, see Firewall Interface Identifiers in SNMP Managers and NetFlow Collectors.

To troubleshoot NetFlow delivery issues, use the operational CLI command **`debug log-receiver netflow statistics`**.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

NetFlow Templates

NetFlow collectors use templates to decipher the fields that the firewall exports. The firewall selects a template based on the type of exported data: IPv4 or IPv6 traffic, with or without NAT, and with standard or enterprise-specific (PAN-OS specific) fields. The firewall periodically refreshes templates to re-evaluate which one to use (in case the type of exported data changes) and to apply any changes to the fields in the selected template. When you Configure NetFlow Exports, set the refresh rate based on a time interval and a number of exported records according to the requirements of your NetFlow collector. The firewall refreshes the templates after either threshold is passed.

The Palo Alto Networks firewall supports the following NetFlow templates:

Template	ID
IPv4 Standard	256
IPv4 Enterprise	257
IPv6 Standard	258
IPv6 Enterprise	259
IPv4 with NAT Standard	260
IPv4 with NAT Enterprise	261
IPv6 with NAT Standard	262
IPv6 with NAT Enterprise	263

The following table lists the NetFlow fields that the firewall can send, along with the templates that define them:

Value	Field	Description	Templates
1	IN_BYTES	Incoming counter with length N * 8 bits for the number of bytes associated with an IP flow. By default, N is 4.	All templates
2	IN_PKTS	Incoming counter with length N * 8 bits for the number of packets associated with an IP glow. By default, N is 4.	All templates
4	PROTOCOL	IP protocol byte.	All templates

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Value	Field	Description	Templates
5	TOS	Type of Service byte setting when entering the ingress interface.	All templates
6	TCP_FLAGS	Total of all the TCP flags in this flow.	All templates
7	L4_SRC_PORT	TCP/UDP source port number (for example, FTP, Telnet, or equivalent).	All templates
8	IPV4_SRC_ADDR	IPv4 source address.	IPv4 standard IPv4 enterprise IPv4 with NAT standard IPv4 with NAT enterprise
10	INPUT_SNMP	Input interface index. The value length is 2 bytes by default, but higher values are possible. For details on how Palo Alto Networks firewalls generate interface indexes, see Firewall Interface Identifiers in SNMP Managers and NetFlow Collectors.	All templates
11	L4_DST_PORT	TCP/UDP destination port number (for example, FTP, Telnet, or equivalent).	All templates
12	IPV4_DST_ADDR	IPv4 destination address.	IPv4 standard IPv4 enterprise IPv4 with NAT standard IPv4 with NAT enterprise
14	OUTPUT_SNMP	Output interface index. The value length is 2 bytes by default, but higher values are possible. For details on how Palo Alto Networks firewalls generate interface indexes, see Firewall Interface	All templates

https://t.me/learningnets

Value	Field	Description	Templates
		Identifiers in SNMP Managers and NetFlow Collectors.	
21	LAST_SWITCHED	System uptime in milliseconds when the last packet of this flow was switched.	All templates
22	FIRST_SWITCHED	System uptime in milliseconds when the first packet of this flow was switched.	All templates
27	IPV6_SRC_ADDR	IPv6 source address.	IPv6 standard IPv6 enterprise IPv6 with NAT standard IPv6 with NAT enterprise
28	IPV6_DST_ADDR	IPv6 destination address.	IPv6 standard IPv6 enterprise IPv6 with NAT standard IPv6 with NAT enterprise
32	ICMP_TYPE	Internet Control Message Protocol (ICMP) packet type. This is reported as: ICMP Type * 256 + ICMP code	All templates
61	DIRECTION	Flow direction: • 0 = ingress • 1 = egress	All templates
148	flowId	An identifier of a flow that is unique within an observation domain. You can use this information element to distinguish between different flows if flow keys such as IP addresses and port numbers are not reported or are reported in separate records. The	All templates

https://t.me/learningnets

Value	Field	Description	Templates
		flowID corresponds to the session ID field in Traffic and Threat logs.	
233	firewallEvent	Indicates a firewall event: • 0 = Ignore (invalid)—Not used. • 1 = Flow created—The NetFlow data record is for a new flow. • 2 = Flow deleted—The NetFlow data record is for the end of a flow. • 3 = Flow denied—The NetFlow data record indicates a flow that firewall policy denied. • 4 = Flow alert—Not used. • 5 = Flow update—The NetFlow data record is sent for a *long-lasting* flow, which is a flow that lasts longer than the **Active Timeout** period configured in the NetFlow server profile.	All templates
225	postNATSourceIPv4Address	The definition of this information element is identical to that of sourceIPv4Address, except that it reports a modified value that the firewall produced during network address translation after the packet traversed the interface.	IPv4 with NAT standard IPv4 with NAT enterprise
226	postNATDestinationIPv4Address	The definition of this information element is identical to that of destinationIPv4Address, except that it reports a modified value that the firewall produced during network address translation after the packet traversed the interface.	IPv4 with NAT standard IPv4 with NAT enterprise
227	postNAPTSourceTransportPort	The definition of this information element is identical to that of sourceTransportPort, except that it reports a modified value that the firewall produced during network address port translation after the packet traversed the interface.	IPv4 with NAT standard IPv4 with NAT enterprise

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Value	Field	Description	Templates
228	postNAPTDestinationTransportPort	The definition of this information element is identical to that of destinationTransportPort, except that it reports a modified value that the firewall produced during network address port translation after the packet traversed the interface.	IPv4 with NAT standard IPv4 with NAT enterprise
281	postNATSourceIPv6Address	The definition of this information element is identical to the definition of information element sourceIPv6Address, except that it reports a modified value that the firewall produced during NAT64 network address translation after the packet traversed the interface. See RFC 2460 for the definition of the source address field in the IPv6 header. See RFC 6146 for NAT64 specification.	IPv6 with NAT standard IPv6 with NAT enterprise
282	postNATDestinationIPv6Address	The definition of this information element is identical to the definition of information element destinationIPv6Address, except that it reports a modified value that the firewall produced during NAT64 network address translation after the packet traversed the interface. See RFC 2460 for the definition of the destination address field in the IPv6 header. See RFC 6146 for NAT64 specification.	IPv6 with NAT standard IPv6 with NAT enterprise
346	privateEnterpriseNumber	This is a unique private enterprise number that identifies Palo Alto Networks: 25461.	IPv4 enterprise IPv4 with NAT enterprise IPv6 enterprise IPv6 with NAT enterprise

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Value	Field	Description	Templates
56701	App-ID	The name of an application that App-ID identified. The name can be up to 32 bytes.	IPv4 enterprise IPv4 with NAT enterprise IPv6 enterprise IPv6 with NAT enterprise
56702	User-ID	A username that User-ID identified. The name can be up to 64 bytes.	IPv4 enterprise IPv4 with NAT enterprise IPv6 enterprise IPv6 with NAT enterprise

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Firewall Interface Identifiers in SNMP Managers and NetFlow Collectors

When you use a NetFlow collector (see NetFlow Monitoring) or SNMP manager (see SNMP Monitoring and Traps) to monitor the Palo Alto Networks firewall, an interface index (SNMP ifindex object) identifies the interface that carried a particular flow (see Interface Indexes in an SNMP Manager). In contrast, the firewall web interface uses interface names as identifiers (for example, ethernet1/1), not indexes. To understand which statistics that you see in a NetFlow collector or SNMP manager apply to which firewall interface, you must be able to match the interface indexes with interface names.

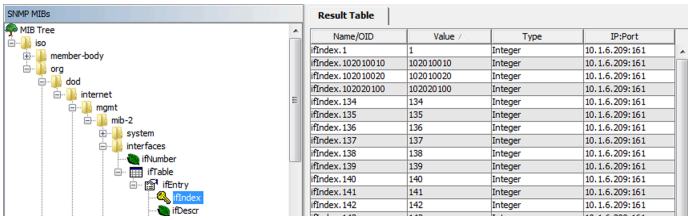


Figure 3: Interface Indexes in an SNMP Manager

You can match the indexes with names by understanding the formulas that the firewall uses to calculate indexes. The formulas vary by platform and interface type: physical or logical.

Physical interface indexes have a range of 1-9999, which the firewall calculates as follows:

Firewall Platform	Calculation	Example Interface Index
VM-Series	Number of management ports + physical port offset • **Number of management ports**—This is a constant of 1. • **Physical port offset**—This is the physical port number.	VM-100 firewall, Eth1/4 = 1 (number of management ports) + 4 (physical port) = **5**
PA-220, PA-220R, PA-800 Series	Number of management ports + physical port offset • **Number of management ports**—This is a constant of 5. • **Physical port offset**—This is the physical port number.	PA-5200 Series firewall, Eth1/4 = 5 (number of management ports) + 4 (physical port) = **9**
PA-3200 Series, PA-5200 Series	Number of management ports + physical port offset • **Number of management ports**—This is a constant of 4.	PA-5200 Series firewall, Eth1/4 = 4 (number of management ports) + 4 (physical port) = **8**

https://t.me/learningnets

Firewall Platform	Calculation	Example Interface Index
	• **Physical port offset**—This is the physical port number.	
PA-7000 Series	(Max. ports * slot) + physical port offset + number of management ports • **Maximum ports**—This is a constant of 64. • **Slot**—This is the chassis slot number of the network interface card. • **Physical port offset**—This is the physical port number. • **Number of management ports**—This is a constant of 5.	PA-7000 Series firewall, Eth3/9 = [64 (max. ports) * 3 (slot)] + 9 (physical port) + 5 (number of management ports) = **206**

Logical interface indexes for all platforms are nine-digit numbers that the firewall calculates as follows:

Interface Type	Range	Digit 9	Digits 7-8	Digits 5-6	Digits 1-4	Example Interface Index
Layer 3 subinterface	101010001-199999999	Type: 1	Interface slot: 1-9 (01-09)	Interface port: 1-9 (01-09)	Subinterface: suffix 1-9999 (0001-9999)	Eth1/5.22 = 100000000 (type) + 100000 (slot) + 50000 (port) + 22 (suffix) = **101050022**
Layer 2 subinterface	101010001-199999999	Type: 1	Interface slot: 1-9 (01-09)	Interface port: 1-9 (01-09)	Subinterface: suffix 1-9999 (0001-9999)	Eth2/3.6 = 100000000 (type) + 200000 (slot) + 30000 (port) + 6 (suffix) = **102030006**
Vwire subinterface	101010001-199999999	Type: 1	Interface slot: 1-9 (01-09)	Interface port: 1-9 (01-09)	Subinterface: suffix 1-9999 (0001-9999)	Eth4/2.312 = 100000000 (type) + 400000 (slot) + 20000 (port) + 312 (suffix) = **104020312**
VLAN	200000001-200009999	Type: 2	00	00	VLAN suffix: 1-9999 (0001-9999)	VLAN.55 = 200000000 (type) + 55 (suffix) = **200000055**

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Interface Type	Range	Digit 9	Digits 7-8	Digits 5-6	Digits 1-4	Example Interface Index
Loopback	300000001-300009999	Type: 3	00	00	Loopback suffix: 1-9999 (0001-9999)	Loopback.55 = 300000000 (type) + 55 (suffix) = **300000055**
Tunnel	400000001-400009999	Type: 4	00	00	Tunnel suffix: 1-9999 (0001-9999)	Tunnel.55 = 400000000 (type) + 55 (suffix) = **400000055**
Aggregate group	500010001-500089999	Type: 5	00	AE suffix: 1-8 (01-08)	Subinterface: suffix 1-9999 (0001-9999)	AE5.99 = 500000000 (type) + 50000 (AE Suffix) + 99 (suffix) = **500050099**

©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Monitor Transceivers

You can monitor the status of transceivers in your physical appliance or device to enable easier installation and troubleshooting. Diagnostics that can be viewed are transmitted bias current, transmitted power, received power, transceiver temperature, and power supply voltage. See below for a list of devices that support transceiver monitoring.

- PA-800 Series
- PA-3200 Series
- PA-5200 Series
- PA-5450 Firewall
- PA-7000 Series

Use the Command Line Interface to run transceiver monitoring. See the following table for all available CLI commands.

If you run commands on an incompatible transceiver, the CLI will return 'n/a' for any diagnostic information it cannot read.

CLI	Definition
`show transceiver <interface name>`	View a summary of the specified transceiver with values for each diagnostic. Example: `admin@PA-7080> show transceiver ethernet11/25` The CLI will return values for Temperature, Voltage, Current, Tx Power, and Rx Power.
`show transceiver-detail <interface name>`	Receive more detailed transceiver specifications, including vendor information and link lengths. The CLI will also provide more detailed diagnostic information.
`show transceiver all`	View a list of all active transceivers as well as a summary of each of their diagnostics.
`show transceiver-detail all`	Get comprehensive details on each transceiver in the device.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

User-ID

The user identity, as opposed to an IP address, is an integral component of an effective security infrastructure. Knowing who is using each of the applications on your network, and who may have transmitted a threat or is transferring files, can strengthen security policies and reduce incident response times. User-ID™, a standard feature on the Palo Alto Networks firewall, enables you to leverage user information stored in a wide range of repositories. The following topics provide more details about User-ID and how to configure it:

- User-ID Overview
- User-ID Concepts
- Enable User-ID
- Map Users to Groups
- Map IP Addresses to Users
- Enable User- and Group-Based Policy
- Enable Policy for Users with Multiple Accounts
- Verify the User-ID Configuration
- Deploy User-ID in a Large-Scale Network

https://t.me/learningnets

User-ID Overview

User-ID™ enables you to identify all users on your network using a variety of techniques to ensure that you can identify users in all locations using a variety of access methods and operating systems, including Microsoft Windows, Apple iOS, Mac OS, Android, and Linux®/UNIX. Knowing who your users are instead of just their IP addresses enables:

- **Visibility**—Improved visibility into application usage based on users gives you a more relevant picture of network activity. The power of User-ID becomes evident when you notice a strange or unfamiliar application on your network. Using either ACC or the log viewer, your security team can discern what the application is, who the user is, the bandwidth and session consumption, along with the source and destination of the application traffic, as well as any associated threats.
- **Policy control**—Tying user information to Security policy rules improves safe enablement of applications traversing the network and ensures that only those users who have a business need for an application have access. For example, some applications, such as SaaS applications that enable access to Human Resources services (such as Workday or Service Now) must be available to any known user on your network. However, for more sensitive applications you can reduce your attack surface by ensuring that only users who need these applications can access them. For example, while IT support personnel may legitimately need access to remote desktop applications, the majority of your users do not.
- **Logging, reporting, forensics**—If a security incident occurs, forensics analysis and reporting based on user information rather than just IP addresses provides a more complete picture of the incident. For example, you can use the pre-defined User/Group Activity to see a summary of the web activity of individual users or user groups, or the SaaS Application Usage report to see which users are transferring the most data over unsanctioned SaaS applications.

To enforce user- and group-based policies, the firewall must be able to map the IP addresses in the packets it receives to usernames. User-ID provides many mechanisms to collect this User Mapping information. For example, the User-ID agent monitors server logs for login events and listens for syslog messages from authenticating services. To identify mappings for IP addresses that the agent didn't map, you can configure Authentication Policy to redirect HTTP requests to an Authentication Portal login. You can tailor the user mapping mechanisms to suit your environment, and even use different mechanisms at different sites to ensure that you are safely enabling access to applications for all users, in all locations, all the time.

https://t.me/learningnets

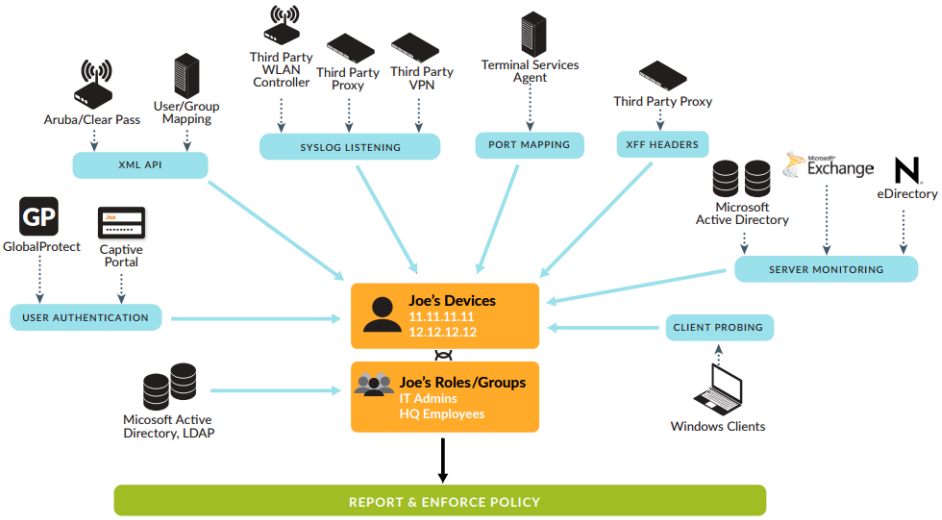


Figure 4: User-ID

To enable user- and group-based policy enforcement, the firewall requires a list of all available users and their corresponding group memberships so that you can select groups when defining your policy rules. The firewall collects Group Mapping information by connecting directly to your LDAP directory server, or using XML API integration with your directory server.

See User-ID Concepts for information on how User-ID works and Enable User-ID for instructions on setting up User-ID.

User-ID does not work in environments where the source IP addresses of users are subject to NAT translation before the firewall maps the IP addresses to usernames.

https://t.me/learningnets

User-ID Concepts

- Group Mapping
- User Mapping

Group Mapping

To define policy rules based on user or group, first you create an LDAP server profile that defines how the firewall connects and authenticates to your directory server. The firewall supports a variety of directory servers, including Microsoft Active Directory (AD), Novell eDirectory, and Sun ONE Directory Server. The server profile also defines how the firewall searches the directory to retrieve the list of groups and the corresponding list of members. If you are using a directory server that is not natively supported by the firewall, you can integrate the group mapping function using the XML API. You can then create a group mapping configuration to Map Users to Groups and Enable User- and Group-Based Policy.

Defining policy rules based on group membership rather than on individual users simplifies administration because you don't have to update the rules whenever new users are added to a group. When configuring group mapping, you can limit which groups will be available in policy rules. You can specify groups that already exist in your directory service or define custom groups based on LDAP filters. Defining custom groups can be quicker than creating new groups or changing existing ones on an LDAP server, and doesn't require an LDAP administrator to intervene. User-ID maps all the LDAP directory users who match the filter to the custom group. For example, you might want a security policy that allows contractors in the Marketing Department to access social networking sites. If no Active Directory group exists for that department, you can configure an LDAP filter that matches users for whom the LDAP attribute Department is set to Marketing. Log queries and reports that are based on user groups will include custom groups.

User Mapping

Knowing user and groups names is only one piece of the puzzle. The firewall also needs to know which IP addresses map to which users so that security rules can be enforced appropriately. User-ID Overview illustrates the different methods that are used to identify users and groups on your network and shows how user mapping and group mapping work together to enable user- and group-based security enforcement and visibility. The following topics describe the different methods of user mapping:

- Server Monitoring
- Port Mapping
- Syslog
- XFF Headers
- Username Header Insertion
- Authentication Policy and Authentication Portal
- GlobalProtect
- XML API

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

- Client Probing

Server Monitoring

With server monitoring a User-ID agent—either a Windows-based agent running on a domain server in your network, or the PAN-OS integrated User-ID agent running on the firewall—monitors the security event logs for specified Microsoft Exchange Servers, Domain Controllers, or Novell eDirectory servers for login events. For example, in an AD environment, you can configure the User-ID agent to monitor the security logs for Kerberos ticket grants or renewals, Exchange server access (if configured), and file and print service connections. For these events to be recorded in the security log, the AD domain must be configured to log successful account login events. In addition, because users can log in to any of the servers in the domain, you must set up server monitoring for all servers to capture all user login events. See Configure User Mapping Using the Windows User-ID Agent or Configure User Mapping Using the PAN-OS Integrated User-ID Agent for details.

Port Mapping

In environments with multi-user systems—such as Microsoft Terminal Server or Citrix environments—many users share the same IP address. In this case, the user-to-IP address mapping process requires knowledge of the source port of each client. To perform this type of mapping, you must install the Palo Alto Networks Terminal Server Agent on the Windows/Citrix terminal server itself to intermediate the assignment of source ports to the various user processes. For terminal servers that do not support the Terminal Server agent, such as Linux terminal servers, you can use the XML API to send user mapping information from login and logout events to User-ID. See Configure User Mapping for Terminal Server Users for configuration details.

XFF Headers

If you have a proxy server deployed between the users on your network and the firewall, the firewall might see the proxy server IP address as the source IP address in HTTP/HTTPS traffic that the proxy forwards rather than the IP address of the client that requested the content. In many cases, the proxy server adds an X-Forwarded-For (XFF) header to traffic packets that includes the actual IPv4 or IPv6 address of the client that requested the content or from whom the request originated. In such cases, you can configure the firewall to extract the end user IP address from the XFF so that User-ID can map the IP address to a username. This enables you to Use XFF Values for Policies and Logging Source Users so that you can enforce user-based policy to safely enable access to web-based for your users behind a proxy server.

Username Header Insertion

When you configure a secondary enforcement device with your Palo Alto Networks firewall to enforce user-based policy, the secondary device may not have the IP address-to-username mapping from the firewall. Transmitting the user's identity to downstream devices may require deployment of additional devices such as proxies or negatively impact the user's experience (for example, users having to log in multiple times). You can dynamically add the domain and username to the HTTP header of the user's outgoing traffic, allowing any secondary devices that you use with your Palo Alto Networks firewall to receive the user's information and enforce user-based policy. Including the user's identity by inserting the username and domain in the traffic headers enables enforcement of user-based policy without negatively impacting the user's experience or deployment of additional infrastructure.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Authentication Policy and Authentication Portal

In some cases, the User-ID agent can't map an IP address to a username using server monitoring or other methods—for example, if the user isn't logged in or uses an operating system such as Linux that your domain servers don't support. In other cases, you might want users to authenticate when accessing sensitive applications regardless of which methods the User-ID agent uses to perform user mapping. For all these cases, you can configure Configure Authentication Policy and Map IP Addresses to Usernames Using Authentication Portal. Any web traffic (HTTP or HTTPS) that matches an Authentication policy rule prompts the user to authenticate through Authentication Portal. You can use the following Authentication Portal Authentication Methods:

- Browser challenge—Use Kerberos single sign-on if you want to reduce the number of login prompts that users must respond to.
- Web form—Use Multi-Factor Authentication, SAML single sign-on, Kerberos, TACACS+, RADIUS, LDAP, or Local Authentication.
- Client Certificate Authentication.

Syslog

Your environment might have existing network services that authenticate users. These services include wireless controllers, 802.1x devices, Apple Open Directory servers, proxy servers, and other Network Access Control (NAC) mechanisms. You can configure these services to send syslog messages that contain information about login and logout events and configure the User-ID agent to parse those messages. The User-ID agent parses for login events to map IP addresses to usernames and parses for logout events to delete outdated mappings. Deleting outdated mappings is particularly useful in environments where IP address assignments change often.

Both the PAN-OS integrated User-ID agent and Windows-based User-ID agent use Syslog Parse profiles to parse syslog messages. In environments where services send the messages in different formats, you can create a custom profile for each format and associate multiple profiles with each syslog sender. If you use the PAN-OS integrated User-ID agent, you can also use predefined Syslog Parse profiles that Palo Alto Networks provides through Applications content updates.

Syslog messages must meet the following criteria for a User-ID agent to parse them:

- Each message must be a single-line text string. The allowed delimiters for line breaks are a new line (**\n**) or a carriage return plus a new line (**\r\n**).
- The maximum size for individual messages is 8,000 bytes.
- Messages sent over UDP must be contained in a single packet; messages sent over SSL can span multiple packets. A single packet might contain multiple messages.

See Configure User-ID to Monitor Syslog Senders for User Mapping for configuration details.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

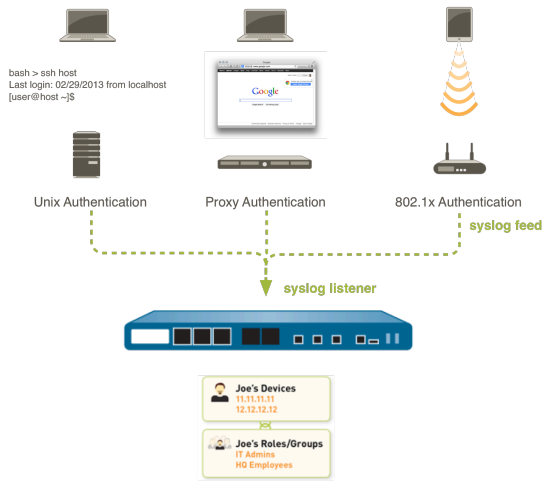


Figure 5: User-ID Integration with Syslog

GlobalProtect

For mobile or roaming users, the GlobalProtect endpoint provides the user mapping information to the firewall directly. In this case, every GlobalProtect user has an app running on the endpoint that requires the user to enter login credentials for VPN access to the firewall. This login information is then added to the User-ID user mapping table on the firewall for visibility and user-based security policy enforcement. Because GlobalProtect users must authenticate to gain access to the network, the IP address-to-username mapping is explicitly known. This is the best solution in sensitive environments where you must be certain of who a user is in order to allow access to an application or service. For more information on setting up GlobalProtect, refer to the GlobalProtect Administrator's Guide.

XML API

Authentication Portal and the other standard user mapping methods might not work for certain types of user access. For example, the standard methods cannot add mappings of users connecting from a third-party VPN solution or users connecting to a 802.1x-enabled wireless network. For such cases, you can use the PAN-OS XML API to capture login events and send them to the PAN-OS integrated User-ID agent. See Send User Mappings to User-ID Using the XML API for details.

Client Probing

Palo Alto Networks strongly recommends disabling client probing because it is not a recommended method of obtaining User-ID information in a high-security network.

Palo Alto Networks does not recommend using client probing due to the following potential risks:

- Because client probing trusts data reported back from the endpoint, it can expose you to security risks when misconfigured. If you enable it on external, untrusted interfaces, this would cause the agent to send client probes containing sensitive information such as the username, domain name, and password hash of the User-ID agent service account outside

https://t.me/learningnets

of your network. If you do not configure the service account correctly, the credentials could potentially be exploited by an attacker to penetrate the network to gain further access.

- Client probing was designed for legacy networks where most users were on Windows workstations on the internal network, but is not ideal for today's more modern networks that support a roaming and mobile user base on a variety of devices and operating systems.
- Client probing can generate a large amount of network traffic (based on the total number of mapped IP addresses).

Instead, Palo Alto Networks strongly recommends using the following alternate methods for user mapping:

- Using more isolated and trusted sources, such as domain controllers and integrations with Syslog or the XML API, to safely capture user mapping information from any device type or operating system.
- Configuring Authentication Policy and Authentication Portal to ensure that you only allow access to authorized users.

The User-ID agent supports two types of client probing:

- NetBIOS probing, which uses the Windows User-ID agent.
- WMI probing, which uses either the PAN-OS integrated User-ID agent or the Windows User-ID agent.



Client probing is not recommended as a user mapping method, but if you plan to enable it, Palo Alto Networks strongly recommends using WMI probing over NetBIOS probing.

In a Microsoft Windows environment, you can configure the User-ID agent to probe client systems using Windows Management Instrumentation (WMI) or NetBIOS probing at regular intervals to verify that an existing user mapping is still valid or to obtain the username for an IP address that is not yet mapped.

If you do choose to enable probing in your trusted zones, the agent will probe each learned IP address periodically (every 20 minutes by default, but this is configurable) to verify that the same user is still logged in. In addition, when the firewall encounters an IP address for which it has no user mapping, it will send the address to the agent for an immediate probe.

See Configure User Mapping Using the Windows User-ID Agent or Configure User Mapping Using the PAN-OS Integrated User-ID Agent for details.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Enable User-ID

The user identity, as opposed to an IP address, is an integral component of an effective security infrastructure. Knowing who is using each of the applications on your network, and who may have transmitted a threat or is transferring files, can strengthen your security policy and reduce incident response times. User-ID enables you to leverage user information stored in a wide range of repositories for visibility, user- and group-based policy control, and improved logging, reporting, and forensics:

STEP 1 | Enable User-ID on the source zones that contain the users who will send requests that require user-based access controls.



Enable User-ID on trusted zones only. If you enable User-ID and client probing on an external untrusted zone (such as the internet), probes could be sent outside your protected network, resulting in an information disclosure of the User-ID agent service account name, domain name, and encrypted password hash, which could allow an attacker to gain unauthorized access to protected services and applications.

1. Select **Network** > **Zones** and click the **Name** of the zone.
2. **Enable User Identification** and click **OK**.

STEP 2 | Create a Dedicated Service Account for the User-ID Agent.



As a best practice, create a service account with the minimum set of permissions required to support the User-ID options you enable to reduce your attack surface in the event that the service account is compromised.

This is required if you plan to use the Windows-based User-ID agent or the PAN-OS integrated User-ID agent to monitor domain controllers, Microsoft Exchange servers, or Windows clients for user login and logout events.

STEP 3 | Map Users to Groups.

This enables the firewall to connect to your LDAP directory and retrieve Group Mapping information so that you will be able to select usernames and group names when creating policy.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 4 | Map IP Addresses to Users.



As a best practice, do not enable client probing as a user mapping method on high-security networks. Client probing can generate a large amount of network traffic and can pose a security threat when misconfigured.

The way you do this depends on where your users are located and what types of systems they are using, and what systems on your network are collecting login and logout events for your users. You must configure one or more User-ID agents to enable User Mapping:

- Configure User Mapping Using the Windows User-ID Agent.
- Configure User Mapping Using the PAN-OS Integrated User-ID Agent.
- Configure User-ID to Monitor Syslog Senders for User Mapping.
- Configure User Mapping for Terminal Server Users.
- Send User Mappings to User-ID Using the XML API.
- Insert Username in HTTP Headers.

STEP 5 | Specify the networks to include and exclude from user mapping.



As a best practice, always specify which networks to include and exclude from User-ID. This allows you to ensure that only your trusted assets are probed and that unwanted user mappings are not created unexpectedly.

The way you specify which networks to include and exclude depends on whether you are using the Windows-based User-ID agent or the PAN-OSintegrated User-ID agent.

STEP 6 | Configure Authentication Policy and Authentication Portal.

The firewall uses Authentication Portal to authenticate end users when they request services, applications, or URL categories that match Authentication Policy rules. Based on user information collected during authentication, the firewall creates new user mappings or updates existing mappings. The mapping information collected during authentication overrides information collected through other User-ID methods.

1. Configure Authentication Portal.
2. Configure Authentication Policy.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 7 | Enable user- and group-based policy enforcement.



Create rules based on group rather than user whenever possible. This prevents you from having to continually update your rules (which requires a commit) whenever your user base changes.

After configuring User-ID, you will be able to choose a username or group name when defining the source or destination of a security rule:

1. Select **Policies** > **Security** and **Add** a new rule or click an existing rule name to edit.
2. Select **User** and specify which users and groups to match in the rule in one of the following ways:
   - If you want to select specific users or groups as matching criteria, click **Add** in the Source User section to display a list of users and groups discovered by the firewall group mapping function. Select the users or groups to add to the rule.
   - If you want to match any user who has or has not authenticated and you don't need to know the specific user or group name, select **known-user** or **unknown** from the drop-down above the Source User list.
3. Configure the rest of the rule as appropriate and then click **OK** to save it. For details on other fields in the security rule, see Set Up a Basic Security Policy.

STEP 8 | Create the Security policy rules to safely enable User-ID within your trusted zones and prevent User-ID traffic from egressing your network.

Follow the Best Practice Internet Gateway Security Policy to ensure that the User-ID application (`paloalto-userid-agent`) is only allowed in the zones where your agents (both your Windows agents and your PAN-OS integrated agents) are monitoring services and distributing mappings to firewalls. Specifically:

- Allow the `paloalto-userid-agent` application between the zones where your agents reside and the zones where the monitored servers reside (or even better, between the specific systems that host the agent and the monitored servers).
- Allow the `paloalto-userid-agent` application between the agents and the firewalls that need the user mappings and between firewalls that are redistributing user mappings and the firewalls they are redistributing the information to.
- Deny the `paloalto-userid-agent` application to any external zone, such as your internet zone.

STEP 9 | Configure the firewall to obtain user IP addresses from X-Forwarded-For (XFF) headers.

When the firewall is between the Internet and a proxy server, the IP addresses in the packets that the firewall sees are for the proxy server rather than users. To enable visibility of user IP addresses instead, configure the firewall to use the XFF headers for user mapping. With this option enabled, the firewall matches the IP addresses with usernames referenced in policy to

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

enable control and visibility for the associated users and groups. For details, see Identify Users Connected through a Proxy Server.

1. Select **Device** > **Setup** > **Content-ID** and edit the X-Forwarded-For Headers settings.
2. Select **X-Forwarded-For Header in User-ID**.

   *Selecting **Strip-X-Forwarded-For Header** doesn't disable the use of XFF headers for user attribution in policy rules; the firewall zeroes out the XFF value only after using it for user attribution.*
3. Click **OK** to save your changes.

STEP 10 | If you use a high availability (HA) configuration, enable synchronization.

*As a best practice, always enable the **Enable Config Sync** option for an HA configuration to ensure that the group mappings and user mappings are synchronized between the active and passive firewall.*

1. Select **Device** > **High Availability** > **General** and edit the Setup section.
2. Select **Enable HA**.
3. Select **Enable Config Sync**.
4. Enter the **Peer HA1 IP Address**, which is the IP address of the HA1 control link on the peer firewall.
5. (Optional) Enter a **Backup Peer HA1 IP Address**, which is the IP address of the backup control link on the peer firewall.
6. Click **OK**.

STEP 11 | Commit your changes.

Commit your changes to activate them.

STEP 12 | Verify the User-ID Configuration.

After you configure user mapping and group mapping, verify that the configuration works properly and that you can safely enable and monitor user and group access to your applications and services.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Map Users to Groups

Defining policy rules based on user group membership rather than individual users simplifies administration because you don't have to update the rules whenever group membership changes. The number of distinct user groups that each firewall or Panorama can reference across all policies varies by model. For more information, refer to the Compatibility Matrix.

Use the following procedure to enable the firewall to connect to your LDAP directory and retrieve Group Mapping information. You can then Enable User- and Group-Based Policy.



The following are best practices for group mapping in an Active Directory (AD) environment:

- *If you have a single domain, you need only one group mapping configuration with an LDAP server profile that connects the firewall to the domain controller with the best connectivity. You can add up to four domain controllers to the LDAP server profile for redundancy. Note that you cannot increase redundancy beyond four domain controllers for a single domain by adding multiple group mapping configurations for that domain.*
- *If you have multiple domains and/or multiple forests, you must create a group mapping configuration with an LDAP server profile that connects the firewall to a domain server in each domain/forest. Take steps to ensure unique usernames in separate forests.*
- *If you have Universal Groups, create an LDAP server profile to connect to the root domain of the Global Catalog server on port 3268 or 3269 for SSL, then create another LDAP server profile to connect to the root domain controllers on port 389. This helps ensure that users and group information is available for all domains and subdomains.*
- *Before using group mapping, configure a* ***Primary Username*** *for user-based security policies, since this attribute will identify users in the policy configuration, logs, and reports.*

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 1 | Add an LDAP server profile.

The profile defines how the firewall connects to the directory servers from which it collects group mapping information.

If you create multiple group mapping configurations that use the same base distinguished name (DN) or LDAP server, the group mapping configurations cannot contain overlapping groups (for example, the Include list for one group mapping configuration cannot contain a group that is also in a different group mapping configuration).

1. Select **Device** > **Server Profiles** > **LDAP** and **Add** a server profile.
2. Enter a **Profile Name** to identify the server profile.
3. **Add** the LDAP servers. You can add up to four servers to the profile but they must be the same **Type**. For each server, enter a **Name** (to identify the server), **LDAP Server** IP address or FQDN, and server **Port** (default 389).
4. Select the server **Type**.

   Based on your selection (such as **active-directory**), the firewall automatically populates the correct LDAP attributes in the group mapping settings. However, if you customized your LDAP schema, you might need to modify the default settings.
5. For the **Base DN**, enter the Distinguished Name (DN) of the LDAP tree location where you want the firewall to start searching for user and group information.
6. For the **Bind DN**, **Password** and **Confirm Password**, enter the authentication credentials for binding to the LDAP tree.

   The **Bind DN** can be a fully qualified LDAP name (such as `cn=administrator,cn=users,dc=acme,dc=local`) or a user principal name (such as `administrator@acme.local`).
7. Enter the **Bind Timeout** and **Search Timeout** in seconds (default is 30 for both).
8. Click **OK** to save the server profile.

STEP 2 | Configure the server settings in a group mapping configuration.

1. Select **Device** > **User Identification** > **Group Mapping Settings**.
2. **Add** the group mapping configuration.
3. Enter a unique **Name** to identify the group mapping configuration.
4. Select the LDAP **Server Profile** you just created.
5. (Optional) Specify the **Update Interval** (in seconds). Enter a value (range is 60—86400, default is 3600) based on how often the firewall should be check the LDAP source for updates to the group mapping configuration. If the LDAP source contains many groups, a value that is too low may not allow enough time to map all the groups.
6. (Optional) By default, the **User Domain** field is blank: the firewall automatically detects the domain names for Active Directory (AD) servers. If you enter a value, it overrides any domain names that the firewall retrieves from the LDAP source. For most configurations,

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

if you need to enter a value, enter the NetBIOS domain name (for example, **example** not **example.com**).

If you use Global Catalog, entering a value replaces the domain name for all users and groups from this server, including those from other domains.

7. (Optional) To filter the groups that the firewall tracks for group mapping, in the Group Objects section, enter a **Search Filter** (LDAP query) and **Object Class** (group definition).
8. (Optional) To filter the users that the firewall tracks for group mapping, in the User Objects section, enter a **Search Filter** (LDAP query), and **Object Class** (user definition).
9. Make sure the group mapping configuration is **Enabled** (default is enabled).

STEP 3 | (Optional) Define User and Group Attributes to collect for user and group mapping. This step is required if you want to map users based on directory attributes other than the domain.

1. If your User-ID sources only send the username and the username is unique across the organization, select **Device** > **User Identification** > **User Mapping** > **Setup** and **Edit** the Setup section to **Allow matching usernames without domains** to allow the firewall to check if unique usernames collected from the LDAP server during group mapping match the users associated with a policy and avoid overwriting the domain in your source profile.

   Before enabling this option, configure group mapping for the LDAP group containing the User-ID source (such as GlobalProtect or Authentication Portal) that collects the mappings. After you commit the changes, the User-ID source populates the usernames without domains. Only usernames collected during group mapping can be matched without a domain. If your User-ID sources send user information in multiple formats and you enable this option, verify that the attributes collected by the firewall have a unique prefix. To ensure users are identified correctly if you enable this option, all attributes for group mapping should be unique. If the username is not unique, the firewall logs an error in the Debug logs.

2. Select **Device** > **User Identification** > **Group Mapping Settings** > **Add** > **User and Group Attributes** > **User Attributes** and enter the **Directory Attribute** you want to collect for user identification. Specify a **Primary Username** to identify the user on the firewall and

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

to represent the user in reports and logs that will override any other format the firewall receives from the User-ID source.

When you select the Server Profile **Type**, the firewall auto-populates the values for the user and group attributes. Based on the user information that your User-ID sources send, you may need to configure the correct attributes:

- **User Principal Name (UPN)**: **`userPrincipalName`**
- **NetBios Name**: **`sAMAccountName`**
- **Email ID**: Directory attribute for that email
- **Multiple formats**: Retrieve the user mapping attributes from the user directory before enabling your User-ID sources.

If you do not specify a primary username, the firewall uses the following default values for each server profile type:

Attribute	Active Directory	Novell eDirectory or Sun ONE Directory Server
Primary Username	`sAMAccountName`	`uid`
E-Mail	`mail`	`mail`
Alternate Username 1	`userPrincipalName`	None.
Group Name	`name`	`cn`
Group Member	`member`	`member`

3. (Optional) Specify an **E-Mail** address format and up to three **Alternate Username** formats.
4. Select **Device** > **User Identification** > **Group Mapping Settings** > **Add** > **User and Group Attributes** > **Group Attributes** and specify the **Group Name**, **Group Member**, and **E-Mail** address formats.

   You must commit before the firewall collects the directory attributes from the LDAP server.

STEP 4 | Limit which groups will be available in policy rules.

Required only if you want to limit policy rules to specific groups. The combined maximum for the **Group Include List** and **Custom Group** list is 640 entries per group mapping configuration.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Each entry can be a single group or a list of groups. By default, if you don't specify groups, all groups are available in policy rules.

Any custom groups you create will also be available in the Allow List of authentication profiles (Configure an Authentication Profile and Sequence).

1. Add existing groups from the directory service:
   1. Select **Group Include List**.
   2. Select the Available Groups you want to appear in policy rules and add (⊕) them to the Included Groups.
2. If you want to base policy rules on user attributes that don't match existing user groups, create custom groups based on LDAP filters:
   1. Select **Custom Group** and **Add** the group.
   2. Enter a group **Name** that is unique in the group mapping configuration for the current firewall or virtual system.

      If the **Name** has the same value as the Distinguished Name (DN) of an existing AD group domain, the firewall uses the custom group in all references to that name (such as in policies and logs).
   3. Specify an **LDAP Filter** of up to 2,048 UTF-8 characters and click **OK**.

      The firewall doesn't validate LDAP filters, so it's up to you to ensure they are accurate.

      To minimize the performance impact on the LDAP directory server, use only indexed attributes in the filter.
3. Click **OK** to save your changes.

   You must commit before custom groups will be available in policies and objects.

STEP 5 | **Commit** your changes.

You must commit before you can use custom groups in policies and objects and before the firewall can collect the attributes from the LDAP server.

After configuring the firewall to retrieve group mapping information from an LDAP server, but before configuring policies based on the groups it retrieves, the best practice is to either wait for the firewall to refresh its group mappings cache or refresh the cache manually. To verify which groups you can currently use in policies, access the firewall CLI and run the ***`show user group`*** *command. To determine when the firewall will next refresh the group mappings cache, run the* ***`show user group-mapping statistics`*** *command and check the `Next Action`. To manually refresh the cache, run the* ***`debug user-id refresh group-mapping all`*** *command.*

https://t.me/learningnets

STEP 6 | Verify that the user and group mapping has correctly identified users.

1. Select **Device** > **User Identification** > **Group Mapping** > **Group Include List** to confirm the firewall has fetched all of the groups.
2. To verify that all of the user attributes have been correctly captured, use the following CLI command:

   ```
   show user user-attributes user all
   ```

   The normalized format for the User Principal Name (UPN), primary username, email attributes, and any configured alternate usernames display for all users:

   ```
   admin@PA-VM-8.1> show user user-attributes user all
   ```

   ```
   Primary: nam\sam-user   Email: sam-user@nam.com
   ```

   ```
   Alt User Names:1) nam.com\sam-user
   ```

   ```
   2) nam\sam-user-upn
   ```

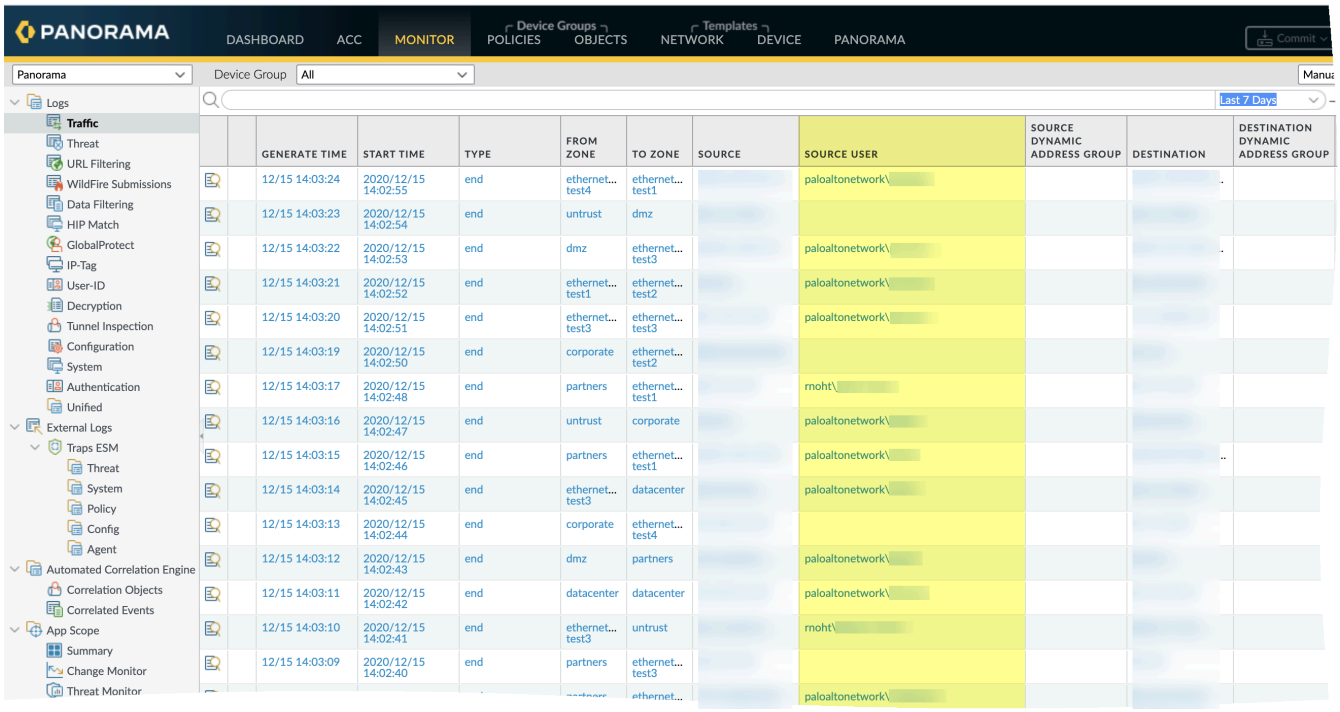
   ```
   3) sam-user-upn@nam.local
   ```

   ```
   4) sam-user@nam.com
   ```

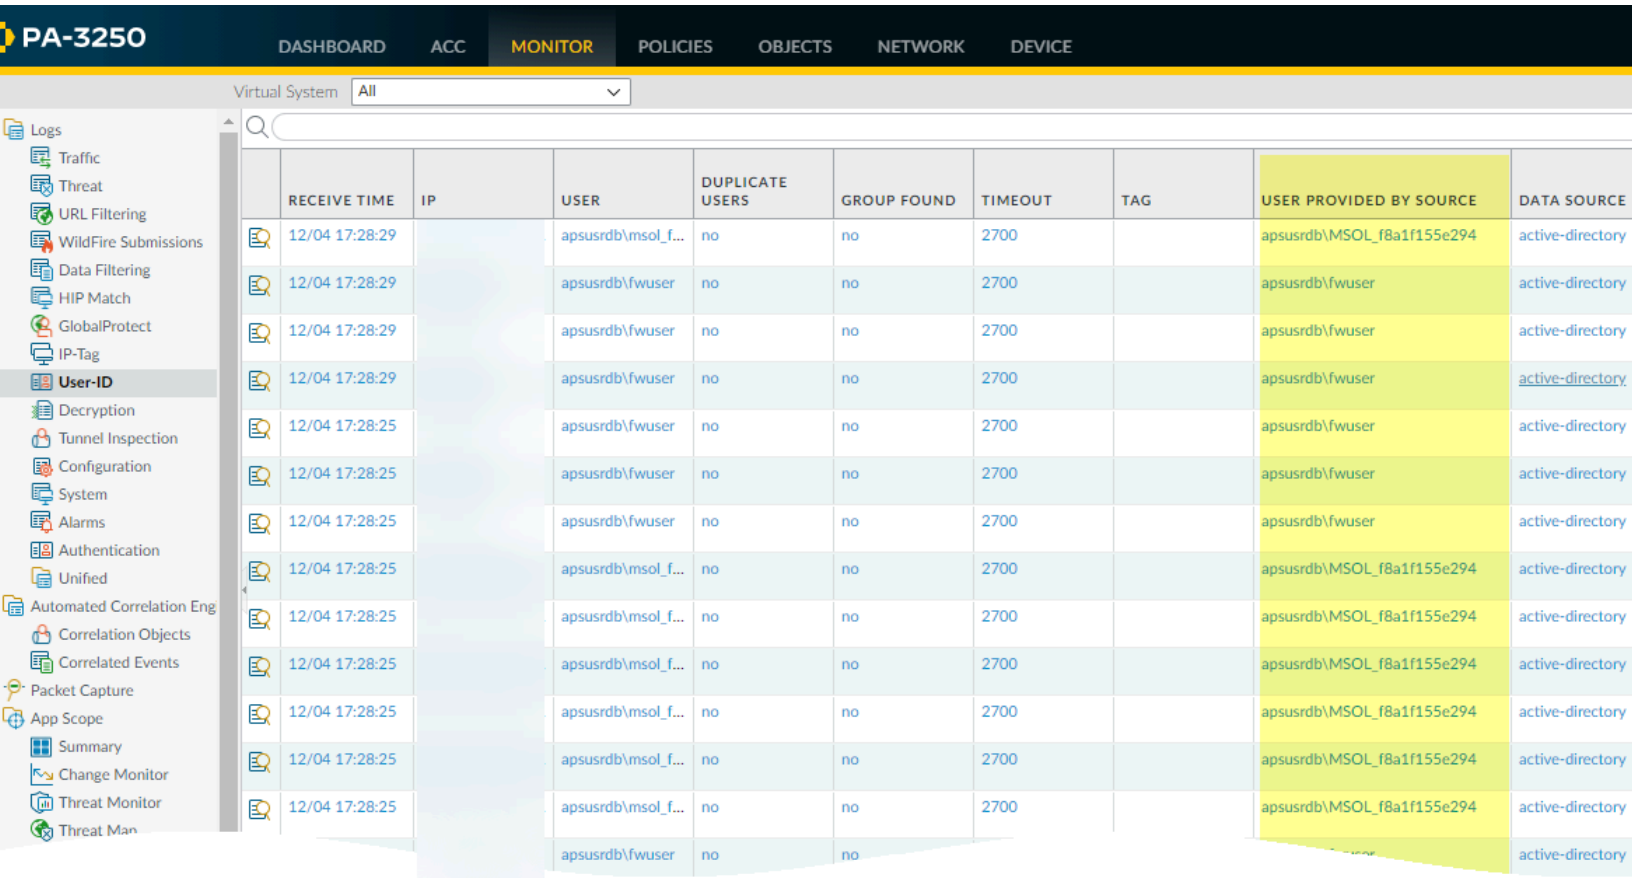
3. Verify that the usernames are correctly displayed in the **Source User** column under **Monitor** > **Logs** > **Traffic**.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets



4. Verify that the users are mapped to the correct usernames in the **User Provided by Source** column under **Monitor** > **Logs** > **User-ID**.



https://t.me/learningnets

Map IP Addresses to Users

User-ID provides many different methods for mapping IP addresses to usernames. Before you begin configuring user mapping, consider where your users are logging in from, what services they are accessing, and what applications and data you need to control access to. This will inform which types of agents or integrations would best allow you to identify your users.

Once you have your plan, you can begin configuring user mapping using one or more of the following methods as needed to enable user-based access and visibility to applications and resources:

- If you have users with client systems that aren't logged in to your domain servers—for example, users running Linux clients that don't log in to the domain—you can Map IP Addresses to Usernames Using Authentication Portal. Using Authentication Portal in conjunction with Authentication Policy also ensures that all users authenticate to access your most sensitive applications and data.
- To map users as they log in to your Exchange servers, domain controllers, eDirectory servers, or Windows clients you must configure a User-ID agent:
  - Configure User Mapping Using the PAN-OS Integrated User-ID Agent
  - Configure User Mapping Using the Windows User-ID Agent
- If you have clients running multi-user systems in a Windows environment, such as Microsoft Terminal Server or Citrix Metaframe Presentation Server or XenApp, Configure the Palo Alto Networks Terminal Server (TS) Agent for User Mapping. For a multi-user system that doesn't run on Windows, you can Retrieve User Mappings from a Terminal Server Using the PAN-OS XML API.
- To obtain user mappings from existing network services that authenticate users—such as wireless controllers, 802.1x devices, Apple Open Directory servers, proxy servers, or other Network Access Control (NAC) mechanisms—Configure User-ID to Monitor Syslog Senders for User Mapping.

  While you can configure either the Windows agent or the PAN-OS integrated User-ID agent on the firewall to listen for authentication syslog messages from the network services, because only the PAN-OS integrated agent supports syslog listening over TLS, it is the preferred configuration.
- To include the username and domain in the headers for outgoing traffic so other devices in your network can identify the user and enforce user-based policy, you can Insert Username in HTTP Headers.
- To Share User-ID Mappings Across Virtual Systems, you can configure a virtual system as a User-ID hub.
- For other clients that you can't map using the other methods, you can Send User Mappings to User-ID Using the XML API.
- A large-scale network can have hundreds of information sources that firewalls query for user and group mapping and can have numerous firewalls that enforce policies based on the mapping information. You can simplify User-ID administration for such a network by aggregating the mapping information before the User-ID agents collect it. You can also reduce the resources that the firewalls and information sources use in the querying process by

https://t.me/learningnets

configuring some firewalls to redistribute the mapping information. For details, see Deploy User-ID in a Large-Scale Network.

Create a Dedicated Service Account for the User-ID Agent

To use the Windows-based User-ID agent or the PAN-OS integrated User-ID agent to map users as they log in to your Exchange servers, domain controllers, eDirectory servers, or Windows clients, create a dedicated service account for the User-ID agent on a domain controller in each domain that the agent will monitor.

The User-ID agent maps users based on logs for security events. To ensure that the User-ID agent can successfully map users, verify that the source for your mappings generates logs for Audit Logon, Audit Kerberos Authentication Service, and Audit Kerberos Service Ticket Operations events. At a minimum, the source must generate logs for the following events:

- Logon Success (4624)
- Authentication Ticket Granted (4768)
- Service Ticket Granted (4769)
- Ticket Granted Renewed (4770)

The required permissions for the service account depend on the user mapping methods and settings you plan to use. For example, if you are using the PAN-OS integrated User-ID agent, the service account requires Server Operator privileges to monitor user sessions. If you are using the Windows-based User-ID agent, the service account does not require Server Operator privileges to monitor user sessions. To reduce the risk of compromising the User-ID service account, always configure the account with the minimum set of permissions necessary for the agent.

- If you are installing the Windows-based User-ID agent on a supported Windows server, Configure a Service Account for the Windows User-ID Agent.
- If you are using the PAN-OS integrated User-ID agent on the firewall, Configure a Service Account for the PAN-OS Integrated User-ID Agent.

User-ID provides many methods for safely collecting user mapping information. Some legacy features designed for environments that only required user mapping on Windows desktops attached to the local network require privileged service accounts. If the privileged service account is compromised, this would open your network to attack. As a best practice, avoid using legacy features that require privileges that would pose a threat if compromised, such as client probing and session monitoring.

Configure a Service Account for the Windows User-ID Agent

Create a dedicated Active Directory (AD) service account for the Windows User-ID agent to access the services and hosts it will monitor to collect user mappings. You must create a service account in each domain the agent will monitor. After you enable the required permissions for the service account, Configure User Mapping Using the Windows User-ID Agent.

The following workflow details all required privileges and provides guidance for the User-ID features which require privileges that could pose a threat so that you can decide how to best identify users without compromising your overall security posture.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 1 | Create an AD service account for the User-ID agent.

You must create a service account in each domain the agent will monitor.

1. Log in to the domain controller.
2. Right-click the Windows icon (), **Search** for **`Active Directory Users and Computers`**, and launch the application.
3. In the navigation pane, open the domain tree, right-click **Managed Service Accounts** and select **New** > **User**.
4. Enter the **First Name**, **Last Name**, and **User logon name** of the user and click **Next**.
5. Enter the **Password** and **Confirm Password**, then click **Next** and **Finish**.

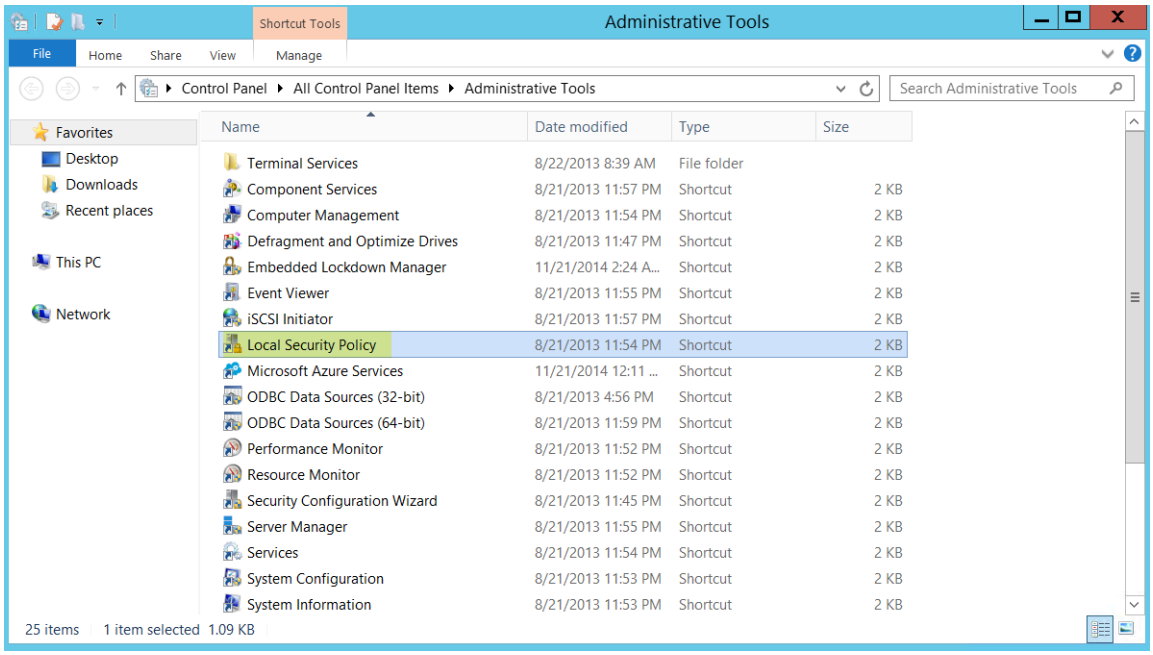
 ©2022 Palo Alto Networks, Inc.

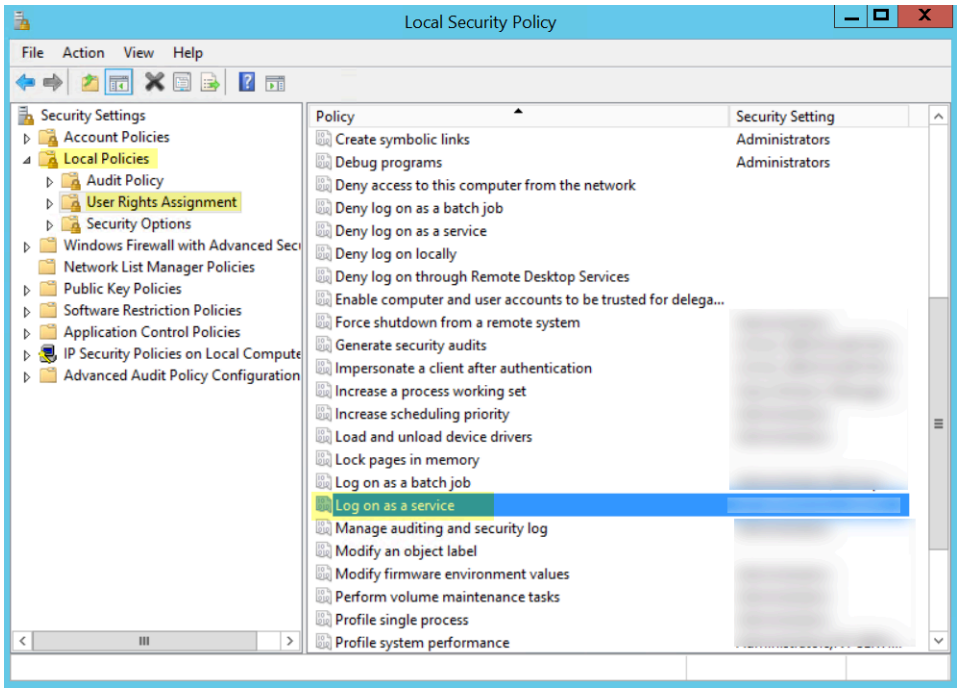
https://t.me/learningnets

STEP 2 | Configure either local or group policy to allow the service account to log on as a service.

The permission to log on as a service is only needed locally on the Windows server that is the agent host.

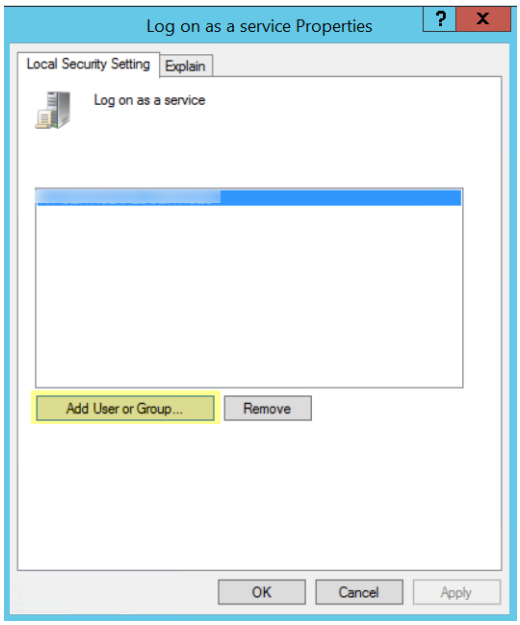
- To assign permissions locally:

1. select **Control Panel** > **Administrative Tools** > **Local Security Policy**.
2. 

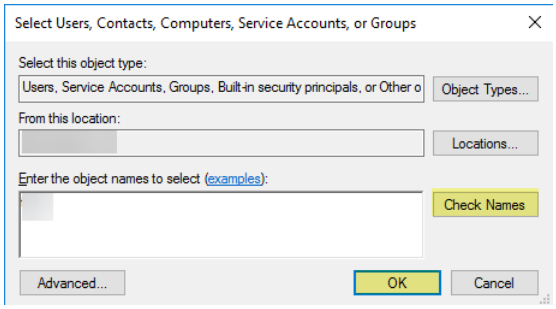
3. Select **Local Policies** > **User Rights Assignment** > **Log on as a service**.


4. **Add User or Group** to add the service account.

https://t.me/learningnets

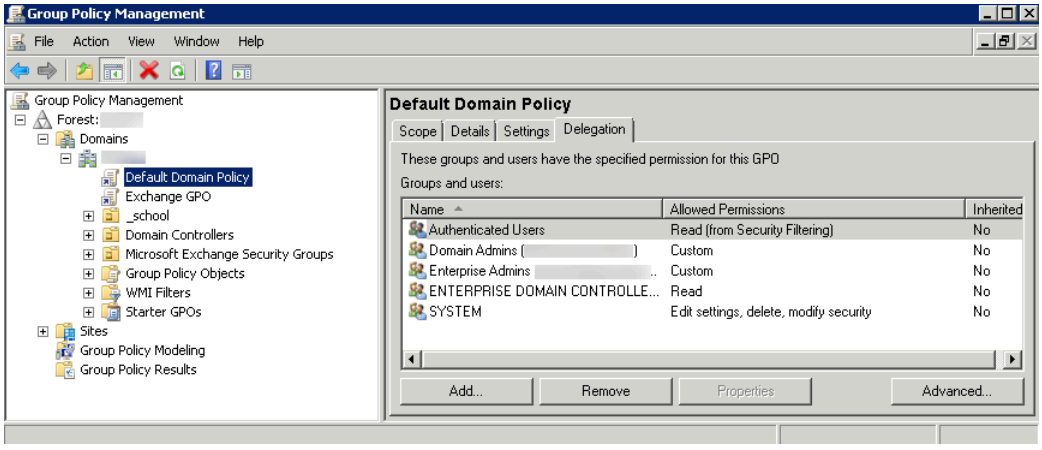


5. **Enter the object names to select** (the service account name) in **`domain\username`** format and click **OK**.



- To configure group policy if you are installing Windows User-ID agents on multiple servers, use the Group Policy Management Editor.

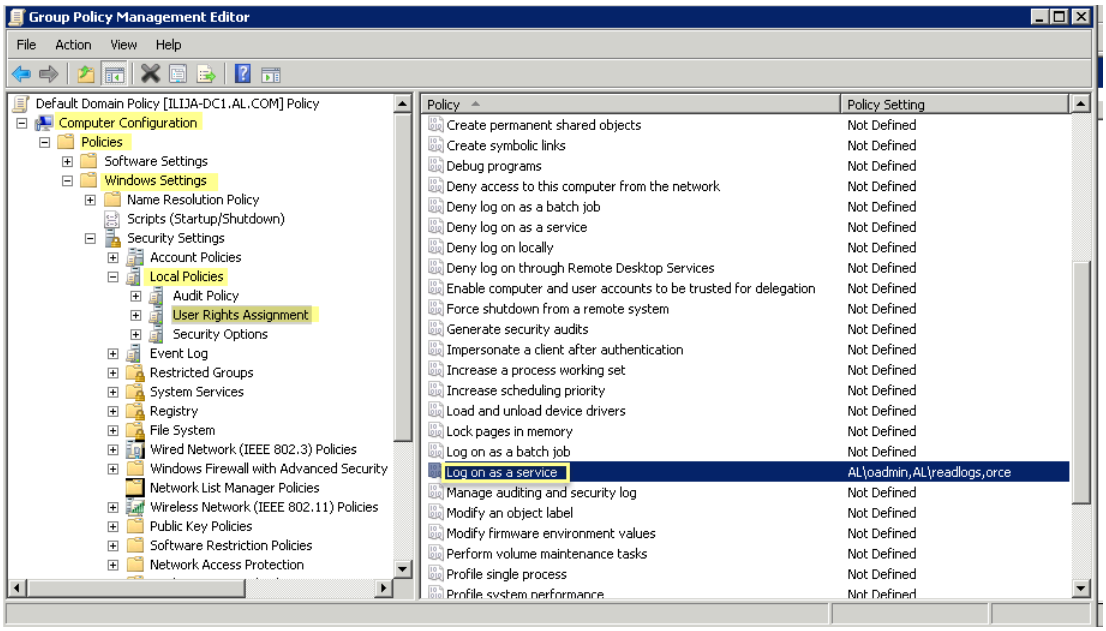
1. Select **Start** > **Group Policy Management** > **<your domain>** > **Default Domain Policy** > **Action** > **Edit** for the Windows server that is the agent host.



2. Select **Computer Configuration** > **Policies** > **Windows Settings** > **Security Settings** > **Local Policies** > **User Rights Assignment**.

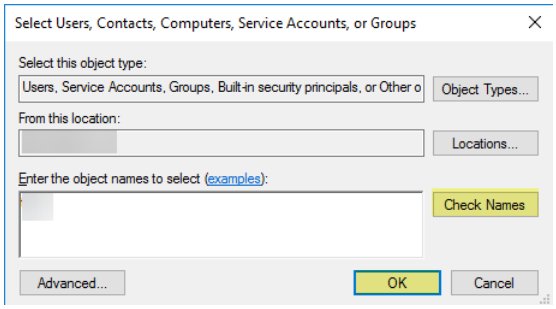
 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets



3. Right-click **Log on as a service**, then select **Properties**.
4. **Add User or Group** to add the service account username or builtin group, then click **OK** twice.

   Administrators have this privilege by default.



STEP 3 | If you want to use WMI to collect user data, assign DCOM privileges to the service account so that it can use WMI queries on monitored servers.

1. Select **Active Directory Users and Computers** > **<your domain>** > **Builtin** > **Distributed COM Users**.
2. Right-click **Properties** > **Members** > **Add** and enter the service account name.

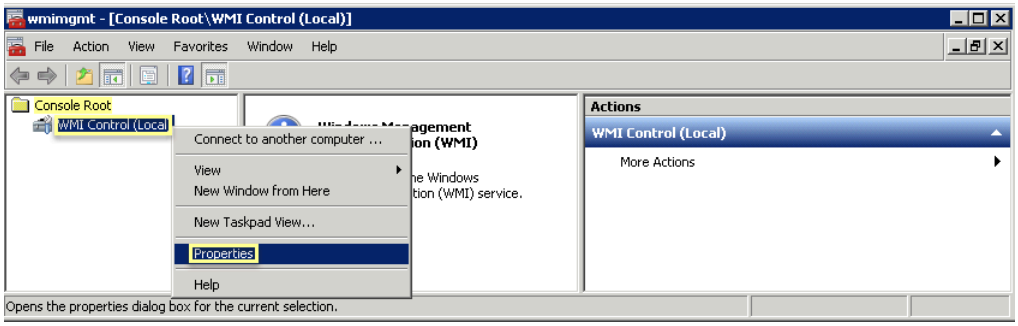
https://t.me/learningnets

STEP 4 | If you plan to use WMI probing, enable the account to read the CIMV2 namespace and assign the required permissions on the client systems to be probed.

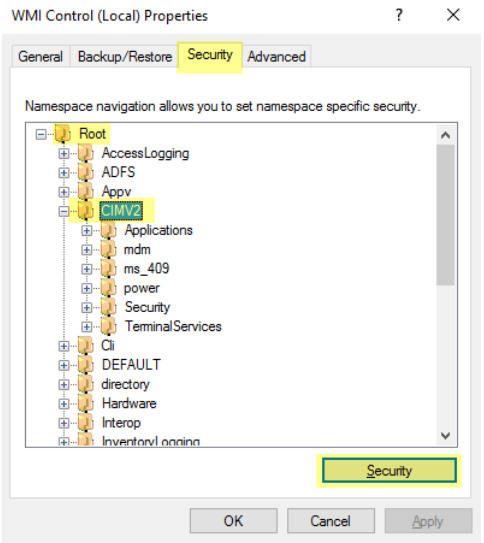
Do not enable client probing on high-security networks. Client probing can generate a large amount of network traffic and can pose a security threat when misconfigured. Instead collect user mapping information from more isolated and trusted sources, such as domain controllers and through integrations with Syslog or the XML API, which have the added benefit of allowing you to safely capture user mapping information from any device type or operating system, instead of just Windows clients.

Perform this task on each client system that the User-ID agent will probe for user mapping information:

1. Right-click the Windows icon (), **Search** for **`wmimgmt.msc`**, and launch the WMI Management Console.
2. In the console tree, right-click **WMI Control** and select **Properties**.



3. Select the **Security** tab, then select **Root** > **CIMV2**, and click the **Security** button.

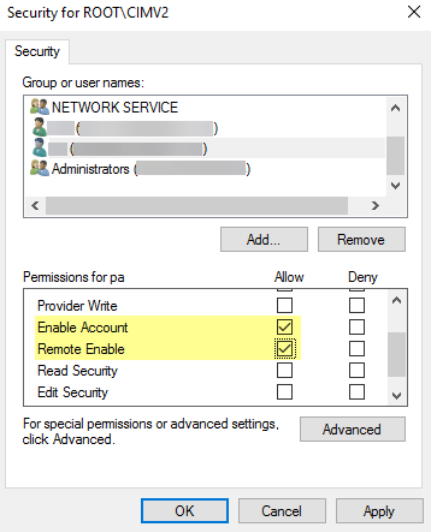


4. **Add** the name of the service account you created, **Check Names** to verify your entry, and click **OK**.

*You might have to change the **Locations** or click **Advanced** to query for account names. See the dialog help for details.*

https://t.me/learningnets

5. In the Permissions for **<Username>** section, **Allow** the **Enable Account** and **Remote Enable** permissions.



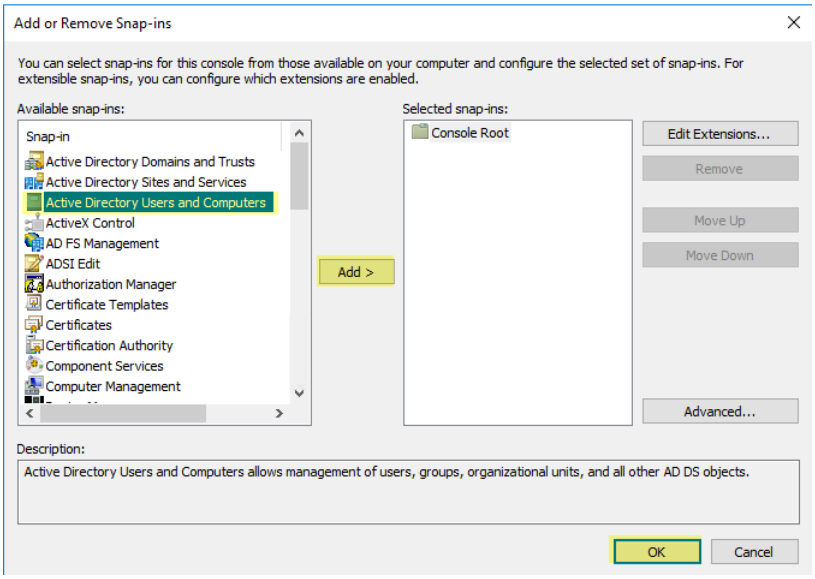
6. Click **OK** twice.
7. Use the Local Users and Groups MMC snap-in (lusrmgr.msc) to add the service account to the local Distributed Component Object Model (DCOM) Users and Remote Desktop Users groups on the system that will be probed.

 ©2022 Palo Alto Networks, Inc.

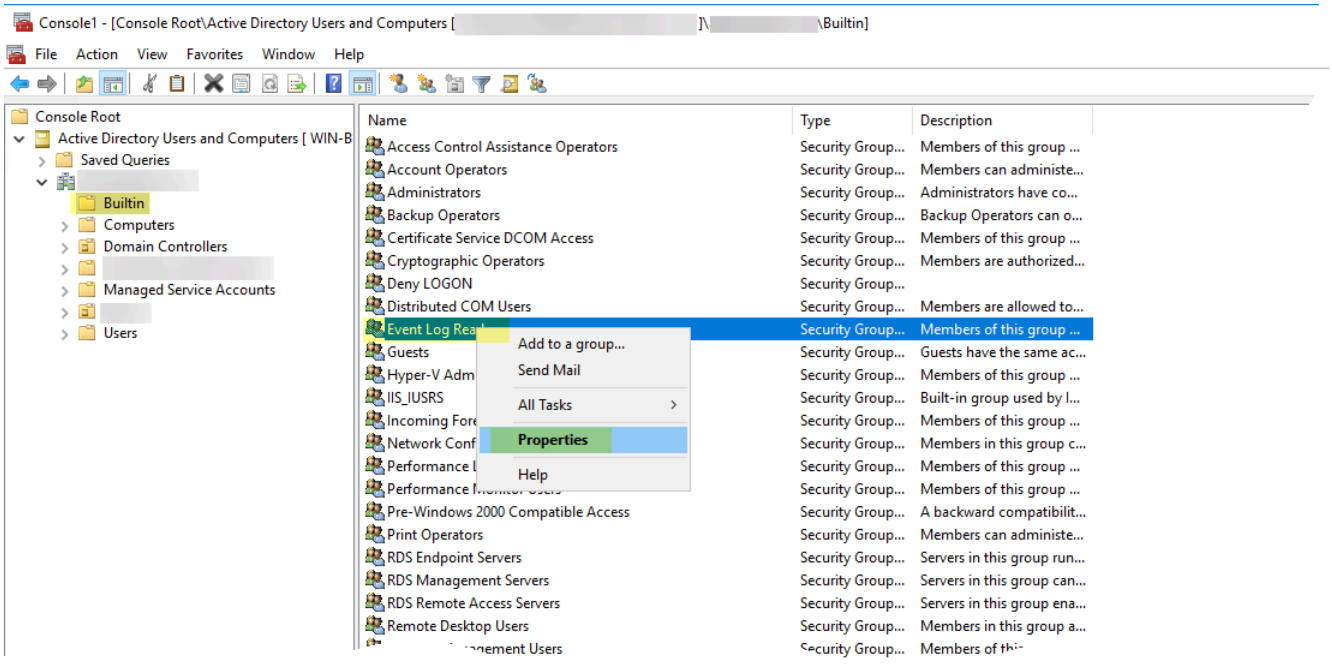
https://t.me/learningnets

STEP 5 | If you want to use Server Monitoring to identify users, add the service account to the Event Log Reader builtin group to allow the service account to read the security log events.

1. On the domain controller or Exchange server that contains the logs you want the User-ID agent to read, or on the member server that receives events from Windows log forwarding, select **Start** > **Run**, enter **MMC**.
2. Select **File** > **Add/Remove Snap-in** > **Active Directory Users and Computers** > **Add**, then click **OK** to run the MMC and launch the Active Directory Users and Computers snap-in.

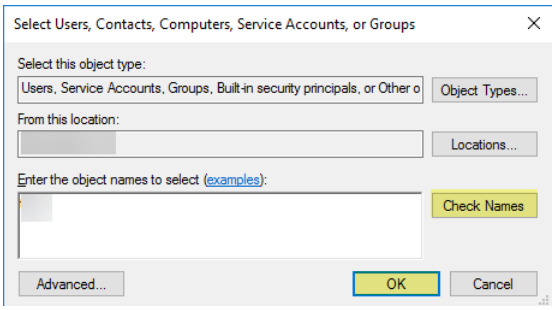


3. Navigate to the Builtin folder for the domain, right-click the **Event Log Readers** group, and select **Properties** > **Members**.



4. **Add** the service account then click **Check Names** to validate that you have the proper object name.

https://t.me/learningnets



5. Click **OK** twice to save the settings.
6. Confirm that the builtin Event Log Reader group lists the service account as a member (**Event Log Readers** > **Properties** > **Members**).

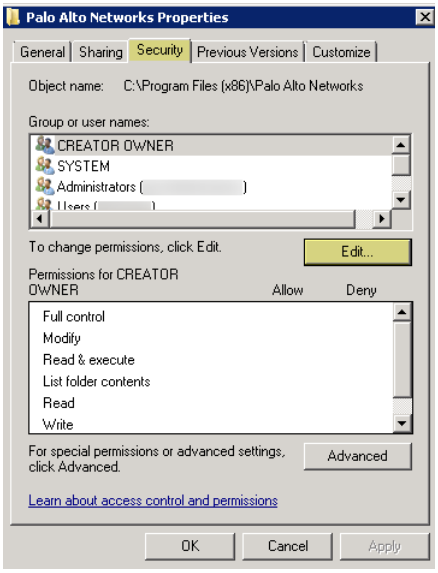
 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

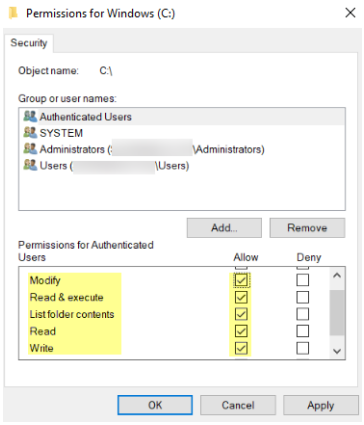
STEP 6 | Assign account permissions to the installation folder to allow the service account to access the agent's installation folder to read the configuration and write logs.

You only need to perform this step if the service account you configured for the User-ID agent is not either a domain administrator or a local administrator on the User-ID agent server host.

1. From the Windows Explorer, navigate to **`C:\Program Files(x86)\Palo Alto Networks`**, right-click the folder, and select **Properties**.
2. On the **Security** tab, click **Edit**.



3. **Add** the User-ID agent service account and **Allow** permissions to **Modify**, **Read & execute**, **List folder contents**, **Read**, and **Write**, and then click **OK** to save the account settings.



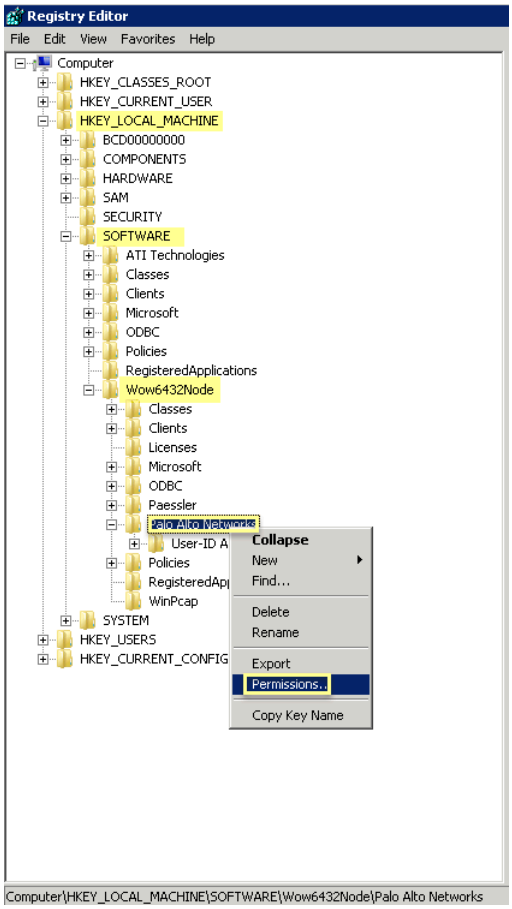
*If you do not want to configure individual permissions, you can **Allow** the **Full Control** permission instead.*

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 7 | To allow the agent to make configuration changes (for example, if you select a different logging level), give the service account permissions to the User-ID agent registry sub-tree.

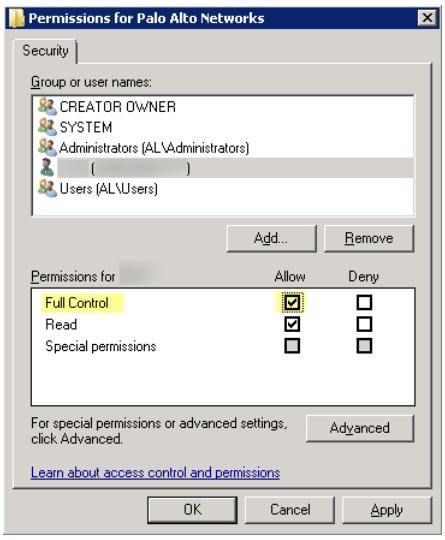
1. Select **Start** > **Run** and enter **`regedt32`** and navigate to the Palo Alto Networks sub-tree in one of the following locations:
   - **32-bit systems**—`HKEY_LOCAL_MACHINE\Software\Palo Alto Networks`
   - **64-bit systems**—`HKEY_LOCAL_MACHINE\Software\WOW6432Node\PaloAlto Networks`
2. Right-click the **Palo Alto Networks** node and select **Permissions**.



3. Assign the User-ID service account **Full Control** and then click **OK** to save the setting.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets



STEP 8 | Disable service account privileges that are not required.

By ensuring that the User-ID service account has the minimum set of account privileges, you can reduce the attack surface should the account be compromised.

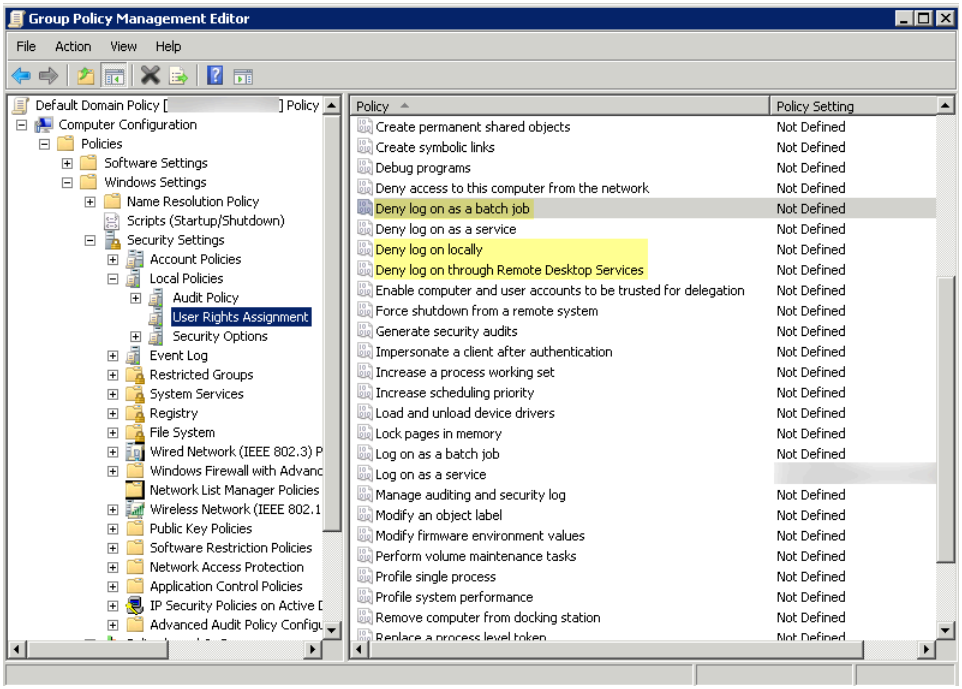
To ensure that the User-ID account has the minimum privileges necessary, deny the following privileges on the account.

- **Deny interactive logon for the User-ID service account**—While the User-ID service account does need permission to read and parse Active Directory security event logs, it does not require the ability to logon to servers or domain systems interactively. You can restrict this

https://t.me/learningnets

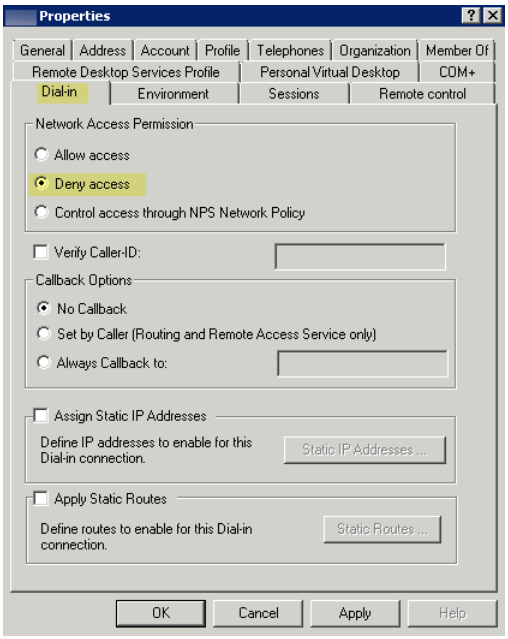
privilege using Group Policies or by using a Managed Service account (refer to Microsoft TechNet for more information).

1. Select **Group Policy Management Editor** > **Default Domain Policy** > **Computer Configuration** > **Policies** > **Windows Settings** > **Security Settings** > **User Rights Assignment**.
2. For **Deny log on as a batch job**, **Deny log on locally**, and **Deny log on through Remote Desktop Services**, right-click **Properties**.
3. Select **Define these policy settings** > **Add User or Group** and add the service account name, then click **OK**.



- **Deny remote access for the User-ID service account**—This prevents an attacker from using the account to access your network from the outside the network.
  1. Select **Start** > **Run**, enter **MMC**, and select **File** > **Add/Remove Snap-in** > **Active Directory Users and Computers** > **Users**.
  2. Right-click the service account name, then select **Properties**.
  3. Select **Dial-in**, then **Deny** the **Network Access Permission**.

https://t.me/learningnets



STEP 9 | As a next step, Configure User Mapping Using the Windows User-ID Agent.

Configure a Service Account for the PAN-OS Integrated User-ID Agent

Create a dedicated Active Directory (AD) service account for the PAN-OS Integrated User-ID agent to access the services and hosts it will monitor to collect user mappings.You must create a service account in each domain the agent will monitor. After you enable the required permissions for the service account, Configure User Mapping Using the PAN-OS Integrated User-ID Agent.

The following workflow details all required privileges and provides guidance for the User-ID features which require privileges that could pose a threat so that you can decide how to best identify users without compromising your overall security posture.

STEP 1 | Create an AD service account for the User-ID agent.

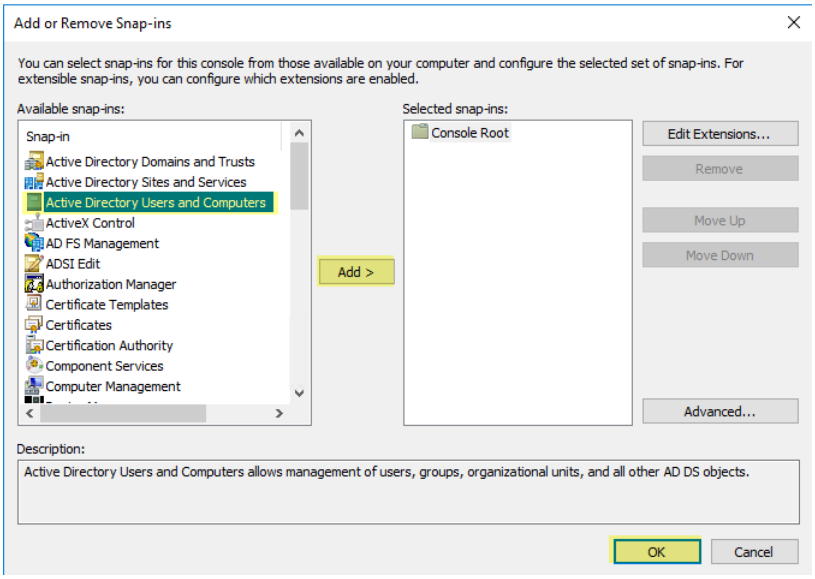
You must create a service account in each domain the agent will monitor.

1. Log in to the domain controller.
2. Right-click the Windows icon (), **Search** for **`Active Directory Users and Computers`**, and launch the application.
3. In the navigation pane, open the domain tree, right-click **Managed Service Accounts** and select **New** > **User**.
4. Enter the **First Name**, **Last Name**, and **User logon name** of the user and click **Next**.
5. Enter the **Password** and **Confirm Password**, then click **Next** and **Finish**.

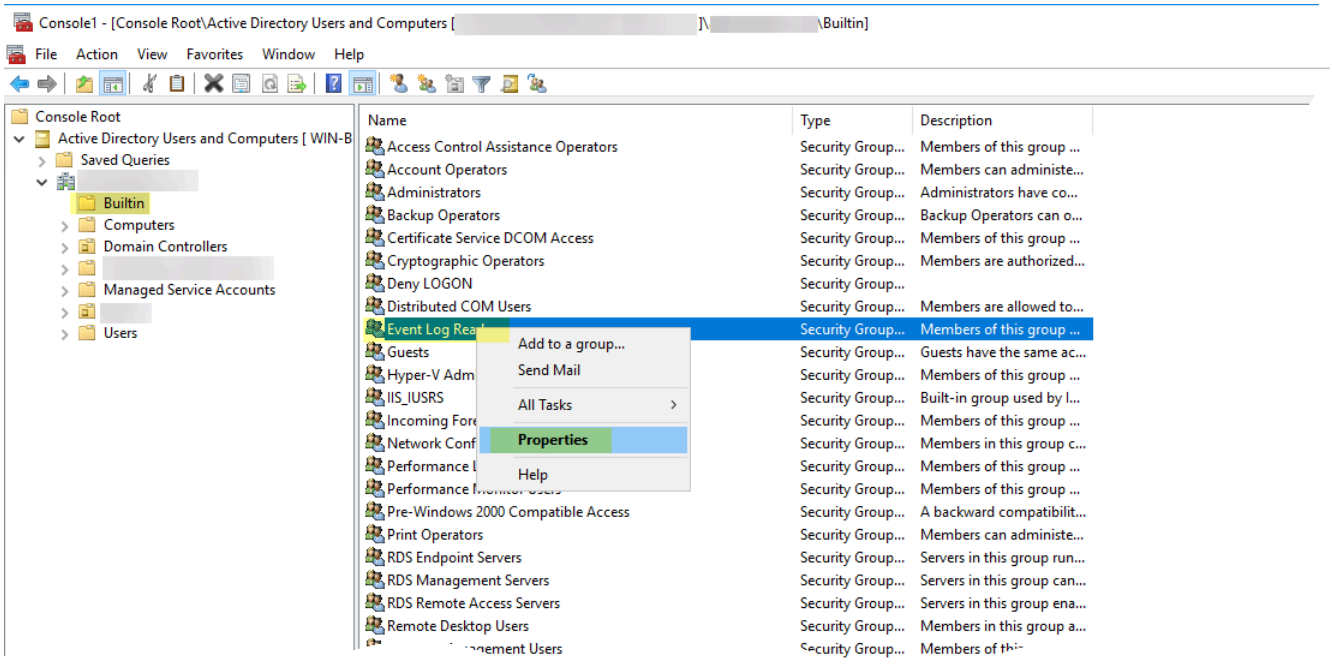
https://t.me/learningnets

STEP 2 | If you want to use Server Monitoring to identify users, add the service account to the Event Log Reader builtin group to allow the service account to read the security log events.

1. On the domain controller or Exchange server that contains the logs you want the User-ID agent to read, or on the member server that receives events from Windows log forwarding, select **Start** > **Run**, enter **MMC**.
2. Select **File** > **Add/Remove Snap-in** > **Active Directory Users and Computers** > **Add**, then click **OK** to run the MMC and launch the Active Directory Users and Computers snap-in.

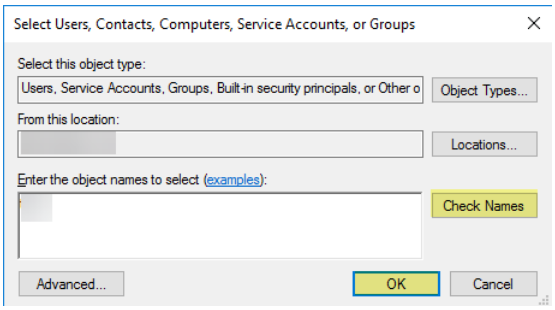


3. Navigate to the Builtin folder for the domain, right-click the **Event Log Readers** group, and select **Properties** > **Members**.



4. **Add** the service account then click **Check Names** to validate that you have the proper object name.

https://t.me/learningnets



5. Click **OK** twice to save the settings.
6. Confirm that the builtin Event Log Reader group lists the service account as a member (**Event Log Readers** > **Properties** > **Members**).

STEP 3 | If you want to use WMI to collect user data, assign DCOM privileges to the service account so that it can use WMI queries on monitored servers.

1. Select **Active Directory Users and Computers** > **<your domain>** > **Builtin** > **Distributed COM Users**.
2. Right-click **Properties** > **Members** > **Add** and enter the service account name.

 ©2022 Palo Alto Networks, Inc.

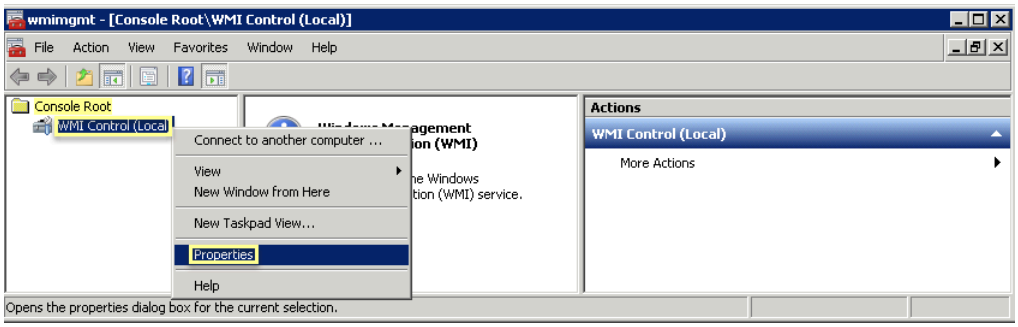
https://t.me/learningnets

STEP 4 | If you plan to use WMI probing, enable the service account to read the CIMV2 namespace on the domain controllers you want to monitor and assign the required permissions on the client systems to be probed.

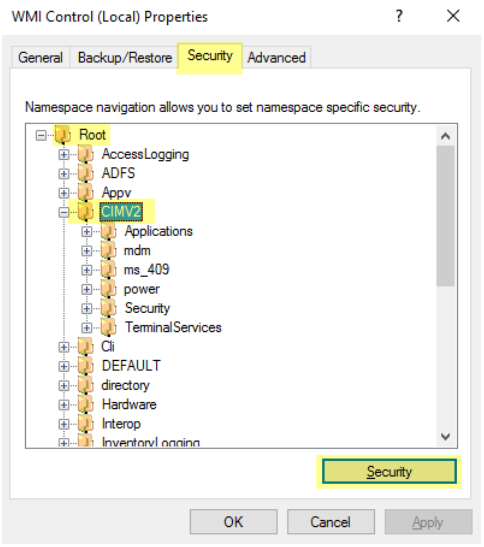
Do not enable client probing on high-security networks. Client probing can generate a large amount of network traffic and can pose a security threat when misconfigured. Instead collect user mapping information from more isolated and trusted sources, such as domain controllers and through integrations with Syslog or the XML API, which have the added benefit of allowing you to safely capture user mapping information from any device type or operating system, instead of just Windows clients.

Perform this task on each client system that the User-ID agent will probe for user mapping information:

1. Right-click the Windows icon (), **Search** for **`wmimgmt.msc`**, and launch the WMI Management Console.
2. In the console tree, right-click **WMI Control** and select **Properties**.



3. Select the **Security** tab, then select **Root** > **CIMV2**, and click the **Security** button.

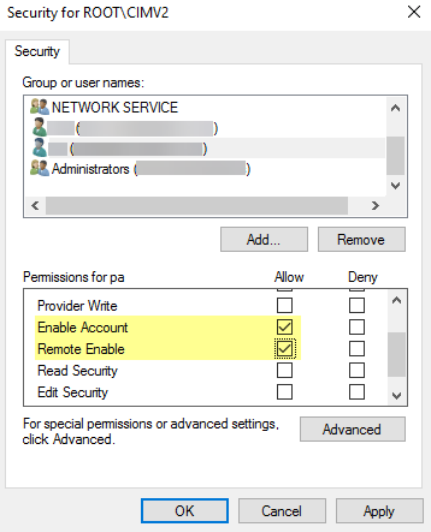


4. **Add** the name of the service account you created, **Check Names** to verify your entry, and click **OK**.

*You might have to change the **Locations** or click **Advanced** to query for account names. See the dialog help for details.*

https://t.me/learningnets

5. In the Permissions for **<Username>** section, **Allow** the **Enable Account** and **Remote Enable** permissions.



6. Click **OK** twice.
7. Use the Local Users and Groups MMC snap-in (lusrmgr.msc) to add the service account to the local Distributed Component Object Model (DCOM) Users and Remote Desktop Users groups on the system that will be probed.

STEP 5 | (Not Recommended) To allow the agent to monitor user sessions to poll Windows servers for user mapping information, assign Server Operator privileges to the service account.

Because this group also has privileges for shutting down and restarting servers, only assign the account to this group if monitoring user sessions is very important.

1. Select **Active Directory Users and Computers** > **<your domain>** > **Builtin** > **Server Operators Group**.
2. Right-click **Properties** > **Members** > **Add** add service account name

STEP 6 | Disable service account privileges that are not required.

By ensuring that the User-ID service account has the minimum set of account privileges, you can reduce the attack surface should the account be compromised.

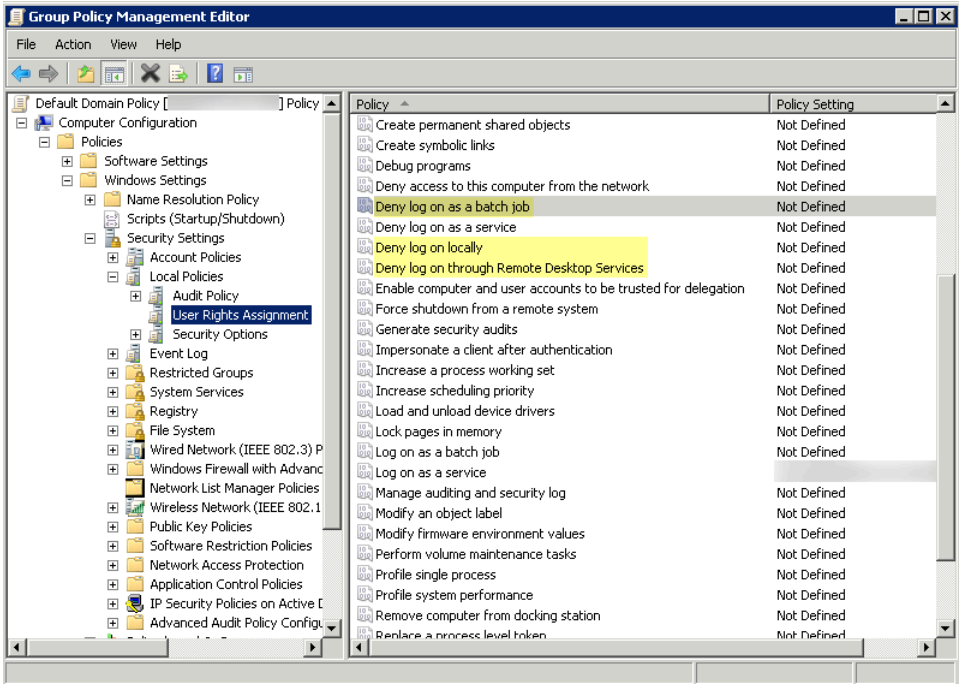
To ensure that the User-ID account has the minimum privileges necessary, deny the following privileges on the account:

- **Deny interactive logon for the User-ID service account**—While the User-ID service account does need permission to read and parse Active Directory security event logs, it does not require the ability to logon to servers or domain systems interactively. You can restrict this

https://t.me/learningnets

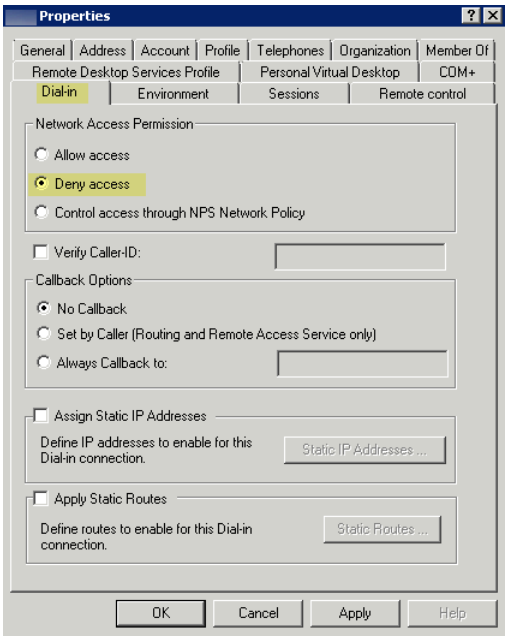
privilege using Group Policies or by using a Managed Service account (refer to Microsoft TechNet for more information).

1. Select **Group Policy Management Editor** > **Default Domain Policy** > **Computer Configuration** > **Policies** > **Windows Settings** > **Security Settings** > **User Rights Assignment**.
2. For **Deny log on as a batch job**, **Deny log on locally**, and **Deny log on through Remote Desktop Services**, right-click **Properties**, then select **Define these policy settings** > **Add User or Group** and add the service account name, then click **OK**.



- **Deny remote access for the User-ID service account**—This prevents an attacker from using the account to access your network from the outside the network.
  1. **Start** > **Run**, enter **MMC**, and select **File** > **Add/Remove Snap-in** > **Active Directory Users and Computers** > **Users**.
  2. Right-click the service account name, then select **Properties**.
  3. Select **Dial-in**, then **Deny** the **Network Access Permission**.

https://t.me/learningnets



STEP 7 | As a next step, Configure User Mapping Using the PAN-OS Integrated User-ID Agent.

Configure User Mapping Using the Windows User-ID Agent

In most cases, the majority of your network users will have logins to your monitored domain services. For these users, the Palo Alto Networks User-ID agent monitors the servers for login events and performs the IP address to username mapping. The way you configure the User-ID agent depends on the size of your environment and the location of your domain servers. As a best practice, locate your User-ID agents near the servers it will monitor (that is, the monitored servers and the Windows User-ID agent should not be across a WAN link from each other). This is because most of the traffic for user mapping occurs between the agent and the monitored server, with only a small amount of traffic—the delta of user mappings since the last update—from the agent to the firewall.

The following topics describe how to install and configure the User-ID Agent and how to configure the firewall to retrieve user mapping information from the agent:

- Install the Windows-Based User-ID Agent
- Configure the Windows User-ID Agent for User Mapping

Install the Windows-Based User-ID Agent

The following procedure shows how to install the User-ID agent on a member server in the domain and set up the service account with the required permissions. If you are upgrading, the installer will automatically remove the older version; however, it is a good idea to back up the config.xml file before running the installer.

For information about the system requirements for installing the Windows-based User-ID agent and for information on supported server OS versions, refer to the User-ID agent release notes *and the* Palo Alto Networks Compatibility Matrix.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 1 | Create a dedicated Active Directory service account for the User-ID agent to access the services and hosts it will monitor to collect user mappings.

Create a Dedicated Service Account for the User-ID Agent and grant the necessary permissions for the Windows User-ID agent.

1. Enable the service account to log on as a service by configuring either local or group policy.
   1. To configure the group policy if you are installing Windows-based User-ID agents on multiple servers, select **Group Policy Management** > **Default Domain Policy** >

 ©2022 Palo Alto Networks, Inc.

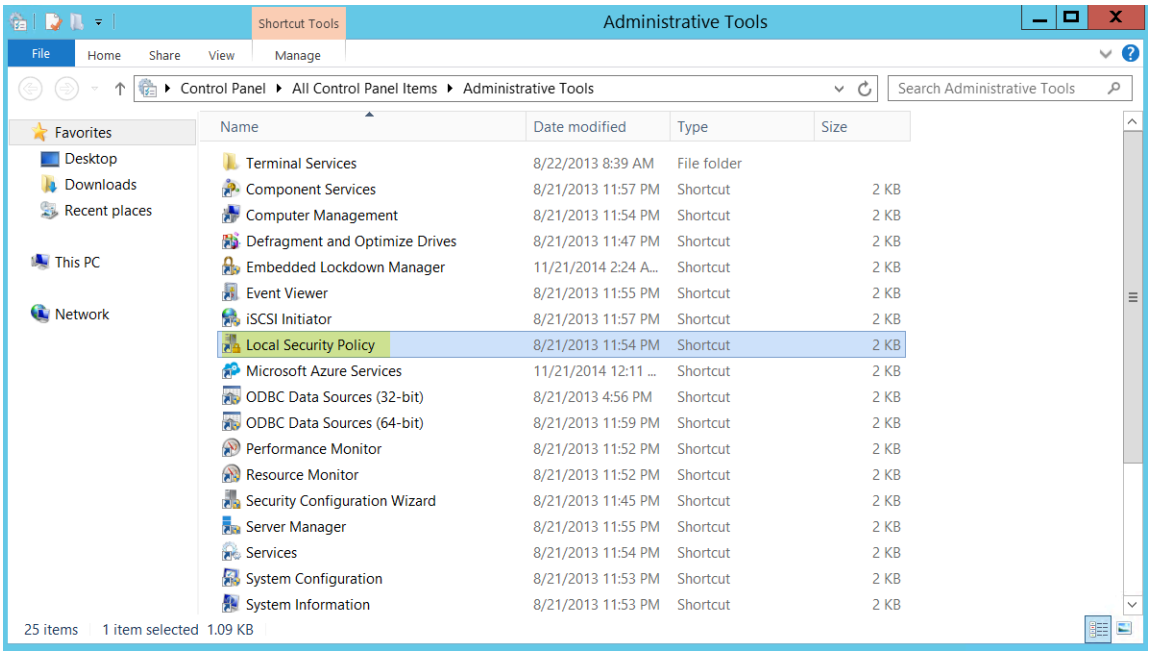
https://t.me/learningnets

Computer Configuration > **Policies** > **Windows Settings** > **Security Settings** > **Local Policies** > **User Rights Assignment** for the Windows server that is the agent host.

2. Right-click **Log on as a service**, then select **Properties**.
3. Add the service account username or builtin group (Administrators have this privilege by default).

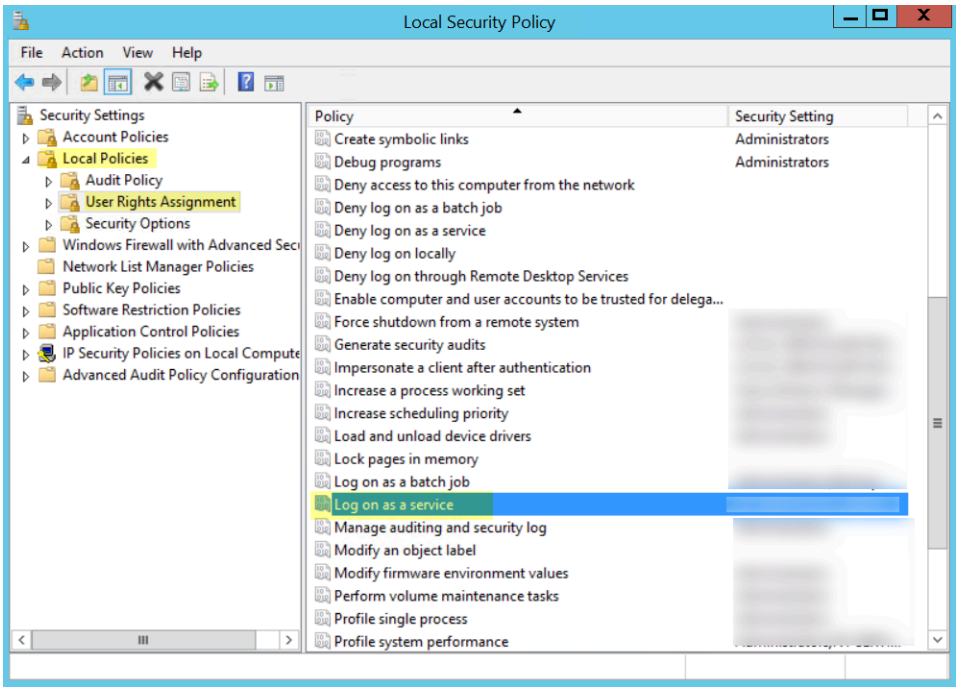
The permission to log on as a service is only needed locally on the Windows server that is the agent host. If you are using only one User-ID agent, you can grant the permissions locally on the agent host using the following instructions.

1. To assign permissions locally, select **Control Panel** > **Administrative Tools** > **Local Security Policy**.

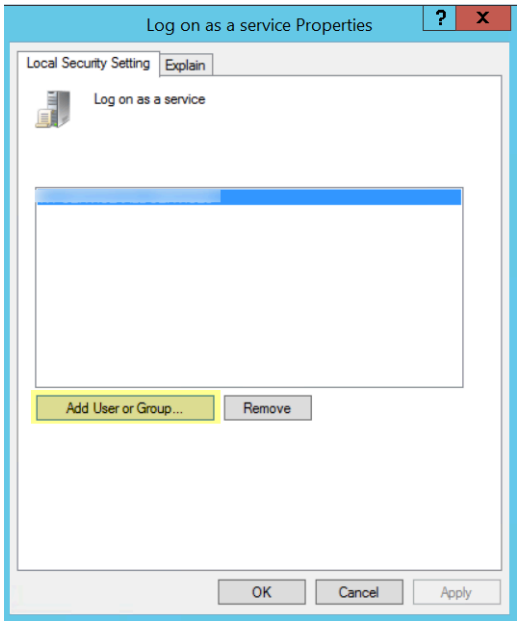


2. Select **Local Policies** > **User Rights Assignment** > **Log on as a service**.

https://t.me/learningnets



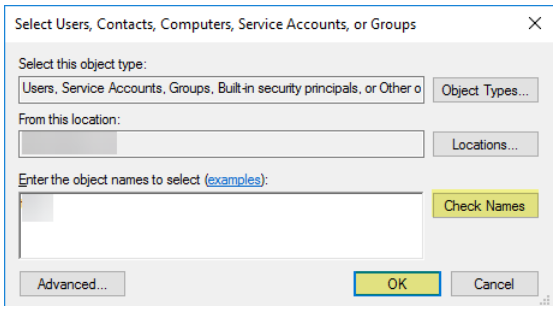
3. **Add User or Group** to add the service account.



4. Enter the service account name in **`domain\username`** format in the **Enter the object names to select** entry field and click **OK**.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets



To confirm the service account name is valid, **Check Names**.

2. If you want to use server monitoring to identify users, add the service account to the Event Log Reader builtin group to enable privileges for reading the security log events.

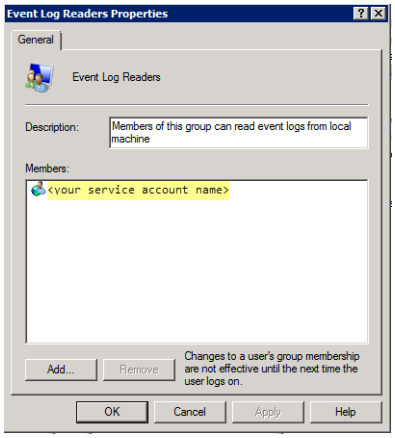
   1. On the domain controller or Exchange server that contains the logs you want the User-ID agent to read, or on the member server that receives events from Windows

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

log forwarding, run the MMC and launch the Active Directory Users and Computers snap-in.

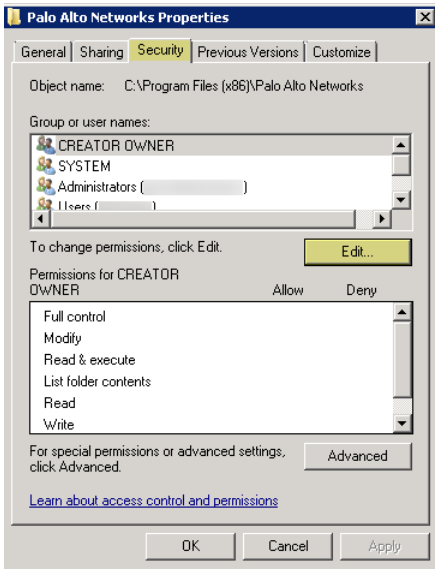
2. Navigate to the Builtin folder for the domain, right-click the **Event Log Reader** group and select **Add to Group** to open the properties dialog.
3. Click **Add** and enter the name of the service account that you configured the User-ID service to use and then click **Check Names** to validate that you have the proper object name.
4. Click **OK** twice to save the settings.
5. Confirm that the builtin Event Log Reader group lists the service account as a member.



3. Assign account permissions to the installation folder to allow the service account to access the agent's installation folder to read the configuration and write logs.

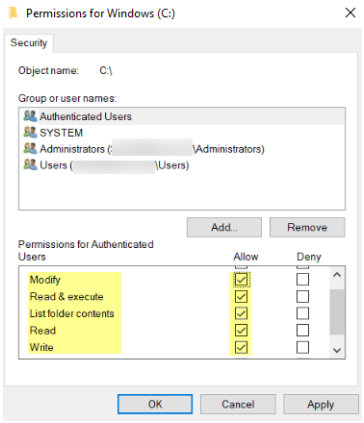
   You only need to perform this step if the service account you configured for the User-ID agent is not either a domain administrator or a local administrator on the User-ID agent server host.

   1. From the Windows Explorer, navigate to **`C:\Program Files(x86)\Palo Alto Networks`** for 32-bit systems, right-click the folder, and select **Properties**.
   2. On the **Security** tab, click **Edit**.



https://t.me/learningnets

3. **Add** the User-ID agent service account and assign it permissions to **Modify**, **Read & execute**, **List folder contents**, **Read**, and **Write**, and then click **OK** to save the account settings.



*If you want to allow the service account to access the User-ID agent's registry keys, **Allow** the **Full Control** permission.*

4. Give the service account permissions to the User-ID Agent registry sub-tree:
   1. Run **`regedt32`** and navigate to the Palo Alto Networks sub-tree in the following location: `HKEY_LOCAL_MACHINE\Software\Palo Alto Networks`.
   2. Right-click the Palo Alto Networks node and select **Permissions**.
   3. Assign the User-ID service account **Full Control** and then click **OK** to save the setting.

STEP 2 | Decide where to install the User-ID agent.

The User-ID agent queries the Domain Controller and Exchange server logs using Microsoft Remote Procedure Calls (MSRPCs). During the initial connection, the agent transfers the most recent 50,000 events from the log to map users. On each subsequent connection, the agent transfers events with a timestamp later than the last communication with the domain controller. Therefore, always install one or more User-ID agents at each site that has servers to be monitored.

- You must install the User-ID agent on a system running one of the supported OS versions: see "Operating System (OS) Compatibility User-ID Agent" in the Compatibility Matrix. The system must also meet the minimum requirements (see the User-ID agent release notes).
- Make sure the system that will host the User-ID agent is a member of the same domain as the servers it will monitor.
- As a best practice, install the User-ID agent close to the servers it will be monitoring: there is more traffic between the User-ID agent and the monitored servers than there is between the User-ID agent and the firewall, so locating the agent close to the monitored servers optimizes bandwidth usage.
- To ensure the most comprehensive mapping of users, you must monitor all domain controllers that process authentication for users you want to map. You might need to install multiple User-ID agents to efficiently monitor all of your resources.
- If you are using the User-ID agent for credential detection, you must install it on the read-only domain controller (RODC). As a best practice deploy a separate agent for this purpose.

https://t.me/learningnets

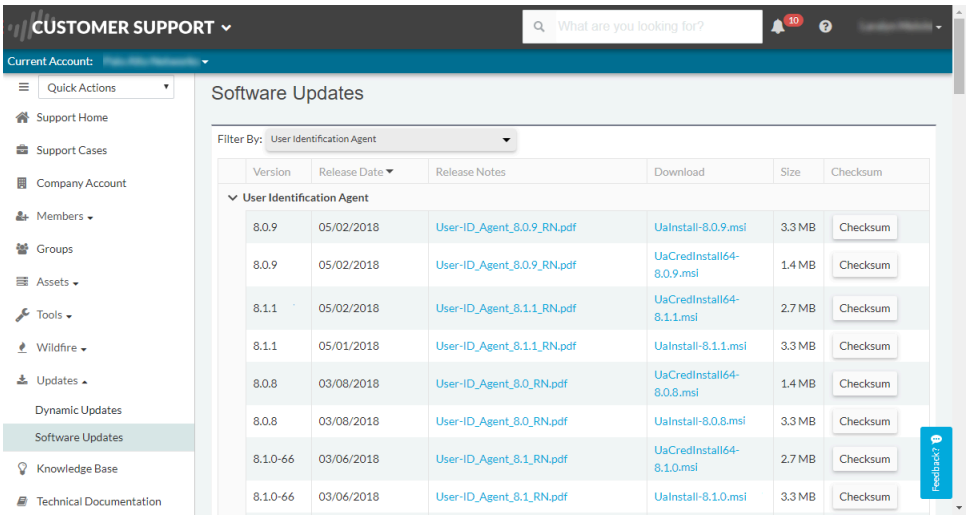
Do not use the User-ID agent installed on the RODC to map IP addresses to users. The User-ID agent installer for credential detection is named UaCredInstall64-x.x.x.msi.

STEP 3 | Download the User-ID agent installer.

Install the User-ID agent version that is the same as the PAN-OS version running on the firewalls. If there is not a User-ID agent version that matches the PAN-OS version, install the latest version that is closest to the PAN-OS version.

1. Log in to the Palo Alto Networks Customer Support Portal.
2. Select **Updates** > **Software Updates**.
3. Set **Filter By** to **User Identification Agent** and select the version of the User-ID agent you want to install from the corresponding Download column. The file name uses the following format: `UaInstall-x.x.x.msi` (where x represents the version number). For example, to download the 10.0 version of the User-ID agent, select **UaInstall-10.0.0-0.msi**.

   If you are using the User-ID agent for credential detection, download the `UaCredInstall64-x.x.x.msi` file instead. Only download and install the `UaCredInstall64-x.x.x.msi` if you are using the User-ID for credential detection.
4. Save the file on the systems where you plan to install the agent.



STEP 4 | Run the installer as an administrator.

1. Open the Windows **Start** menu, right-click the **Command Prompt** program, and select **Run as administrator**.
2. From the command line, run the .msi file you downloaded. For example, if you saved the .msi file to the Desktop, enter the following:

```
C:\Users\administrator.acme>cd Desktop
```


https://t.me/learningnets

```
C:\Users\administrator.acme\Desktop>UaInstall-6.0.0-1.msi
```

3. Follow the setup prompts to install the agent using the default settings. By default, the agent gets installed to **C:\Program Files(x86)\Palo Alto Networks**, but you can **Browse** to a different location.
4. When the installation completes, **Close** the setup window.

STEP 5 | Launch the User-ID Agent application as an administrator.

Open the Windows **Start** menu, right-click the **User-ID Agent** program, and select **Run as administrator**.

You must run the User-ID Agent application as an administrator to install the application, commit configuration changes, or uninstall the application.

STEP 6 | (Optional) Change the service account that the User-ID agent uses to log in.

By default, the agent uses the administrator account used to install the .msi file. To change the account to a restricted account:

1. Select **User Identification** > **Setup** and click **Edit**.
2. Select the **Authentication** tab and enter the service account name that you want the User-ID agent to use in the **User name for Active Directory** field.
3. Enter the **Password** for the specified account.
4. **Commit** the changes to the User-ID agent configuration to restart the service using the service account credentials.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 7 | (Optional) Assign your own certificates for mutual authentication between the Windows User-ID agent and the firewall.

1. Obtain your certificate for the Windows User-ID agent using one of the following methods. Upload the server certificate in Privacy Enhanced Mail (PEM) format and the server certificate's encrypted key.
   - Generate a Certificate and export it for upload to the Windows User-ID agent.
   - Export a certificate from your enterprise certificate authority (CA) and the upload it to the Windows User-ID agent.
2. Add a server certificate to Windows User-ID agent.
   1. On the Windows User-ID agent, select **Server Certificate** and click **Add**.
   2. Enter the path and name of the certificate file received from the CA or browse to the certificate file.
   3. Enter the private key passphrase.
   4. Click **OK** and then **Commit**.
3. Upload a certificate to the firewall to validate the Windows User-ID agent's identity.
4. Configure the certificate profile for the client device (firewall or Panorama).
   1. Select **Device** > **Certificate Management** > **Certificate Profile**.
   2. Configure a Certificate Profile.

   You can only assign one certificate profile for Windows User-ID agents and Terminal Server (TS) agents. Therefore, your certificate profile must include all certificate authorities that issued certificates uploaded to connected User-ID and TS agents.
5. Assign the certificate profile on the firewall.
   1. Select **Device** > **User Identification** > **Connection Security** and click the edit button.
   2. Select the **User-ID Certificate Profile** you configured in the previous step.
   3. Click **OK**.
6. **Commit** your changes.

STEP 8 | Configure Credential Detection with the Windows-based User-ID Agent.

To use the Windows-based User-ID agent to detect credential submissions and Prevent Credential Phishing, you must install the User-ID credential service on the Windows-based User-ID agent. You can only install this add-on on a read-only domain controller (RODC).

Configure the Windows User-ID Agent for User Mapping

The Palo Alto Networks Windows User-ID agent is a Windows service that connects to servers on your network—for example, Active Directory servers, Microsoft Exchange servers, and Novell eDirectory servers—and monitors the logs for login events. The agent uses this information to map IP addresses to usernames. Palo Alto Networks firewalls connect to the User-ID agent to retrieve this user mapping information, enabling visibility into user activity by username rather than IP address and enables user- and group-based security enforcement.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

For information about the server OS versions supported by the User-ID agent, refer to "Operating System (OS) Compatibility User-ID Agent" in the User-ID Agent Release Notes.

STEP 1 | Define the servers the User-ID agent will monitor to collect IP address to user mapping information.

The User-ID agent can monitor up to 100 servers, of which up to 50 can be syslog senders.

To collect all of the required mappings, the User-ID agent must connect to all servers that your users log in to in order to monitor the security log files on all servers that contain login events.

1. Open the Windows **Start** menu and select **User-ID Agent**.
2. Select **User Identification** > **Discovery**.
3. In the **Servers** section of the screen, click **Add**.
4. Enter a **Name** and **Server Address** for the server to be monitored. The network address can be a FQDN or an IP address.
5. Select the **Server Type** (**Microsoft Active Directory**, **Microsoft Exchange**, **Novell eDirectory**, or **Syslog Sender**) and then click **OK** to save the server entry. Repeat this step for each server to be monitored.
6. (Optional) To enable the Windows User-ID agent to automatically discover domain controllers on your network using DNS lookups, click **Auto Discover**. If you have new domain controllers that you want the Windows User-ID agent to discover, click **Auto Discover** each time you want to discover the new domain controllers.

   Auto-discovery locates domain controllers in the local domain only; you must manually add Exchange servers, eDirectory servers, and syslog senders.
7. (Optional) To tune the frequency at which the firewall polls configured servers for mapping information, select **User Identification** > **Setup** and **Edit** the Setup section. On the **Server Monitor** tab, modify the value in the **Server Log Monitor Frequency (seconds)** field. Increase the value in this field to 5 seconds in environments with older Domain Controllers or high-latency links.

   *Ensure that the **Enable Server Session Read** setting is not selected. This setting requires that the User-ID agent have an Active Directory account with Server Operator privileges so that it can read all user sessions. Instead, use a syslog or XML API integration to monitor sources that capture login and logout events for all device types and operating systems (instead of just Windows), such as wireless controllers and Network Access Controllers (NACs).*
8. Click **OK** to save the settings.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 2 | Specify the subnetworks the Windows User-ID agent should include in or exclude from User-ID.

By default, the User-ID maps all users accessing the servers you are monitoring.

As a best practice, always specify which networks to include and exclude from User-ID to ensure that the agent is only communicating with internal resources and to prevent unauthorized users from being mapped. You should only enable User-ID on the subnetworks where users internal to your organization are logging in.

1. Select **User Identification** > **Discovery**.
2. **Add** an entry to the Include/Exclude list of configured networks and enter a **Name** for the entry and enter the IP address range of the subnetwork in as the **Network Address**.
3. Select whether to include or exclude the network:
   - **Include specified network**—Select this option if you want to limit user mapping to users logged in to the specified subnetwork only. For example, if you include 10.0.0.0/8, the agent maps the users on that subnetwork and excludes all others. If you want the agent to map users in other subnetworks, you must repeat these steps to add additional networks to the list.
   - **Exclude specified network**—Select this option only if you want the agent to exclude a subset of the subnetworks you added for inclusion. For example, if you include 10.0.0.0/8 and exclude 10.2.50.0/22, the agent will map users on all the subnetworks of 10.0.0.0/8 except 10.2.50.0/22, and will exclude all subnetworks outside of 10.0.0.0/8.

     If you add Exclude profiles without adding any Include profiles, the User-ID agent excludes all subnetworks, not just the ones you added.
4. Click **OK**.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 3 | (Optional) If you configured the agent to connect to a Novell eDirectory server, you must specify how the agent should search the directory.

1. Select **User Identification** > **Setup** and click **Edit** in the Setup section of the window.
2. Select the **eDirectory** tab and then complete the following fields:
   - **Search Base**—The starting point or root context for agent queries, for example: `dc=domain1,dc=example, dc=com`.
   - **Bind Distinguished Name**—The account to use to bind to the directory, for example: `cn=admin,ou=IT, dc=domain1, dc=example, dc=com`.
   - **Bind Password**—The bind account password. The agent saves the encrypted password in the configuration file.
   - **Search Filter**—The search query for user entries (default is `objectClass=Person`).
   - **Server Domain Prefix**—A prefix to uniquely identify the user. This is only required if there are overlapping name spaces, such as different users with the same name from two different directories.
   - **Use SSL**—Select the check box to use SSL for eDirectory binding.
   - **Verify Server Certificate**—Select the check box to verify the eDirectory server certificate when using SSL.

STEP 4 | (Strongly recommended) Disable client probing.

Palo Alto Networks strongly recommends disabling client probing on high-security networks. Client probing can pose a security threat if not correctly configured. For more information, see client probing.

1. On the **Client Probing** tab, deselect the **Enable WMI Probing** check box if it is enabled.
2. Deselect the **Enable NetBIOS Probing** check box if it is enabled.

Palo Alto Network strongly recommends that you collect user mapping information from isolated and trusted sources, such as domain controllers or integrations with Syslog or the XML API, to safely capture user mapping information from any device type or operating system.

*If you must enable client probing, select the **Enable WMI Probing** check box and on the **Client Probing** tab. Due to the potential security risks of this method, only select the **Enable NetBIOS Probing** check box if the firewall cannot obtain user mappings using any other method. Then add a remote administration exception to the Windows firewall for each probed client to ensure the Windows firewall will allow client probing. Each probed client PC must allow port 139 in the Windows firewall and must also have file and printer sharing services enabled.*

STEP 5 | Save the configuration.

Click **OK** to save the User-ID agent setup settings and then click **Commit** to restart the User-ID agent and load the new settings.

https://t.me/learningnets

STEP 6 | (Optional) Define the set of users for which you do not need to provide IP address-to-username mappings, such as kiosk accounts.

Save the `ignore-user` list as a text document on the agent host using the title `ignore_user_list` and use the .txt file extension to save it to the User-ID Agent folder on the domain server where the agent is installed.

List the user accounts to ignore; there is no limit to the number of accounts you can add to the list. Each user account name must be on a separate line. For example:

```
SPAdmin
SPInstall
TFSReport
```

You can use an asterisk as a wildcard character to match multiple usernames, but only as the last character in the entry. For example, `corpdomain\it-admin*` would match all administrators in the `corpdomain` domain whose usernames start with the string `it-admin`. You can also use the `ignore-user` list to identify users whom you want to force to authenticate using Authentication Portal.

After adding entries to the Ignore User list, you must stop and restart the connection to the service.

STEP 7 | Configure the firewall to connect to the User-ID agent.

The firewall can connect to only one Windows-based User-ID agent that is using the User-ID credential service add-on to detect corporate credential submissions. See Configure Credential Detection with the Windows-based User-ID Agent *for more details on how to use this service for credential phishing prevention.*

Complete the following steps on each firewall you want to connect to the User-ID agent to receive user mappings:

1. Select **Device** > **Data Redistribution** > **Agents** and click **Add**.
2. Enter a **Name** for the agent.
3. **Add an Agent Using** the **Host and Port**.
4. Enter the IP address of the Windows **Host** on which the User-ID Agent is installed.
5. Enter the **Port** number (1-65535) on which the agent will listen for user mapping requests. This value must match the value configured on the User-ID agent. By default, the port is set to 5007 on the firewall and on newer versions of the User-ID agent. However, some older User-ID agent versions use port 2010 as the default.
6. Select **IP User Mappings** as the **Data type**.
7. Make sure that the configuration is **Enabled**, then click **OK**.
8. **Commit** the changes.
9. Verify that the **Connected status** displays as connected (a green light).

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 8 | Verify that the User-ID agent is successfully mapping IP addresses to usernames and that the firewalls can connect to the agent.

1. Launch the User-ID agent and select **User Identification**.
2. Verify that the agent status shows **Agent is running**. If the Agent is not running, click **Start**.
3. To verify that the User-ID agent can connect to monitored servers, make sure the Status for each Server is **Connected**.
4. To verify that the firewalls can connect to the User-ID agent, make sure the Status for each of the Connected Devices is **Connected**.
5. To verify that the User-ID agent is mapping IP addresses to usernames, select **Monitoring** and make sure that the mapping table is populated. You can also **Search** for specific users, or **Delete** user mappings from the list.

Configure User Mapping Using the PAN-OS Integrated User-ID Agent

The following procedure describes how to configure the PAN-OS® integrated User-ID™ agent on the firewall for IP address-to-username mapping. The integrated User-ID agent performs the same tasks as the Windows-based agent with the exception of NetBIOS client probing (WMI probing is supported).

STEP 1 | Create an Active Directory service account for the User-ID agent to access the services and hosts that the firewall will monitor for collecting user mapping information.

Create a Dedicated Service Account for the User-ID Agent.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 2 | Define the servers that the firewall will monitor to collect user mapping information.

Within the total maximum of 100 monitored servers per firewall, you can define no more than 50 syslog senders for any single virtual system.

To collect all the required mappings, the firewall must connect to all servers that your users log in to so that the firewall can monitor the Security log files on all servers that contain login events.

1. Select **Device** > **User Identification** > **User Mapping**.
2. **Add** a server (Server Monitoring section).
3. Enter a **Name** to identify the server.
4. Select the **Type** of server.
   - **Microsoft Active Directory**
   - **Microsoft Exchange**
   - **Novell eDirectory**
   - **Syslog Sender**
5. (Microsoft Active Directory and Microsoft Exchange only) Select the **Transport Protocol** you want to use to monitor security logs and session information on the server.
   - **WMI**—The firewall and the monitored servers use Windows Management Instrumentation (WMI) to communicate.
   - **WinRM-HTTP**—The firewall and the monitored servers use Kerberos for mutual authentication and the monitored server encrypts the communication with the firewall using a negotiated Kerberos session key.
   - **WinRM-HTTPS**—The firewall and the monitored servers use HTTPS to communicate and use basic authentication or Kerberos for mutual authentication.

   If you select a Windows Remote Management (WinRM) option, you must Configure Server Monitoring Using WinRM.
6. (Microsoft Active Directory, Microsoft Exchange, and Novell eDirectory only) Enter the **Network Address** of the server.

   *If you are using WinRM with Kerberos, you must enter a fully qualified domain name (FDQN). If you want to use WinRM with basic authentication or use **WMI** to monitor the server, you can enter an IP address or FQDN.*

   *To monitor servers using WMI, specify an IP address, the service account name (if all server monitoring is in the same domain), or a fully qualified domain name (FQDN). If you specify an FQDN, use the down-level logon name in the (DLN)\sAMAccountName format instead of the FQDN\sAMAccountName format. For example, use **`example\user.services`** not **`example.com\user.services`**. If you specify an FQDN, the firewall will attempt to authenticate using Kerberos, which does not support WMI.*
7. (Syslog Sender only) If you select **Syslog Sender** as the server **Type**, Configure the PAN-OS Integrated User-ID Agent as a Syslog Listener.
8. (Novell eDirectory only) Make sure the **Server Profile** you select is **Enabled** and click **OK**.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

9. (Optional) Configure the firewall to automatically **Discover** domain controllers on your network using DNS lookups.

   The auto-discovery feature is for domain controllers only; you must manually add any Exchange servers or eDirectory servers you want to monitor.

STEP 3 | (Optional) Specify the frequency at which the firewall polls Windows servers for mapping information. This is the interval between the end of the last query and the start of the next query.

If the domain controller is processing many requests, delays between queries may exceed the specified value.

1. **Edit** the **Palo Alto Networks User ID Agent Setup**.
2. Select the **Server Monitor** tab and specify the **Server Log Monitor Frequency** in seconds (range is 1 to 3,600; default is 2). In environments with older domain controllers or high-latency links, set this frequency to a minimum of five seconds.

   *Ensure that the **Enable Session** option is not enabled. This option requires that the User-ID agent have an Active Directory account with Server Operator privileges so that it can read all user sessions. Instead, use a Syslog or XML API integration to monitor sources that capture login and logout events for all device types and operating systems (instead of just Windows), such as wireless controllers and network access control (NAC) devices.*
3. Click **OK** to save your changes.

STEP 4 | Specify the subnetworks that the PAN-OS integrated User-ID agent should include in or exclude from user mapping.

By default, the User-ID maps all users accessing the servers you are monitoring.

As a best practice, always specify which networks to include and, optionally, which networks to exclude from User-ID to ensure that the agent is communicating only with internal resources and to prevent unauthorized users from being mapped. You should enable user mapping only on the subnetworks where users internal to your organization are logging in.

1. Select **Device** > **User Identification** > **User Mapping**.
2. **Add** an entry to the **Include/Exclude Networks** and enter a **Name** for the entry. Ensure that the entry is **Enabled**.
3. Enter the **Network Address** and then select whether to include or exclude it:
   - **Include**—Select this option to limit user mapping to only users logged in to the specified subnetwork. For example, if you include 10.0.0.0/8, the agent maps the users on that subnetwork and excludes all others. If you want the agent to map users in other subnetworks, you must repeat these steps to add additional networks to the list.
   - **Exclude**—Select this option to configure the agent to exclude a subset of the subnetworks you added for inclusion. For example, if you include 10.0.0.0/8 and

https://t.me/learningnets

exclude 10.2.50.0/22, the agent will map users on all the subnetworks of 10.0.0.0/8 except 10.2.50.0/22 and will exclude all subnetworks outside of 10.0.0.0/8.

If you add Exclude profiles without adding any Include profiles, the User-ID agent excludes all subnetworks, not just the ones you added.

4. Click **OK**.

STEP 5 | Set the domain credentials for the account that the firewall will use to access Windows resources. This is required for monitoring Exchange servers and domain controllers as well as for WMI probing.

1. **Edit** the **Palo Alto Networks User-ID Agent Setup**.
2. Select the **Server Monitor Account** tab and enter the **User Name** and **Password** for the service account that the User-ID agent will use to probe the clients and monitor servers. Enter the username using the **`domain\username`** syntax.
3. If you are using WinRM to monitor servers, configure the firewall to authenticate with the server you are monitoring.
   - If you want to use WinRM with basic authentication, enable WinRM on the server, configure basic authentication, and specify the service account **Domain's DNS Name**.
   - If you want to use WinRM with Kerberos, Configure a Kerberos server profile if you have not already done so and then select the **Kerberos Server Profile**.

STEP 6 | (Optional, not recommended) Configure WMI probing (the PAN-OS integrated User-ID agent does not support NetBIOS probing).

Do not enable WMI probing on high-security networks. Client probing can generate a large amount of network traffic and can pose a security threat when misconfigured.

1. On the **Client Probing** tab, **Enable Probing**.
2. (Optional) Specify the **Probe Interval** to define the interval (in minutes) between the end of the last probe request and the start of the next request.

   If necessary, increase the value to ensure the User-ID agent has sufficient time to probe all the learned IP addresses (range is 1 to 1440; default is 20).

   If the request load is high, the observed delay between requests might significantly exceed the specified interval.
3. Click **OK**.
4. Make sure the Windows firewall will allow client probing by adding a remote administration exception to the Windows firewall for each probed client.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 7 | (Optional) Define the set of user accounts that don't require IP address-to-username mappings, such as kiosk accounts.

> *Define the ignore user list on the firewall that is the User-ID agent, not the client. If you define the ignore user list on the client firewall, the users in the list are still mapped during redistribution.*

On the **Ignore User List** tab, **Add** each username you want to exclude from user mapping. You can also use the ignore user list to identify the users you want to force to use Authentication Portal to authenticate. You can use an asterisk as a wildcard character to match multiple usernames but only as the last character in the entry. For example, **`corpdomain\it-admin*`** would match all administrators in the `corpdomain` domain whose usernames start with the string `it-admin`. You can add up to 5,000 entries to exclude from user mapping.

STEP 8 | Activate your configuration changes.

Click **OK** and **Commit**.

STEP 9 | Verify the configuration.

1. Access the firewall CLI.
2. Enter the following operational command:

```
> show user server-monitor state all
```

3. On the **Device** > **User Identification** > **User Mapping** tab in the web interface, verify that the Status of each server you configured for server monitoring is `Connected`.

Configure Server Monitoring Using WinRM

You can configure the PAN-OS integrated User-ID agent to monitor servers using Windows Remote Management (WinRM). Using the WinRM protocol improves speed, efficiency, and security when monitoring server events to map user events to IP addresses. The PAN-OS integrated User-ID agent supports the WinRM protocol on Windows Server 2012 Active Directory and Microsoft Exchange Server 2012 or later versions of both.

There are three ways to configure server monitoring using WinRM:

- Configure WinRM over HTTPS with Basic Authentication—The firewall authenticates to the monitored server using the username and password of the service account for the User-ID agent and the firewall authenticates the monitored server using the User-ID certificate profile.
- Configure WinRM over HTTP with Kerberos—The firewall and the monitored servers use Kerberos for mutual authentication and the monitored server encrypts the communication with the firewall using a negotiated Kerberos session key.
- Configure WinRM over HTTPS with Kerberos—The firewall and the monitored server use HTTPS to communicate and use Kerberos for mutual authentication.

Configure WinRM over HTTPS with Basic Authentication

When you configure WinRM to use HTTPS with basic authentication, the firewall transfers the credentials for the service account in a secure tunnel using SSL.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 1 | Configure the service account with Remote Management User and CIMV2 privileges for the server you want to monitor.

STEP 2 | On the Windows server you are monitoring, obtain the thumbprint from the certificate for the Windows server to use with WinRM and enable WinRM.

Ensure that you use an account with administrator privileges to configure WinRM on the server you want to monitor. As a best practice for security, this account should not be the same account as the service account in Step 1.

1. Verify the certificate is installed in the Local Computer certificate store (**Certificates (Local Computer)** > **Personal** > **Certificates**).

   If you do not see the Local Computer certificate store, launch the Microsoft Management Console (**Start** > **Run** > **MMC**) and add the Certificates snap-in (**File** > **Add/Remove Snap-in** > **Certificates** > **Add** > **Computer account** > **Next** > **Finish**).
2. Open the certificate and select **General** > **Details** > **Show: <All>**.
3. Select the **Thumbprint** and copy it.
4. To enable the firewall to connect to the Windows server using WinRM, enter the following command: **`winrm quickconfig`**.
5. Enter **`y`** to confirm the changes and then confirm the output displays `WinRM service started`.

   If WinRM is enabled, the output displays `WinRM service is already running on this machine.` You will be prompted to confirm any additional required configuration changes.
6. To verify that WinRM is communicating using HTTPS, enter the following command: **`winrm enumerate winrm/config/listener`** and confirm that the output displays `Transport = HTTPS`.

   By default, WinRM/HTTPS uses port 5986.
7. From the Windows server command prompt, enter the following command: **`winrm create winrm/config/Listener?Address=*+Transport=HTTPS @{Hostname="<hostname>";CertificateThumbprint="Certificate Thumbprint"}`**, where *hostname* is the hostname of the Windows server and *Certificate Thumbprint* is the value you copied from the certificate.

   Use the command prompt (not Powershell) and remove any spaces in the Certificate Thumbprint to ensure that WinRM can validate the certificate.
8. From the Windows server command prompt, enter the following command:

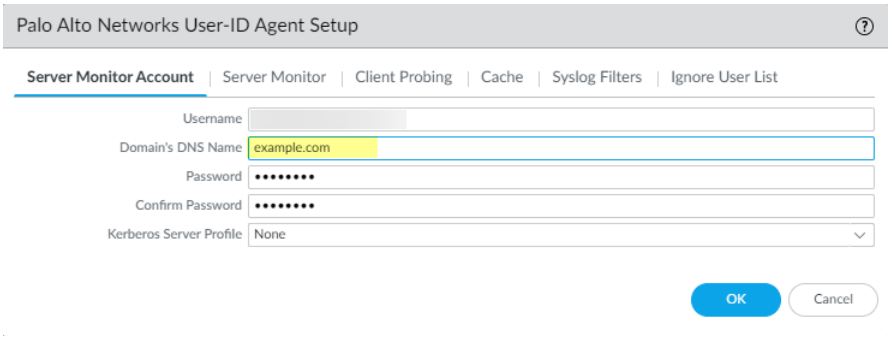
   ```
   c:\> winrm set winrm/config/client/auth @{Basic="true"}
   ```
9. Enter the following command: **`winrm get winrm/config/service/Auth`** and confirm that `Basic = true`.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 3 | Enable Basic Authentication between the PAN-OS integrated User-ID agent and the monitored servers.

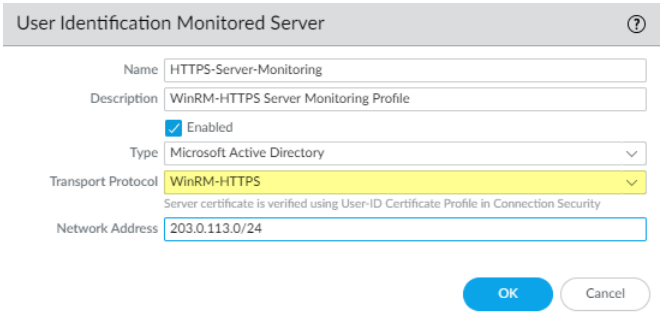
1. Select **Device** > **User Identification** > **User Mapping** > **Palo Alto Networks User-ID Agent Setup** > **Server Monitor Account**.
2. In **`domain\username`** format, enter the **User Name** for the service account that the User-ID agent will use to monitor servers.
3. Enter the **Domain's DNS Name** of the server monitor account.



4. Enter the **Password** and **Confirm Password** for the service account.
5. Click **OK**

STEP 4 | Configure server monitoring for the PAN-OS integrated User-ID agent.

1. Select the Microsoft server **Type** (**Microsoft Active Directory** or **Microsoft Exchange**).
2. Select **Win-RM-HTTPS** as the **Transport Protocol** to use Windows Remote Management (WinRM) over HTTPS to monitor the server security logs and session information.



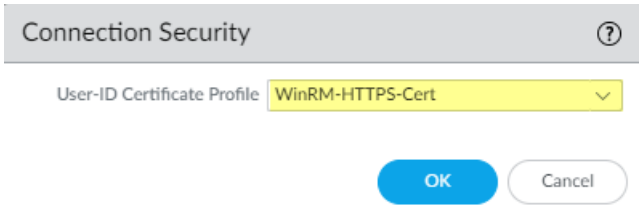
3. Enter the IP address or FQDN **Network Address** of the server.

STEP 5 | To enable the PAN-OS integrated User-ID agent to communicate with the monitored servers using WinRM-HTTPS, verify that you successfully imported the root certificate for the

https://t.me/learningnets

service certificates that the Windows server uses for WinRM on to the firewall and associate the certificate with the User-ID Certificate Profile.

1. Select **Device** > **User Identification** > **Connection Security**.
2. Click **Edit**.
3. Select the Windows server certificate to use for the **User-ID Certificate Profile**.



4. Click **OK**.

STEP 6 | **Commit** your changes.

STEP 7 | Verify that the status of each monitored server is Connected (**Device** > **User Identification** > **User Mapping**).

Configure WinRM over HTTP with Kerberos

When you configure WinRM over HTTP with Kerberos, the firewall and the monitored servers use Kerberos for mutual authentication and the monitored server encrypts the communication with the firewall using a negotiated Kerberos session key.

WinRM with Kerberos supports the aes128-cts-hmac-sha1-96 and aes256-cts-hmac-sha1-96 ciphers. If the server you want to monitor uses RC4, you must download the Windows update and disable RC4 for Kerberos in the registry settings of the server you want to monitor.

STEP 1 | Configure the service account with Remote Management User and CIMV2 privileges for the server you want to monitor.

https://t.me/learningnets

STEP 2 | Confirm that WinRM is enabled on the Windows server you are monitoring.

Ensure that you use an account with administrator privileges to configure WinRM on the server you want to monitor. As a best practice for security, this account should not be the same account as the service account in Step 1.

1. To enable the firewall to connect to the Windows server using WinRM, enter the following command: **`winrm quickconfig`**.
2. Enter **`y`** to confirm the changes and then confirm the output displays `WinRM service started`.

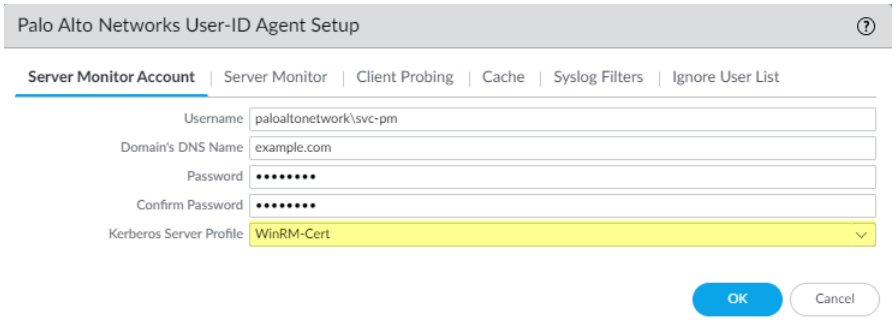
   If WinRM is enabled, the output displays `WinRM service is already running on this machine.` You will be prompted to confirm any additional required configuration changes.
3. To verify that WinRM is communicating using HTTP, enter the following command: **`winrm enumerate winrm/config/listener`** and confirm that the output displays `Transport = HTTP`.

   By default, WinRM/HTTP uses port 5985.
4. Enter the following command: **`winrm get winrm/config/service/Auth`** and confirm that `Kerberos = true`.

STEP 3 | Enable the PAN-OS integrated User-ID agent and the monitored servers to authenticate using Kerberos.

1. If you did not do so during the initial configuration, configure date and time (NTP) settings to ensure successful Kerberos negotiation.
2. Configure a Kerberos server profile on the firewall to authenticate with the server to monitor the security logs and session information.
3. Select **Device** > **User Identification** > **User Mapping** > **Palo Alto Networks User-ID Agent Setup** > **Server Monitor Account**.
4. In **`domain\username`** format, enter the **User Name** for the service account that the User-ID agent will use to monitor servers.
5. Enter the **Domain's DNS Name** of the server monitor account.

   Kerberos uses the domain name to locate the service account.
6. Enter the **Password** and **Confirm Password** for the service account.
7. Select the **Kerberos Server Profile** you configured in Step 3.2.

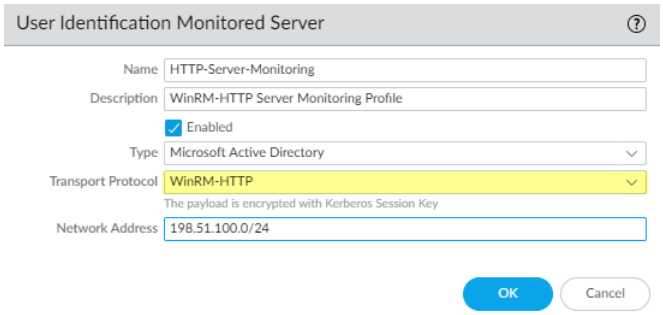


8. Click **OK**.

https://t.me/learningnets

STEP 4 | Configure server monitoring for the PAN-OS integrated User-ID agent.

1. Configure the Microsoft server type (**Microsoft Active Directory** or **Microsoft Exchange**).
2. Select **WinRM-HTTP** as the **Transport Protocol** to use Windows Remote Management (WinRM) over HTTP to monitor the server security logs and session information.



3. Enter the FQDN **Network Address** of the server.

   If you are using Kerberos, the network address must be a fully qualified domain name (FDQN).

STEP 5 | **Commit** your changes.

STEP 6 | Verify that the status of each monitored server is Connected (**Device** > **User Identification** > **User Mapping**).

Configure WinRM over HTTPS with Kerberos

When you configure WinRM over HTTPS with Kerberos, the firewall and the monitored server use HTTPS to communicate and use Kerberos for mutual authentication.

WinRM with Kerberos supports the aes128-cts-hmac-sha1-96 and aes256-cts-hmac-sha1-96 ciphers. If the server you want to monitor uses RC4, you must download the Windows update and disable RC4 for Kerberos in the registry settings of the server you want to monitor.

STEP 1 | Configure the service account with Remote Management User and CIMV2 privileges for the server you want to monitor.

https://t.me/learningnets

STEP 2 | On the Windows server you are monitoring, obtain the thumbprint from the certificate for the Windows server to use with WinRM and enable WinRM.

> *Ensure that you use an account with administrator privileges to configure WinRM on the server you want to monitor. As a best practice for security, this account should not be the same account as the service account in Step 1.*

1. Verify the certificate is installed in the Local Computer certificate store (**Certificates (Local Computer)** > **Personal** > **Certificates**).

   If you do not see the Local Computer certificate store, launch the Microsoft Management Console (**Start** > **Run** > **MMC**) and add the Certificates snap-in (**File** > **Add/Remove Snap-in** > **Certificates** > **Add** > **Computer account** > **Next** > **Finish**).
2. Open the certificate and select **General** > **Details** > **Show: <All>**.
3. Select the **Thumbprint** and copy it.
4. To enable the firewall to connect to the Windows server using WinRM, enter the following command: **`winrm quickconfig`**.
5. Enter **`y`** to confirm the changes and then confirm the output displays `WinRM service started`.

   If WinRM is enabled, the output displays `WinRM service is already running on this machine.` You will be prompted to confirm any additional required configuration changes.
6. To verify that WinRM is communicating using HTTPS, enter the following command: **`winrm enumerate winrm/config/listener`**. Then confirm that the output displays `Transport = HTTPS`.

   By default, WinRM/HTTPS uses 5986.
7. From the Windows server command prompt, enter the following command: **`winrm create winrm/config/Listener?Address=*+Transport=HTTPS @{Hostname="<hostname>";CertificateThumbprint="Certificate Thumbprint"}`**, where *hostname* is the hostname of the Windows server and *Certificate Thumbprint* is the value you copied from the certificate.

   > *Use the command prompt (not Powershell) and remove any spaces in the Certificate Thumbprint to ensure that WinRM can validate the certificate.*
8. Enter the following command: **`winrm get winrm/config/service/Auth`** and confirm that `Basic = false` and `Kerberos= true`.

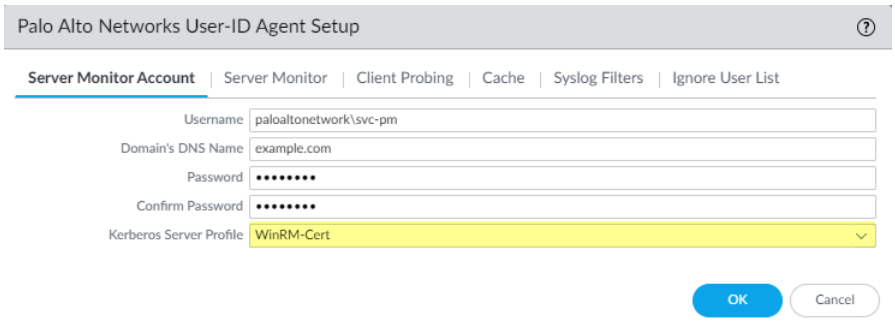
 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 3 | Enable the PAN-OS integrated User-ID agent and the monitored servers to authenticate using Kerberos.

1. If you did not do so during the initial configuration, configure date and time (NTP) settings to ensure successful Kerberos negotiation.
2. Configure a Kerberos server profile on the firewall to authenticate with the server to monitor the security logs and session information.
3. Select **Device** > **User Identification** > **User Mapping** > **Palo Alto Networks User-ID Agent Setup** > **Server Monitor Account**.
4. In **`domain\username`** format, enter the **User Name** for the service account that the User-ID agent will use to monitor servers.
5. Enter the **Domain's DNS Name** of the server monitor account.

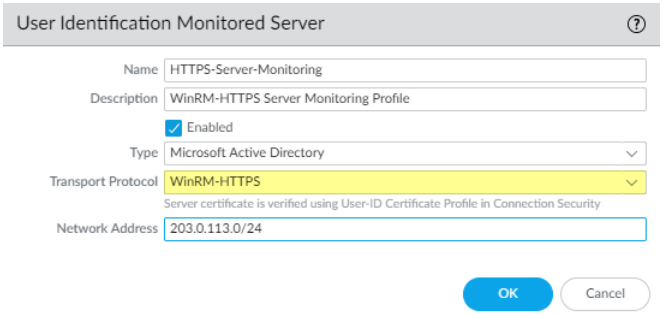
   Kerberos uses the domain name to locate the service account.
6. Enter the **Password** and **Confirm Password** for the service account.
7. Select the **Kerberos Server Profile** you created in Step 3.2.



8. Click **OK**.

STEP 4 | Configure server monitoring for the PAN-OS integrated User-ID agent.

1. Configure the Microsoft server type (**Microsoft Active Directory** or **Microsoft Exchange**).
2. Select **Win-RM-HTTPS** as the **Transport Protocol** to use Windows Remote Management (WinRM) over HTTPS to monitor the server security logs and session information.



3. Enter the FQDN **Network Address** of the server.

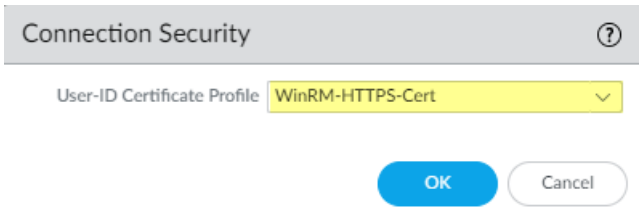
   If you are using Kerberos, the network address must be a fully qualified domain name (FDQN).

https://t.me/learningnets

STEP 5 | To enable the PAN-OS integrated User-ID agent to communicate with the monitored servers using WinRM-HTTPS, verify that you successfully imported the root certificate for the service certificates that the Windows server uses for WinRM on to the firewall and associate the certificate with the User-ID Certificate Profile.

The firewall uses the same certificate to authenticate with all monitored servers.

1. Select **Device** > **User Identification** > **Connection Security**.
2. Click **Edit**.
3. Select the Windows server certificate to use for the **User-ID Certificate Profile**.



4. Click **OK**.
5. **Commit** your changes.

STEP 6 | Verify that the status of each monitored server is Connected (**Device** > **User Identification** > **User Mapping**).

Configure User-ID to Monitor Syslog Senders for User Mapping

To obtain IP address-to-username mappings from existing network services that authenticate users, you can configure the PAN-OS integrated User-ID agent or Windows-based User-ID agent to parse Syslog messages from those services. To keep user mappings up to date, you can also configure the User-ID agent to parse syslog messages for logout events so that the firewall automatically deletes outdated mappings.

- Configure the PAN-OS Integrated User-ID Agent as a Syslog Listener
- Configure the Windows User-ID Agent as a Syslog Listener

Configure the PAN-OS Integrated User-ID Agent as a Syslog Listener

To configure the PAN-OS Integrated User-ID agent to create new user mappings and remove outdated mappings through syslog monitoring, start by defining Syslog Parse profiles. The User-ID agent uses the profiles to find login and logout events in syslog messages. In environments where *syslog senders* (the network services that authenticate users) deliver syslog messages in different formats, configure a profile for each syslog format. Syslog messages must meet certain criteria for a User-ID agent to parse them (see Syslog). This procedure uses examples with the following formats:

- **Login events**—`[Tue Jul 5 13:15:04 2016 CDT] Administratorauthentication success User:johndoe1 Source:192.168.3.212`
- **Logout events**—`[Tue Jul 5 13:18:05 2016CDT] User logout successful User:johndoe1 Source:192.168.3.212`

After configuring the Syslog Parse profiles, you specify syslog senders for the User-ID agent to monitor.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 1 | Determine whether there is a predefined Syslog Parse profile for your particular syslog senders.

Palo Alto Networks provides several predefined profiles through Application content updates. The predefined profiles are global to the firewall, whereas custom profiles apply to a single virtual system only.

Any new Syslog Parse profiles in a given content release is documented in the corresponding release note along with the specific regex used to define the filter.

1. Install the latest Applications or Applications and Threats update:
   1. Select **Device** > **Dynamic Updates** and **Check Now**.
   2. **Download** and **Install** any new update.
2. Determine which predefined Syslog Parse profiles are available:
   1. Select **Device** > **User Identification** > **User Mapping** and click **Add** in the Server Monitoring section.
   2. Set the **Type** to **Syslog Sender** and click **Add** in the Filter section. If the Syslog Parse profile you need is available, skip the steps for defining custom profiles.

STEP 2 | Define custom Syslog Parse profiles to create and delete user mappings.

Each profile filters syslog messages to identify either login events (to create user mappings) or logout events (to delete mappings), but no single profile can do both.

1. Review the syslog messages that the syslog sender generates to identify the syntax for login and logout events. This enables you to define the matching patterns when creating Syslog Parse profiles.

   *While reviewing syslog messages, also determine whether they include the domain name. If they don't, and your user mappings require domain names, enter the **Default Domain Name** when defining the syslog senders that the User-ID agent monitors (later in this procedure).*
2. Select **Device** > **User Identification** > **User Mapping** and edit the Palo Alto Networks User-ID Agent Setup.
3. Select **Syslog Filters** and **Add** a Syslog Parse profile.
4. Enter a name to identify the **Syslog Parse Profile**.
5. Select the **Type** of parsing to find login or logout events in syslog messages:
   - **Regex Identifier**—Regular expressions.
   - **Field Identifier**—Text strings.

   The following steps describe how to configure these parsing types.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 3 | (**Regex Identifier parsing only**) Define the regex matching patterns.

If the syslog message contains a standalone space or tab as a delimiter, use **\s** *for a space and* **\t** *for a tab.*

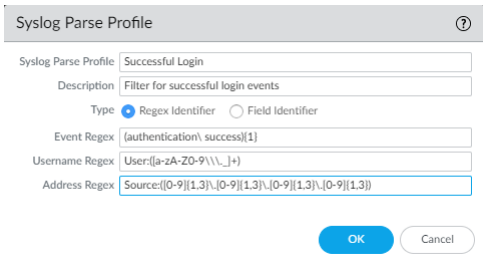
1. Enter the **Event Regex** for the type of events you want to find:
   - **Login events**—For the example message, the regex **`(authentication\ success){1}`** extracts the first **`{1}`** instance of the string `authenticationsuccess`.
   - **Logout events**—For the example message, the regex **`(logout\ successful){1}`** extracts the first **`{1}`** instance of the string `logoutsuccessful`.

   The backslash (\) before the space is a standard regex escape character that instructs the regex engine not to treat the space as a special character.
2. Enter the **Username Regex** to identify the start of the username.

   In the example message, the regex **`User:([a-zA-Z0-9\\\._]+)`** matches the string `User:johndoe1` and identifies `johndoe1` as the username.
3. Enter the **Address Regex** to identify the IP address portion of syslog messages.

   In the example message, the regular expression **`Source:([0-9]{1,3}\.[0-9]{1,3}\.[0-9]{1,3}\.[0-9]{1,3})`** matches the IPv4 address `Source:192.168.3.212`.

   The following is an example of a completed Syslog Parse profile that uses regex to identify login events:



4. Click **OK** twice to save the profile.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 4 | (**Field Identifier parsing only**) Define string matching patterns.

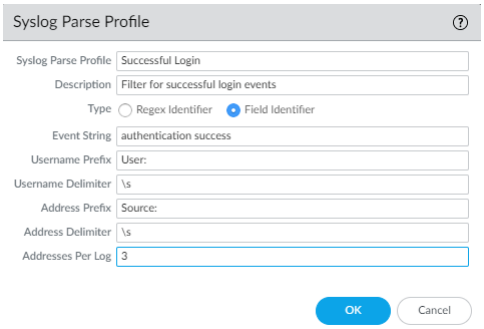
1. Enter an **Event String** to identify the type of events you want to find.
   - **Login events**—For the example message, the string `authentication success` identifies login events.
   - **Logout events**—For the example message, the string `logoutsuccessful` identifies logout events.
2. Enter a **Username Prefix** to identify the start of the username field in syslog messages. The field does not support regex expressions such as \s (for a space) or \t (for a tab).

   In the example messages, `User:` identifies the start of the username field.
3. Enter the **Username Delimiter** that indicates the end of the username field in syslog messages. Use **`\s`** to indicate a standalone space (as in the sample message) and **`\t`** to indicate a tab.
4. Enter an **Address Prefix** to identify the start of the IP address field in syslog messages. The field does not support regex expressions such as \s (for a space) or \t (for a tab).

   In the example messages, `Source:` identifies the start of the address field.
5. Enter the **Address Delimiter** that indicates the end of the IP address field in syslog messages.

   For example, enter **`\n`** to indicate the delimiter is a line break.

   The following is an example of a completed Syslog Parse profile that uses string matching to identify login events:



6. Click **OK** twice to save the profile.

https://t.me/learningnets

STEP 5 | Specify the syslog senders that the firewall monitors.

Within the total maximum of 100 monitored servers per firewall, you can define no more than 50 syslog senders for any single virtual system.

The firewall discards any syslog messages received from senders that are not on this list.

1. Select **Device** > **User Identification** > **User Mapping** and **Add** an entry to the Server Monitoring list.
2. Enter a **Name** to identify the sender.
3. Make sure the sender profile is **Enabled** (default is enabled).
4. Set the **Type** to **Syslog Sender**.
5. Enter the **Network Address** (IP address) of the syslog sender.
6. Select **SSL** (default) or **UDP** as the **Connection Type**.

   *To select the TLS certificate that the firewall uses to receive syslog messages, select **Device** > **User Identification** > **User Mapping** > **Palo Alto Networks User-ID Agent Setup**. **Edit** the settings and select **Server Monitor**, then select the **Syslog Service Profile** that contains the TLS certificate you want to the firewall to use to receive syslog messages.*

   The PAN-OS integrated User-ID agent accepts syslogs over SSL and UDP only. However, you must use caution when using UDP to receive syslog messages because it is an unreliable protocol and as such there is no way to verify that a message was sent from a trusted syslog sender. Although you can restrict syslog messages to specific source IP addresses, an attacker can still spoof the IP address, potentially allowing the injection of unauthorized syslog messages into the firewall.

   Always use SSL to listen for syslog messages because the traffic is encrypted (UDP sends the traffic in cleartext). If you must use UDP, make sure that the syslog sender and client are both on a dedicated, secure network to prevent untrusted hosts from sending UDP traffic to the firewall.

   A syslog sender using SSL to connect will show a Status of Connected only when there is an active SSL connection. Syslog senders using UDP will not show a Status value.
7. For each syslog format that the sender supports, **Add** a Syslog Parse profile to the Filter list. Select the **Event Type** that each profile is configured to identify: **login** (default) or **logout**.
8. (Optional) If the syslog messages don't contain domain information and your user mappings require domain names, enter a **Default Domain Name** to append to the mappings.
9. Click **OK** to save the settings.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 6 | Enable syslog listener services on the interface that the firewall uses to collect user mappings.

1. Select **Network** > **Network Profiles** > **Interface Mgmt** and edit an existing Interface Management profile or **Add** a new profile.
2. Select **User-ID Syslog Listener-SSL** or **User-ID Syslog Listener-UDP** or both, based on the protocols you defined for the syslog senders in the Server Monitoring list.

   The listening ports (514 for UDP and 6514 for SSL) are not configurable; they are enabled through the management service only.
3. Click **OK** to save the interface management profile.

   Even after enabling the User-ID Syslog Listener service on the interface, the interface only accepts syslog connections from senders that have a corresponding entry in the User-ID monitored servers configuration. The firewall discards connections or messages from senders that are not on the list.
4. Assign the Interface Management profile to the interface that the firewall uses to collect user mappings:
   1. Select **Network** > **Interfaces** and edit the interface.
   2. Select **Advanced** > **Other info**, select the Interface **Management Profile** you just added, and click **OK**.
5. **Commit** your changes.

STEP 7 | Verify that the firewall adds and deletes user mappings when users log in and out.

You can use CLI commands to see additional information about syslog senders, syslog messages, and user mappings.

1. Log in to a client system for which a monitored syslog sender generates login and logout event messages.
2. Log in to the firewall CLI.
3. Verify that the firewall mapped the login username to the client IP address:

```
> show user ip-user-mapping ip <ip-address>
IP address:    192.0.2.1 (vsys1)
User:          localdomain\username
From:          SYSLOG
```

4. Log out of the client system.
5. Verify that the firewall deleted the user mapping:

```
> show user ip-user-mapping ip <ip-address>
No matched record
```

Configure the Windows User-ID Agent as a Syslog Listener

To configure the Windows-based User-ID agent to create new user mappings and remove outdated mappings through syslog monitoring, start by defining Syslog Parse profiles. The User-ID

https://t.me/learningnets

agent uses the profiles to find login and logout events in syslog messages. In environments where *syslog senders* (the network services that authenticate users) deliver syslog messages in different formats, configure a profile for each syslog format. Syslog messages must meet certain criteria for a User-ID agent to parse them (see Syslog). This procedure uses examples with the following formats:

- Login events—`[Tue Jul 5 13:15:04 2016 CDT] Administrator authentication success User:johndoe1 Source:192.168.3.212`
- Logout events—`[Tue Jul 5 13:18:05 2016 CDT] User logout successful User:johndoe1 Source:192.168.3.212`

After configuring the Syslog Parse profiles, you specify the syslog senders that the User-ID agent monitors.

The Windows User-ID agent accepts syslogs over TCP and UDP only. However, you must use caution when using UDP to receive syslog messages because it is an unreliable protocol and as such there is no way to verify that a message was sent from a trusted syslog sender. Although you can restrict syslog messages to specific source IP addresses, an attacker can still spoof the IP address, potentially allowing the injection of unauthorized syslog messages into the firewall. As a best practice, use TCP instead of UDP. In either case, make sure that the syslog sender and client are both on a dedicated, secure VLAN to prevent untrusted hosts from sending syslogs to the User-ID agent.

STEP 1 | Deploy the Windows-based User-ID agents if you haven't already.

1. Install the Windows-Based User-ID Agent.
2. Configure the firewall to connect to the User-ID agent.

STEP 2 | Define custom Syslog Parse profiles to create and delete user mappings.

Each profile filters syslog messages to identify either login events (to create user mappings) or logout events (to delete mappings), but no single profile can do both.

1. Review the syslog messages that the syslog sender generates to identify the syntax for login and logout events. This enables you to define the matching patterns when creating Syslog Parse profiles.

   *While reviewing syslog messages, also determine whether they include the domain name. If they don't, and your user mappings require domain names, enter the **Default Domain Name** when defining the syslog senders that the User-ID agent monitors (later in this procedure).*
2. Open the Windows **Start** menu and select **User-ID Agent**.
3. Select **User Identification** > **Setup** and **Edit** the Setup.
4. Select **Syslog**, **Enable Syslog Service**, and **Add** a Syslog Parse profile.
5. Enter a **Profile Name** and **Description**.
6. Select the **Type** of parsing to find login and logout events in syslog messages:
   - **Regex**—Regular expressions.
   - **Field**—Text strings.

   The following steps describe how to configure these parsing types.

https://t.me/learningnets

STEP 3 | (**Regex parsing only**) Define the regex matching patterns.

If the syslog message contains a standalone space or tab as a delimiter, use **\s** for a space and **\t** for a tab.

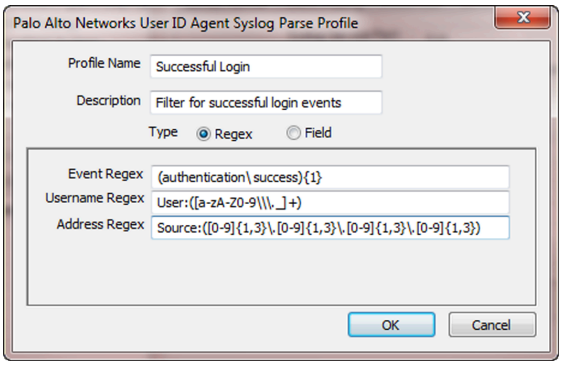
1. Enter the **Event Regex** for the type of events you want to find:
   - **Login events**—For the example message, the regex **(authentication\ success){1}** extracts the first **{1}** instance of the string `authentication success`.
   - **Logout events**—For the example message, the regex **(logout\ successful){1}** extracts the first **{1}** instance of the string `logout successful`.

   The backslash before the space is a standard regex escape character that instructs the regex engine not to treat the space as a special character.
2. Enter the **Username Regex** to identify the start of the username.

   In the example message, the regex **User:([a-zA-Z0-9\\\._]+)** matches the string `User:johndoe1` and identifies `johndoe1` as the username.
3. Enter the **Address Regex** to identify the IP address portion of syslog messages.

   In the example message, the regular expression **Source:([0-9]{1,3}\.[0-9]{1,3}\.[0-9]{1,3}\.[0-9]{1,3})** matches the IPv4 address `Source:192.168.3.212`.

   The following is an example of a completed Syslog Parse profile that uses regex to identify login events:



4. Click **OK** twice to save the profile.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 4 | (**Field Identifier parsing only**) Define string matching patterns.

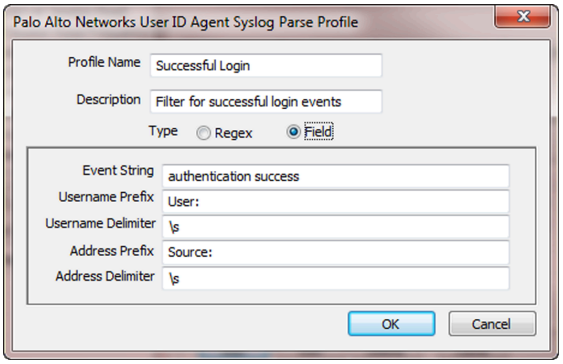
1. Enter an **Event String** to identify the type of events you want to find.
   - **Login events**—For the example message, the string `authentication success` identifies login events.
   - **Logout events**—For the example message, the string `logout successful` identifies logout events.
2. Enter a **Username Prefix** to identify the start of the username field in syslog messages. The field does not support regex expressions such as \s (for a space) or \t (for a tab).

   In the example messages, `User:` identifies the start of the username field.
3. Enter the **Username Delimiter** that indicates the end of the username field in syslog messages. Use **`\s`** to indicate a standalone space (as in the sample message) and **`\t`** to indicate a tab.
4. Enter an **Address Prefix** to identify the start of the IP address field in syslog messages. The field does not support regex expressions such as \s (for a space) or \t (for a tab).

   In the example messages, `Source:` identifies the start of the address field.
5. Enter the **Address Delimiter** that indicates the end of the IP address field in syslog messages.

   For example, enter **`\n`** to indicate the delimiter is a line break.

   The following is an example of a completed Syslog Parse profile that uses string matching to identify login events:



6. Click **OK** twice to save the profile.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 5 | Specify the syslog senders that the User-ID agent monitors.

Within the total maximum of 100 servers of all types that the User-ID agent can monitor, up to 50 can be syslog senders.

The User-ID agent discards any syslog messages received from senders that are not on this list.

1. Select **User Identification** > **Discovery** and **Add** an entry to the Servers list.
2. Enter a **Name** to identify the sender.
3. Enter the **Server Address** of the syslog sender (IP address or FQDN).
4. Set the **Server Type** to **Syslog Sender**.
5. (**Optional**) If you want to override the current domain in the username of your syslog message or prepend the domain to the username if your syslog message doesn't contain a domain, enter a **Default Domain Name**.
6. For each syslog format that the sender supports, **Add** a Syslog Parse profile to the Filter list. Select the **Event Type** that you configured each profile to identify—**login** (default) or **logout**—and then click **OK**.
7. Click **OK** to save the settings.
8. **Commit** your changes to the User-ID agent configuration.

STEP 6 | Verify that the User-ID agent adds and deletes user mappings when users log in and out.



You can use CLI commands to see additional information about syslog senders, syslog messages, and user mappings.

1. Log in to a client system for which a monitored syslog sender generates login and logout event messages.
2. Verify that the User-ID agent mapped the login username to the client IP address:
   1. In the User-ID agent, select **Monitoring**.
   2. Enter the username or IP address in the filter field, **Search**, and verify that the list displays the mapping.
3. Verify that the firewall received the user mapping from the User-ID agent:
   1. Log in to the firewall CLI.
   2. Run the following command:

```
> show user ip-user-mapping ip <ip-address>
```

If the firewall received the user mapping, the output resembles the following:

```
IP address:    192.0.2.1 (vsys1)
User:          localdomain\username
```

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

```
From:          SYSLOG
```

4. Log out of the client system.
5. Verify that the User-ID agent removed the user mapping:
   1. In the User-ID agent, select **Monitoring**.
   2. Enter the username or IP address in the filter field, **Search**, and verify that the list does not display the mapping.
6. Verify that the firewall deleted the user mapping:
   1. Access the firewall CLI.
   2. Run the following command:

      ```
      > show user ip-user-mapping ip <ip-address>
      ```

      If the firewall deleted the user mapping, the output displays:

      ```
      No matched record
      ```

Map IP Addresses to Usernames Using Authentication Portal

When a user initiates web traffic (HTTP or HTTPS) that matches an Authentication Policy rule, the firewall prompts the user to authenticate through Authentication Portal. This ensures that you know exactly who is accessing your most sensitive applications and data. Based on user information collected during authentication, the firewall creates a new IP address-to-username mapping or updates the existing mapping for that user. This method of user mapping is useful in environments where the firewall cannot learn mappings through other methods such as monitoring servers. For example, you might have users who are not logged in to your monitored domain servers, such as users on Linux clients.

- Authentication Portal Authentication Methods
- Authentication Portal Modes
- Configure Authentication Portal

Authentication Portal Authentication Methods

Authentication Portal uses the following methods to authenticate users whose web requests match Authentication Policy rules:

Authentication Method	Description
Kerberos SSO	The firewall uses Kerberos single sign-on (SSO) to transparently obtain user credentials from the browser. To use this method, your network requires a Kerberos infrastructure, including a key distribution center (KDC) with an authentication server and ticket granting service. The firewall must have a Kerberos account.

https://t.me/learningnets

Authentication Method	Description
	If Kerberos SSO authentication fails, the firewall falls back to web form or client certificate authentication, depending on your Authentication policy and Authentication Portal configuration.
Web Form	The firewall redirects web requests to a web form for authentication. For this method, you can configure Authentication policy to use Multi-Factor Authentication (MFA), SAML, Kerberos, TACACS+, RADIUS, or LDAP authentication. Although users have to manually enter their login credentials, this method works with all browsers and operating systems.
Client Certificate Authentication	The firewall prompts the browser to present a valid client certificate to authenticate the user. To use this method, you must provision client certificates on each user system and install the trusted certificate authority (CA) certificate used to issue those certificates on the firewall.

Authentication Portal Modes

The Authentication Portal mode defines how the firewall captures web requests for authentication:

Mode	Description
Transparent	The firewall intercepts the browser traffic per the Authentication policy rule and impersonates the original destination URL, issuing an HTTP 401 to invoke authentication. However, because the firewall does not have the real certificate for the destination URL, the browser displays a certificate error to users attempting to access a secure site. Therefore, use this mode only when absolutely necessary, such as in Layer 2 or virtual wire deployments.
Redirect	The firewall intercepts unknown HTTP or HTTPS sessions and redirects them to a Layer 3 interface on the firewall using an HTTP 302 redirect to perform authentication. This is the preferred mode because it provides a better end-user experience (no certificate errors). However, it does require additional Layer 3 configuration. Another benefit of the Redirect mode is that it provides for the use of session cookies, which enable the user to continue browsing to authenticated sites without requiring re-mapping each time the timeouts expire. This is especially useful for users who roam from one IP address to another (for example, from the corporate LAN to the wireless network) because they won't need to re-authenticate when the IP address changes as long as the session stays open.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Mode	Description
	If you use Kerberos SSO, you must use Redirect mode because the browser will provide credentials only to trusted sites. Redirect mode is also required if you use Multi-Factor Authentication to authenticate Authentication Portal users.

Configure Authentication Portal

The following procedure shows how to set up Authentication Portal authentication by configuring the PAN-OS integrated User-ID agent to redirect web requests that match an Authentication Policy rule to a firewall interface (redirect host).

SSL Inbound Inspection *does not support Authentication Portal redirect. To use Authentication Portal redirect and decryption, you must use* SSL Forward Proxy.

Based on their sensitivity, the applications that users access through Authentication Portal require different authentication methods and settings. To accommodate all authentication requirements, you can use default and custom authentication enforcement objects. Each object associates an Authentication rule with an authentication profile and an Authentication Portal authentication method.

- **Default authentication enforcement objects**—Use the default objects if you want to associate multiple Authentication rules with the same global authentication profile. You must configure this authentication profile before configuring Authentication Portal, and then assign it in the Authentication Portal Settings. For Authentication rules that require Multi-Factor Authentication (MFA), you cannot use default authentication enforcement objects.
- **Custom authentication enforcement objects**—Use a custom object for each Authentication rule that requires an authentication profile that differs from the global profile. Custom objects are mandatory for Authentication rules that require MFA. To use custom objects, create authentication profiles and assign them to the objects after configuring Authentication Portal—when you Configure Authentication Policy.

Keep in mind that authentication profiles are necessary only if users authenticate through a Authentication Portal Web Form or Kerberos SSO. Alternatively, or in addition to these methods, the following procedure also describes how to implement Client Certificate Authentication.

If you use Authentication Portal without the other User-ID functions (user mapping and group mapping), you don't need to configure a User-ID agent.

STEP 1 | Configure the interfaces that the firewall will use for incoming web requests, authenticating users, and communicating with directory servers to map usernames to IP addresses.

When the firewall connects to authentication servers or User-ID agents, it uses the management interface by default. As a best practice, isolate your management network by configuring service routes to connect to the authentication servers or User-ID agents.

1. (MGT interface only) Select **Device** > **Setup** > **Interfaces**, edit the **Management** interface, select **User-ID**, and click **OK**.
2. (Non-MGT interface only) Assign an Interface Management Profile to the Layer 3 interface that the firewall will use for incoming web requests and communication

https://t.me/learningnets

with directory servers. You must enable **Response Pages** and **User-ID** in the Interface Management profile.

3. (Non-MGT interface only) Configure a service route for the interface that the firewall will use to authenticate users. If the firewall has more than one virtual system (vsys), the service route can be global or vsys-specific. The services must include **LDAP** and potentially the following:
   - **Kerberos**, **RADIUS**, **TACACS+**, or **Multi-Factor Authentication**—Configure a service route for any authentication services that you use.
   - **UID Agent**—Configure this service only if you Enable User- and Group-Based Policy.
4. (Redirect mode only) Create a DNS address (A) record that maps the IP address on the Layer 3 interface to the redirect host. If you will use Kerberos SSO, you must also add a DNS pointer (PTR) record that performs the same mapping.

If your network doesn't support access to the directory servers from any firewall interface, you must Configure User Mapping Using the Windows User-ID Agent.

STEP 2 | Make sure Domain Name System (DNS) is configured to resolve your domain controller addresses.

To verify proper resolution, ping the server FQDN. For example:

```
admin@PA-220> ping host dc1.acme.com
```

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 3 | Configure clients to trust Authentication Portal certificates.

Required for redirect mode—to transparently redirect users without displaying certificate errors. You can generate a self-signed certificate or import a certificate that an external certificate authority (CA) signed.

To use a self-signed certificate, create a root CA certificate and use it to sign the certificate you will use for Authentication Portal:

1. Select **Device** > **Certificate Management** > **Certificates** > **Device Certificates**.
2. Create a Self-Signed Root CA Certificate or import a CA certificate (see Import a Certificate and Private Key).
3. Generate a Certificate to use for Authentication Portal. Be sure to configure the following fields:
   - **Common Name**—Enter the DNS name of the intranet host for the Layer 3 interface.
   - **Signed By**—Select the CA certificate you just created or imported.
   - Certificate Attributes—Click **Add**, for the **Type** select **IP** and, for the **Value**, enter the IP address of the Layer 3 interface to which the firewall will redirect requests.
4. Configure an SSL/TLS Service Profile. Assign the Authentication Portal certificate you just created to the profile.

   If you don't assign an SSL/TLS Service Profile, the firewall uses TLS 1.2 by default. To use a different TLS version, configure an SSL/TLS Service Profile for the TLS version you want to use.
5. Configure clients to trust the certificate:
   1. Export the CA certificate you created or imported.
   2. Import the certificate as a trusted root CA into all client browsers, either by manually configuring the browser or by adding the certificate to the trusted roots in an Active Directory (AD) Group Policy Object (GPO).

STEP 4 | (Optional) Configure Client Certificate Authentication.

You don't need an authentication profile or sequence for client certificate authentication. If you configure both an authentication profile/sequence and certificate authentication, users must authenticate using both.

1. Use a root CA certificate to generate a client certificate for each user who will authenticate through Authentication Portal. The CA in this case is usually your enterprise CA, not the firewall.
2. Export the CA certificate in PEM format to a system that the firewall can access.
3. Import the CA certificate onto the firewall: see Import a Certificate and Private Key. After the import, click the imported certificate, select **Trusted Root CA**, and click **OK**.
4. Configure a Certificate Profile.
   - In the **Username Field** drop-down, select the certificate field that contains the user identity information.
   - In the **CA Certificates** list, click **Add** and select the CA certificate you just imported.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 5 | (Optional) Configure Authentication Portal for the Apple Captive Network Assistant.

This step is only required if you are using Authentication Portal with the Apple Captive Network Assistant (CNA). To use Authentication Portal with CNA, perform the following steps.

1. Verify you have specified an FQDN for the redirect host (not just an IP address).
2. Select an SSL/TLS service profile that uses a publicly-signed certificate for the specified FQDN.
3. Enter the following command to adjust the number of requests supported for Authentication Portal: **set deviceconfig setting ctd cap-portal-ask-requests <*threshold-value*>**

   By default, the firewall has a rate limit threshold for Authentication Portal that limits the number of requests to one request every two seconds. The CNA sends multiple requests that can exceed this limit, which can result in a TCP reset and an error from the CNA. The recommended threshold value is 5 (default is one). This value will allow up to 5 requests every two seconds. Based on your environment, you may need to configure a different value. If the current value is not sufficient to handle the number of requests, increase the value.

STEP 6 | Configure the Authentication Portal settings.

1. Select **Device** > **User Identification** > **Authentication Portal Settings** and edit the settings.
2. **Enable Authentication Portal** (default is enabled).
3. Specify the **Timer**, which is the maximum time in minutes that the firewall retains an IP address-to-username mapping for a user after that user authenticates through Authentication Portal (default is 60; range is 1 to 1,440). After the **Timer** expires, the

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

firewall removes the mapping and any associated Authentication Timestamps used to evaluate the **Timeout** in Authentication policy rules.

*When evaluating the Authentication Portal **Timer** and the **Timeout** value in each Authentication policy rule, the firewall prompts the user to re-authenticate for whichever setting expires first. Upon re-authenticating, the firewall resets the time count for the Authentication Portal **Timer** and records new authentication timestamps for the user. Therefore, to enable different **Timeout** periods for different Authentication rules, set the Authentication Portal **Timer** to a value the same as or higher than any rule **Timeout**.*

4. Select the **SSL/TLS Service Profile** you created for redirect requests over TLS. See Configure an SSL/TLS Service Profile.
5. Select the **Mode** (in this example, **Redirect**).
6. (Redirect mode only) Specify the **Redirect Host**, which is the intranet hostname (a hostname with no period in its name) that resolves to the IP address of the Layer 3 interface on the firewall to which web requests are redirected.

   If users authenticate through Kerberos single sign-on (SSO), the **Redirect Host** must be the same as the hostname specified in the Kerberos keytab.
7. Select the fall back authentication method to use:
   - To use client certificate authentication, select the **Certificate Profile** you created.
   - To use global settings for interactive or SSO authentication, select the **Authentication Profile** you configured.
   - To use Authentication policy rule-specific settings for interactive or SSO authentication, assign authentication profiles to authentication enforcement objects when you Configure Authentication Policy.
8. Click **OK** and **Commit** the Authentication Portal configuration.

STEP 7 | Next steps...

The firewall does not display the Authentication Portal web form to users until you Configure Authentication Policy rules that trigger authentication when users request services or applications.

Configure User Mapping for Terminal Server Users

Individual terminal server users appear to have the same IP address and therefore an IP address-to-username mapping is not sufficient to identify a specific user. To identify specific users on Windows-based terminal servers, the Palo Alto Networks Terminal Server agent (TS agent) allocates a port range to each user. The TS agent then notifies every connected firewall about the allocated port range, which allows the firewall to create an IP address-port-user mapping table and enable user- and group-based security policy enforcement. For non-Windows terminal servers, configure the PAN-OS XML API to extract user mapping information. The following values apply for both methods:

- Default port range: 1025 to 65534
- Per user block size: 200
- Maximum number of multi-user systems: 2,500

https://t.me/learningnets

For information about the terminal servers supported by the TS agent and the number of TS agents supported on each firewall model, refer to the Palo Alto Networks Compatibility Matrix and the Product Comparison Tool.

The following sections describe how to configure user mapping for terminal server users:

- Configure the Palo Alto Networks Terminal Server (TS) Agent for User Mapping
- Retrieve User Mappings from a Terminal Server Using the PAN-OS XML API

Configure the Palo Alto Networks Terminal Server (TS) Agent for User Mapping

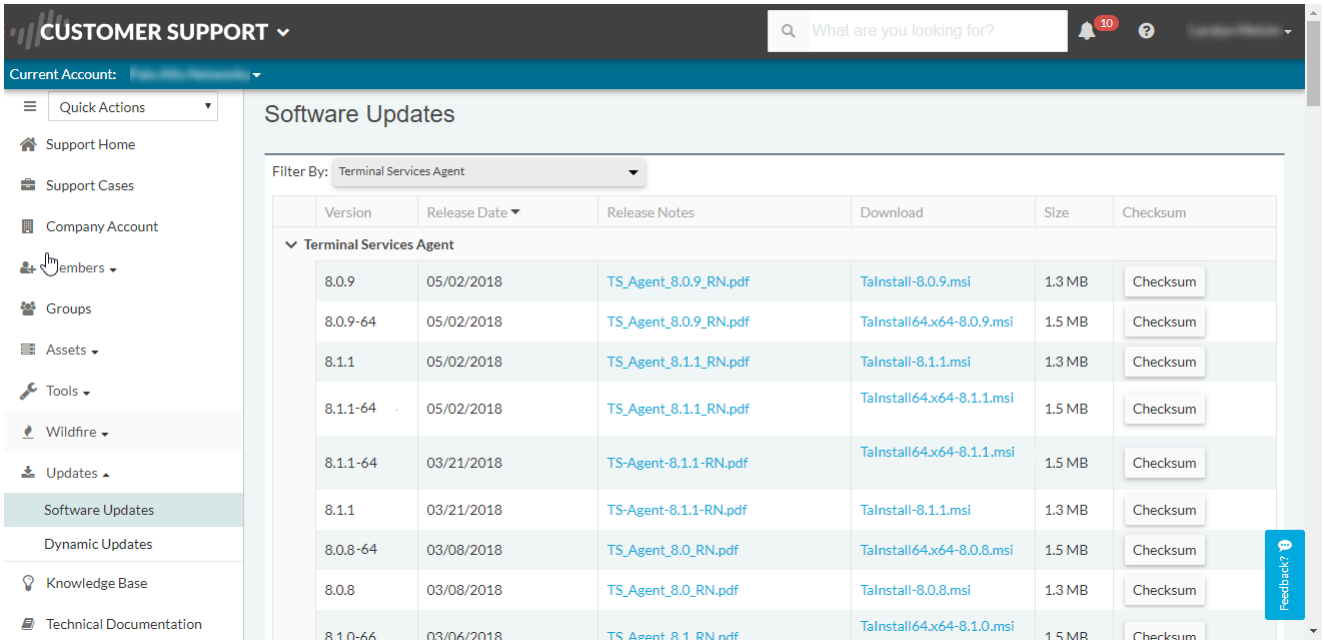
Use the following procedure to install and configure the TS agent on the terminal server. To map all your users, you must install the TS agent on all terminal servers to which your users log in.

If you are using TS agent 7.0 or a later version, disable any Sophos antivirus software on the TS agent host. Otherwise, the antivirus software overwrites the source ports that the TS agent allocates.

For information about default values, ranges, and other specifications, refer to Configure User Mapping for Terminal Server Users. For information about the terminal servers supported by the TS agent and the number of TS agents supported on each firewall model, refer to the Palo Alto Networks Compatibility Matrix.

STEP 1 | Download the TS agent installer.

1. Log in to the Palo Alto Networks Customer Support Portal.
2. Select **Updates** > **Software Updates**.
3. Set **Filter By** to **Terminal Services Agent** and select the version of the agent you want to install from the corresponding Download column. For example, to download TS agent 9.0, select **TaInstall-9.0.msi**.
4. Save the `TaInstall.x64-x.x.x-xx.msi` or `TaInstall-x.x.x-xx.msi` file on the systems where you plan to install the agent; be sure to select the appropriate version based on whether the Windows system is running a 32-bit or a 64-bit OS.



https://t.me/learningnets

STEP 2 | Run the installer as an administrator.

1. Open the Windows **Start** menu, right-click the **Command Prompt** program, and **Run as administrator**.
2. From the command line, run the .msi file you downloaded. For example, if you saved the `TaInstall-9.0.msi` file to the Desktop, then enter the following:

```
C:\Users\administrator.acme>cd Desktop
C:\Users\administrator.acme\Desktop>TaInstall-9.0.0-1.msi
```

3. Follow the setup prompts to install the agent using the default settings. The setup installs the agent in `C:\ProgramFiles\Palo Alto Networks\Terminal Server Agent`.

   To ensure correct port allocation, you must use the default Terminal Server agent installation folder location.

4. When the installation completes, **Close** the setup dialog.

   If you are upgrading to a TS agent version that has a newer driver than the existing installation, the installation wizard prompts you to reboot the system after you upgrade.

STEP 3 | Define the range of ports for the TS agent to allocate to end users.

*The **System Source Port Allocation Range** and **System Reserved Source Ports** specify the range of ports that are allocated to non-user sessions. Make sure the values in these fields do not overlap with the ports you designate for user traffic. These values can be changed only by editing the corresponding Windows registry settings. The TS agent does not allocate ports for network traffic emitted by session 0.*

1. Open the Windows **Start** menu and select **Terminal Server Agent** to launch the Terminal Server agent application.
2. **Configure** (side menu) the agent.
3. Enter the **Source Port Allocation Range** (default is 20,000-39,999). This is the full range of port numbers that the TS agent will allocate for user mapping. The port range you specify cannot overlap the **System Source Port Allocation Range**.
4. (Optional) If there are ports or port ranges within the source port allocation that you do not want the TS agent to allocate to user sessions, specify them as **Reserved Source Ports**. To include multiple ranges, use commas with no spaces (for example: **`2000-3000,3500,4000-5000`**).
5. Specify the number of ports to allocate to each individual user upon login to the terminal server (**Port Allocation Start Size Per User**); default is 200.
6. Specify the **Port Allocation Maximum Size Per User**, which is the maximum number of ports the Terminal Server agent can allocate to an individual user.
7. Specify whether to continue processing traffic from the user if the user runs out of allocated ports. The **Fail port binding when available ports are used up** option is enabled by default, which indicates that the application will fail to send traffic when all ports are

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

used. To enable users to continue using applications when they run out of ports, disable (clear) this option, but if you do, this traffic may not be identified with User-ID.

8. If the terminal server stops responding when you attempt to shut it down, enable the **Detach agent driver at shutdown** option.

STEP 4 | (Optional) Assign your own certificates for mutual authentication between the TS agent and the firewall.

1. Obtain your certificate for the TS agent from your enterprise PKI or generate one on your firewall. The private key of the server certificate must be encrypted and the certificate must be uploaded in PEM file format. Perform one of the following tasks to upload a certificate:
   - Generate a Certificate and export it.
   - Export a certificate from your enterprise certificate authority (CA).
2. Add a server certificate to the TS agent.
   1. On the TS agent, select **Server Certificate** and **Add** a new certificate.
   2. Enter the path and name of the certificate file received from the CA or browse to the certificate file.
   3. Enter the private key password.
   4. Click **OK**.
   5. **Commit** your changes.

   The TS agent uses a self-signed certificate on port 5009 with following information:Issuer: CN=Terminal Server Agent, OU=Engineering, O=Palo Alto Networks, L=Santa Clara, S=California, C=USSubject: CN=Terminal Server Agent, OU=Engineering, O=Palo Alto Networks, L=Santa Clara, S=California, C=US
3. Configure and assign the certificate profile for the firewall.
   1. Select **Device** > **Certificate Management** > **Certificate Profile** to Configure a Certificate Profile.

      You can assign only one certificate profile for Windows User-ID agents and TS agents. Therefore, your certificate profile must include all certificate authorities that issued certificates uploaded to connected Windows User-ID and TS agents.
   2. Select **Device** > **User Identification** > **Connection Security**.
   3. Edit (⚙) and select the certificate profile you configured in the previous step as the **User-ID Certificate Profile**.
   4. Click **OK**.
   5. **Commit** your changes.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 5 | Configure the firewall to connect to the Terminal Server agent.

Complete the following steps on each firewall you want to connect to the Terminal Server agent to receive user mappings:

1. Select **Device** > **User Identification** > **Terminal Server Agents** and **Add** a new TS agent.
2. Enter a **Name** for the Terminal Server agent.
3. Enter the hostname or IP address of the Windows **Host** on which the Terminal Server agent is installed.

   The hostname or IP address must resolve to a static IP address. If you change the existing hostname, the TS agent resets when you commit the changes to resolve the new hostname. If the hostname resolves to multiple IP addresses, the TS agent uses the first address in the list.
4. (Optional) Enter the hostname or IP address for any **Alternative IP Addresses** that can appear as the source IP address for the outgoing traffic.

   The hostname or IP address must resolve to a static IP address. You can enter up to 8 IP addresses or hostnames.
5. Enter the **Port** number on which the agent will listen for user mapping requests. This value must match the value configured on the Terminal Server agent. By default, the port is set to 5009 on the firewall and on the agent. If you change it on the firewall, you must also change the **Listening Port** on the Terminal Server agent **Configure** dialog to the same port.
6. Make sure that the configuration is **Enabled** and then click **OK**.
7. **Commit** your changes.
8. Verify that the **Connected** status displays as connected (a green light).

STEP 6 | Verify that the Terminal Server agent is successfully mapping IP addresses to usernames and that the firewalls can connect to the agent.

1. Open the Windows **Start** menu and select **Terminal Server Agent**.
2. Verify that the firewalls can connect by making sure the **Connection Status** of each firewall in the Connection List is `Connected`.
3. Verify that the Terminal Server agent is successfully mapping port ranges to usernames (**Monitor** in the side menu) and confirm that the mapping table is populated.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 7 | (Windows 2012 R2 servers only) Disable Enhanced Protected Mode in Microsoft Internet Explorer for each user who uses that browser.

This task is not necessary for other browsers, such as Google Chrome or Mozilla Firefox.

To disable Enhanced Protected Mode for all users, use Local Security Policy.

Perform these steps on the Windows Server:

1. Start Internet Explorer.
2. Select **Settings** > **Internet options** > **Advanced** and scroll to the Security section.
3. Disable (clear) the **Enable Enhanced Protected Mode** option.
4. Click **OK**.

In Internet Explorer, Palo Alto Networks recommends that you do not disable Protected Mode, which differs from Enhanced Protected Mode.

Retrieve User Mappings from a Terminal Server Using the PAN-OS XML API

The PAN-OS XML API uses standard HTTP requests to send and receive data. API calls can be made directly from command line utilities such as cURL or using any scripting or application framework that supports RESTful services.

To enable a non-Windows terminal server to send user mapping information directly to the firewall, create scripts that extract the user login and logout events and use them for input to the PAN-OS XML API request format. Then define the mechanisms for submitting the XML API request(s) to the firewall using cURL or wget and providing the firewall's API key for secure communication. Creating user mappings from multi-user systems such as terminal servers requires use of the following API messages:

- **<multiusersystem>**—Sets up the configuration for an XML API Multi-user System on the firewall. This message allows for definition of the terminal server IP address (this will be the source address for all users on that terminal server). In addition, the **<multiusersystem>** setup message specifies the range of source port numbers to allocate for user mapping and the number of ports to allocate to each individual user upon login (called the *block size*). If you want to use the default source port allocation range (1025-65534) and block size (200), you do not need to send a **<multiusersystem>** setup event to the firewall. Instead, the firewall will automatically generate the XML API Multi-user System configuration with the default settings upon receipt of the first user login event message.
- **<blockstart>**—Used with the **<login>** and **<logout>** messages to indicate the starting source port number allocated to the user. The firewall then uses the block size to determine the actual range of port numbers to map to the IP address and username in the login message. For example, if the **<blockstart>** value is 13200 and the block size configured for the multi-user system is 300, the actual source port range allocated to the user is 13200 through 13499. Each connection initiated by the user should use a unique source port number within the allocated range, enabling the firewall to identify the user based on its IP address-port-user mappings for enforcement of user- and group-based security rules. When a user exhausts all the ports allocated, the terminal server must send a new **<login>** message allocating a new port range for the user so that the firewall can update the IP address-port-user mapping. In addition, a single username can have multiple blocks of ports mapped simultaneously. When the firewall

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

receives a **<logout>** message that includes a **<blockstart>** parameter, it removes the corresponding IP address-port-user mapping from its mapping table. When the firewall receives a **<logout>** message with a username and IP address, but no **<blockstart>**, it removes the user from its table. And, if the firewall receives a **<logout>** message with an IP address only, it removes the multi-user system and all mappings associated with it.

The XML files that the terminal server sends to the firewall can contain multiple message types and the messages do not need to be in any particular order within the file. However, upon receiving an XML file that contains multiple message types, the firewall will process them in the following order: multiusersystem requests first, followed by logins, then logouts.

The following workflow provides an example of how to use the PAN-OS XML API to send user mappings from a non-Windows terminal server to the firewall.

STEP 1 | Generate the API key that will be used to authenticate the API communication between the firewall and the terminal server. To generate the key you must provide login credentials for an administrative account; the API is available to all administrators (including role-based administrators with XML API privileges enabled).

Any special characters in the password must be URL/ percent-encoded.

From a browser, log in to the firewall. Then, to generate the API key for the firewall, open a new browser window and enter the following URL:

```
https://<Firewall-IPaddress>/api/?
type=keygen&user=<username>&password=<password>
```

Where **<Firewall-IPaddress>** is the IP address or FQDN of the firewall and **<username>** and **<password>** are the credentials for the administrative user account on the firewall. For example:

```
https://10.1.2.5/api/?type=keygen&user=admin&password=admin
```

The firewall responds with a message containing the key, for example:

```
<response status="success">
   <result>
      <key>k7J335J6hI7nBxIqyfa62sZugWx7ot%2BgzEA9UOnlZRg=</key>
   </result>
</response>
```

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 2 | (Optional) Generate a setup message that the terminal server will send to specify the port range and block size of ports per user that your Terminal Server agent uses.

If the Terminal Server agent does not send a setup message, the firewall will automatically create a Terminal Server agent configuration using the following default settings upon receipt of the first login message:

- Default port range: 1025 to 65534
- Per user block size: 200
- Maximum number of multi-user systems: 1,000

The following shows a sample setup message:

```
<uid-message>
  <payload>
<multiusersystem>
<entry ip="10.1.1.23" startport="20000"          endport="39999"
 blocksize="100/">
</multiusersystem>
</payload>
<type>update</type>
<version>1.0</version>
</uid-message>
```

where `entry ip` specifies the IP address assigned to terminal server users, `startport` and `endport` specify the port range to use when assigning ports to individual users, and `blocksize` specifies the number of ports to assign to each user. The maximum blocksize is 4000 and each multi-user system can allocate a maximum of 1000 blocks.

If you define a custom blocksize and or port range, keep in mind that you must configure the values such that every port in the range gets allocated and that there are no gaps or unused ports. For example, if you set the port range to 1000–1499, you could set the block size to 100, but not to 200. This is because if you set it to 200, there would be unused ports at the end of the range.

STEP 3 | Create a script that will extract the login events and create the XML input file to send to the firewall.

Make sure the script enforces assignment of port number ranges at fixed boundaries with no port overlaps. For example, if the port range is 1000–1999 and the block size is 200, acceptable blockstart values would be 1000, 1200, 1400, 1600, or 1800. Blockstart values of 1001, 1300, or 1850 would be unacceptable because some of the port numbers in the range would be left unused.



The login event payload that the terminal server sends to the firewall can contain multiple login events.

The following shows the input file format for a PAN-OS XML login event:

```
<uid-message>
<payload>
<login>
```

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

```
<entry name="acme\jjaso" ip="10.1.1.23" blockstart="20000">
<entry name="acme\jparker" ip="10.1.1.23" blockstart="20100">
<entry name="acme\ccrisp" ip="10.1.1.23" blockstart="21000">
</login>
</payload>
<type>update</type>
<version>1.0</version>
</uid-message>
```

The firewall uses this information to populate its user mapping table. Based on the mappings extracted from the example above, if the firewall received a packet with a source address and port of 10.1.1.23:20101, it would map the request to user jparker for policy enforcement.

Each multi-user system can allocate a maximum of 1,000 port blocks.

STEP 4 | Create a script that will extract the logout events and create the XML input file to send to the firewall.

Upon receipt of a `logout` event message with a `blockstart` parameter, the firewall removes the corresponding IP address-port-user mapping. If the `logout` message contains a username and IP address, but no `blockstart` parameter, the firewall removes all mappings for the user. If the `logout` message contains an IP address only, the firewall removes the multi-user system and all associated mappings.

The following shows the input file format for a PAN-OS XML logout event:

```
<uid-message>
<payload>
<logout>
<entry name="acme\jjaso" ip="10.1.1.23" blockstart="20000">
<entry name="acme\ccrisp" ip="10.1.1.23">
<entry ip="10.2.5.4">
</logout>
</payload>
<type>update</type>
<version>1.0</version>
</uid-message>
```

You can also clear the multiuser system entry from the firewall using the following CLI command: ***clear xml-api multiusersystem***

STEP 5 | Make sure that the scripts you create include a way to dynamically enforce that the port block range allocated using the XML API matches the actual source port assigned to the user on the terminal server and that the mapping is removed when the user logs out or the port allocation changes.

One way to do this would be to use netfilter NAT rules to hide user sessions behind the specific port ranges allocated via the XML API based on the uid. For example, to ensure that a

https://t.me/learningnets

user with the user ID jjaso is mapped to a source network address translation (SNAT) value of 10.1.1.23:20000-20099, the script you create should include the following:

```
[root@ts1 ~]# iptables -t nat -A POSTROUTING -m owner --uid-owner
 jjaso -p tcp -j SNAT --to-source 10.1.1.23:20000-20099
```

Similarly, the scripts you create should also ensure that the IP table routing configuration dynamically removes the SNAT mapping when the user logs out or the port allocation changes:

```
[root@ts1 ~]# iptables -t nat -D POSTROUTING 1
```

STEP 6 | Define how to package the XML input files containing the setup, login, and logout events into wget or cURL messages for transmission to the firewall.

To apply the files to the firewall using wget:

```
> wget --post file <filename> “https://<Firewall-
IPaddress>/api/?type=user-id&key=<key>&file-
name=<input_filename.xml>&client=wget&vsys=<VSYS_name>”
```

For example, the syntax for sending an input file named login.xml to the firewall at 10.2.5.11 using key k7J335J6hI7nBxIqyfa62sZugWx7ot%2BgzEA9UOnlZRg using wget would look as follows:

```
> wget --post file login.xml “https://10.2.5.11/api/?type=user-
id&key=k7J335J6hI7nBxIqyfa62sZugWx7ot%2BgzEA9UOnlZRg&file-
name=login.xml&client=wget&vsys=vsys1”
```

To apply the file to the firewall using cURL:

```
> curl --form file=@<filename> https://<Firewall-IPaddress>/api/?
type=user-id&key=<key>&vsys=<VSYS_name>
```

For example, the syntax for sending an input file named login.xml to the firewall at 10.2.5.11 using key k7J335J6hI7nBxIqyfa62sZugWx7ot%2BgzEA9UOnlZRg using cURL would look as follows:

```
> curl --form file@login.xml “https://10.2.5.11/api/?type=user-
id&key=k7J335J6hI7nBxIqyfa62sZugWx7ot%2BgzEA9UOnlZRg&vsys=vsys1”
```

STEP 7 | Verify that the firewall is successfully receiving login events from the terminal servers.

Verify the configuration by opening an SSH connection to the firewall and then running the following CLI commands:

To verify if the terminal server is connecting to the firewall over XML:

```
admin@PA-5250> show user xml-api multiusersystem
Host            Vsys    Users   Blocks
---------------------------------------
```

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

```
10.5.204.43   vsys1      5          2
```

To verify that the firewall is receiving mappings from a terminal server over XML:

```
admin@PA-5250> show user ip-port-user-mapping all

Global max host index 1, host hash count 1

XML API Multi-user System 10.5.204.43
Vsys 1, Flag 3
Port range: 20000 - 39999
Port size: start 200; max 2000
Block count 100, port count 20000
20000-20199: acme\administrator

Total host: 1
```

Send User Mappings to User-ID Using the XML API

User-ID provides many out-of-the box methods for obtaining user mapping information. However, you might have applications or devices that capture user information but cannot natively integrate with User-ID. For example, you might have a custom, internally developed application or a device that no standard user mapping method supports. In such cases, you can use the PAN-OS XML API to create custom scripts that send the information to the PAN-OS integrated User-ID agent or directly to the firewall. The PAN-OS XML API uses standard HTTP requests to send and receive data. API calls can be made directly from command line utilities such as cURL or using any scripting or application framework that supports POST and GET requests.

To enable an external system to send user mapping information to the PAN-OS integrated User-ID agent, create scripts that extract user login and logout events and use the events as input to the PAN-OS XML API request. Then define the mechanisms for submitting the XML API requests to the firewall (using cURL, for example) and use the API key of the firewall for secure communication. For more details, refer to the PAN-OS XML API Usage Guide.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Enable User- and Group-Based Policy

After you Enable User-ID, you will be able to configure Security Policy that applies to specific users and groups. User-based policy controls can also include application information (including which category and subcategory it belongs in, its underlying technology, or what the application characteristics are). You can define policy rules to safely enable applications based on users or groups of users, in either outbound or inbound directions.

Examples of user-based policies include:

- Enable only the IT department to use tools such as SSH, telnet, and FTP on standard ports.
- Allow the Help Desk Services group to use Slack.
- Allow all users to read Facebook, but block the use of Facebook apps, and restrict posting to employees in marketing.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Enable Policy for Users with Multiple Accounts

If a user in your organization has multiple responsibilities, that user might have multiple usernames (accounts), each with distinct privileges for accessing a particular set of services, but with all the usernames sharing the same IP address (the client system of the user). However, the User-ID agent can map any one IP address (or IP address and port range for terminal server users) to only one username for enforcing policy, and you can't predict which username the agent will map. To control access for all the usernames of a user, you must make adjustments to the rules, user groups, and User-ID agent.

For example, say the firewall has a rule that allows username corp_user to access email and a rule that allows username admin_user to access a MySQL server. The user logs in with either username from the same client IP address. If the User-ID agent maps the IP address to corp_user, then whether the user logs in as corp_user or admin_user, the firewall identifies that user as corp_user and allows access to email but not the MySQL server. On the other hand, if the User-ID agent maps the IP address to admin_user, the firewall always identifies the user as admin_user regardless of login and allows access to the MySQL server but not email. The following steps describe how to enforce both rules in this example.

STEP 1 | Configure a user group for each service that requires distinct access privileges.

In this example, each group is for a single service (email or MySQL server). However, it is common to configure each group for a set of services that require the same privileges (for example, one group for all basic user services and one group for all administrative services).

If your organization already has user groups that can access the services that the user requires, simply add the username that is used for less restricted services to those groups. In this example, the email server requires less restricted access than the MySQL server, and corp_user is the username for accessing email. Therefore, you add corp_user to a group that can access email (corp_employees) and to a group that can access the MySQL server (network_services).

If adding a username to a particular existing group would violate your organizational practices, you can create a custom group based on an LDAP filter. For this example, say network_services is a custom group, which you configure as follows:

1. Select **Device** > **User Identification** > **Group Mapping Settings** and **Add** a group mapping configuration with a unique **Name**.
2. Select an LDAP **Server Profile** and ensure the **Enabled** check box is enabled.
3. Select the **Custom Group** tab and **Add** a custom group with network_services as a **Name**.
4. Specify an **LDAP Filter** that matches an LDAP attribute of corp_user and click **OK**.
5. Click **OK** and **Commit**.



Later, if other users that are in the group for less restricted services are given additional usernames that access more restricted services, you can add those usernames to the group for more restricted services. This scenario is more common than the inverse; a user with access to more restricted services usually already has access to less restricted services.

https://t.me/learningnets

STEP 2 | Configure the rules that control user access based on the groups you just configured.

For more information, refer to Enable user- and group-based policy enforcement.

1. Configure a security rule that allows the corp_employees group to access email.
2. Configure a security rule that allows the network_services group to access the MySQL server.

STEP 3 | Configure the ignore list of the User-ID agent.

This ensures that the User-ID agent maps the client IP address only to the username that is a member of the groups assigned to the rules you just configured. The ignore list must contain all the usernames of the user that are not members of those groups.

In this example, you add admin_user to the ignore list of the Windows-based User-ID agent to ensure that it maps the client IP address to corp_user. This guarantees that, whether the user logs in as corp_user or admin_user, the firewall identifies the user as corp_user and applies both rules that you configured because corp_user is a member of the groups that the rules reference.

1. Create an `ignore_user_list.txt` file.
2. Open the file and add admin_user.

   If you later add more usernames, each must be on a separate line.
3. Save the file to the User-ID agent folder on the domain server where the agent is installed.



If you use the PAN-OS integrated User-ID agent, see Configure User Mapping Using the PAN-OS Integrated User-ID Agent for instructions on how to configure the ignore list.

STEP 4 | Configure endpoint authentication for the restricted services.

This enables the endpoint to verify the credentials of the user and preserves the ability to enable access for users with multiple usernames.

In this example, you have configured a firewall rule that allows corp_user, as a member of the network_services group, to send a service request to the MySQL server. You must now configure the MySQL server to respond to any unauthorized username (such as corp_user) by prompting the user to enter the login credentials of an authorized username (admin_user).



If the user logs in to the network as admin_user, the user can then access the MySQL server without it prompting for the admin_user credentials again.

In this example, both corp_user and admin_user have email accounts, so the email server won't prompt for additional credentials regardless of which username the user entered when logging in to the network.

The firewall is now ready to enforce rules for a user with multiple usernames.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Verify the User-ID Configuration

After you configure user and group mapping, enable User-ID in your Security policy, and configure Authentication policy, you should verify that User-ID works properly.

STEP 1 | Access the firewall CLI.

STEP 2 | Verify that group mapping is working.

From the CLI, enter the following operational command:

```
> show user group-mapping statistics
```

STEP 3 | Verify that user mapping is working.

If you are using the PAN-OS integrated User-ID agent, you can verify this from the CLI using the following command:

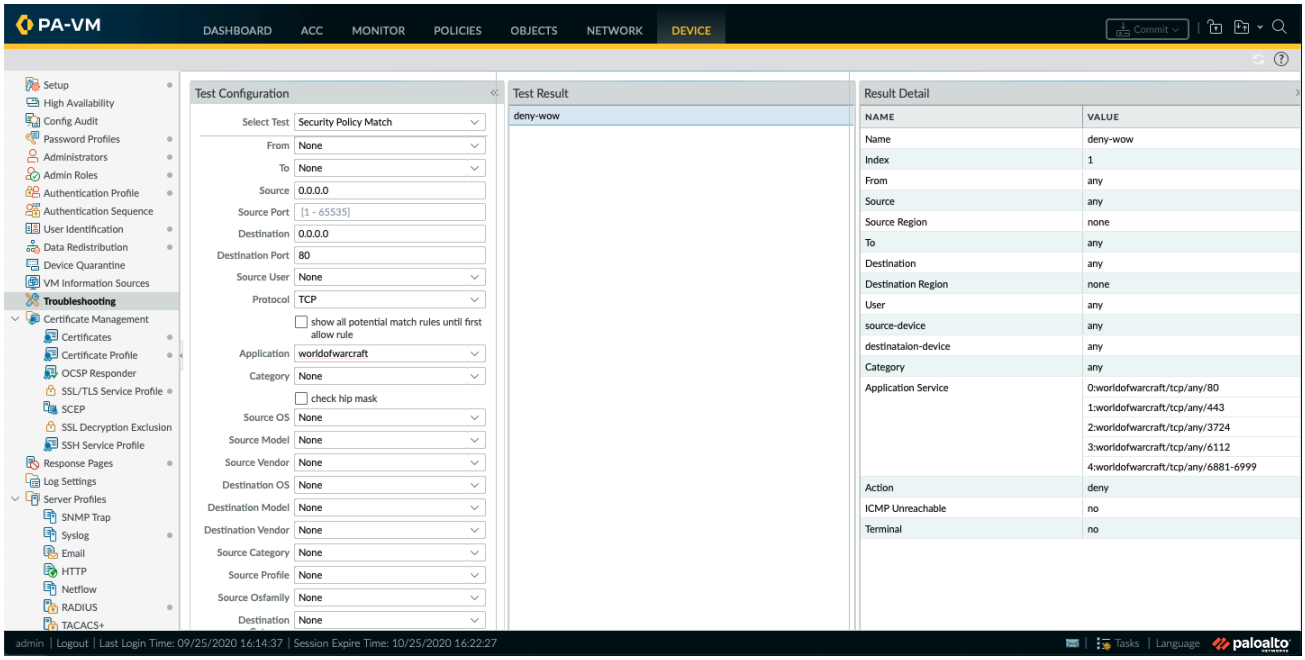
```
> show user ip-user-mapping-mp all
IP              Vsys  From  User         Timeout (sec)
------------------------------------------------------
192.168.201.1   vsys1 UIA   acme\george            210
192.168.201.11  vsys1 UIA   acme\duane             210
192.168.201.50  vsys1 UIA   acme\betsy             210
192.168.201.10  vsys1 UIA   acme\administrator     210
192.168.201.100 vsys1 AD    acme\administrator     748
Total: 5 users
*: WMI probe succeeded
```

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 4 | Test your Security policy rule.

- From a machine in the zone where User-ID is enabled, attempt to access sites and applications to test the rules you defined in your policy and ensure that traffic is allowed and denied as expected.
- You can also troubleshoot the running configuration to determine whether the policy is configured correctly. For example, suppose you have a rule that blocks users from playing World of Warcraft; you could test the policy as follows:
  1. Select **Device** > **Troubleshooting**, and select **Security Policy Match** from the Select Test drop-down.
  2. Enter **`0.0.0.0`** as the Source and Destination IP addresses. This executes the policy match test against any source and destination IP addresses.
  3. Enter the Destination Port.
  4. Enter the Protocol.
  5. **Execute** the security policy match test.



https://t.me/learningnets

STEP 5 | Test your Authentication policy and Authentication Portal configuration.

1. From the same zone, go to a machine that is not a member of your directory, such as a Mac OS system, and try to ping to a system external to the zone. The ping should work without requiring authentication.
2. From the same machine, open a browser and navigate to a web site in a destination zone that matches an Authentication rule you defined. The Authentication Portal web form should display and prompt you for login credentials.
3. Log in using the correct credentials and confirm that you are redirected to the requested page.
4. You can also test your Authentication policy using the **test authentication-policy-match** operational command as follows:

```
> test authentication-policy-match from corporate to internet
 source 192.168.201.10 destination 8.8.8.8
Matched rule: 'authentication portal' action: web-form
```

STEP 6 | Verify that the log files display usernames.

Select a logs page (such as **Monitor** > **Logs** > **Traffic**) and verify that the Source User column displays usernames.

STEP 7 | Verify that reports display usernames.

1. Select **Monitor** > **Reports**.
2. Select a report type that includes usernames. For example, the Denied Applications report, Source User column, should display a list of the users who attempted to access the applications.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Deploy User-ID in a Large-Scale Network

A large-scale network can have hundreds of information sources that firewalls query to map IP addresses to usernames and to map usernames to user groups. You can simplify User-ID administration for such a network by aggregating the user mapping and group mapping information before the User-ID agents collect it, thereby reducing the number of required agents.

A large-scale network can also have numerous firewalls that use the mapping information to enforce policies. You can reduce the resources that the firewalls and information sources use in the querying process by configuring some firewalls to acquire mapping information through redistribution instead of direct querying. Redistribution also enables the firewalls to enforce user-based policies when users rely on local sources for authentication (such as regional directory services) but need access to remote services and applications (such as global data center applications).

If you Configure Authentication Policy, your firewalls must also redistribute the Authentication Timestamps associated with user responses to authentication challenges. Firewalls use the timestamps to evaluate the timeouts for Authentication policy rules. The timeouts allow a user who successfully authenticates to later request services and applications without authenticating again within the timeout periods. Redistributing timestamps enables you to enforce consistent timeouts for each user even if the firewall that initially grants a user access is not the same firewall that later controls access for that user.

If you have configured multiple virtual systems, you can share IP address-to-username mapping information across virtual systems by selecting a virtual system as a User-ID hub.

- Deploy User-ID for Numerous Mapping Information Sources
- Redistribute Data and Authentication Timestamps
- Share User-ID Mappings Across Virtual Systems

Deploy User-ID for Numerous Mapping Information Sources

You can use Windows Log Forwarding and Global Catalog servers to simplify user mapping and group mapping in a large-scale network of Microsoft Active Directory (AD) domain controllers or Exchange servers. These methods simplify User-ID administration by aggregating the mapping information before the User-ID agents collect it, thereby reducing the number of required agents.

- Windows Log Forwarding and Global Catalog Servers
- Plan a Large-Scale User-ID Deployment
- Configure Windows Log Forwarding
- Configure User-ID for Numerous Mapping Information Sources

Windows Log Forwarding and Global Catalog Servers

Because each User-ID agent can monitor up to 100 servers, the firewall needs multiple User-ID agents to monitor a network with hundreds of AD domain controllers or Exchange servers. Creating and managing numerous User-ID agents involves considerable administrative overhead, especially in expanding networks where tracking new domain controllers is difficult. Windows Log Forwarding enables you to minimize the administrative overhead by reducing the number of servers to monitor and thereby reducing the number of User-ID agents to manage. When you

 ©2022 Palo Alto Networks, Inc.

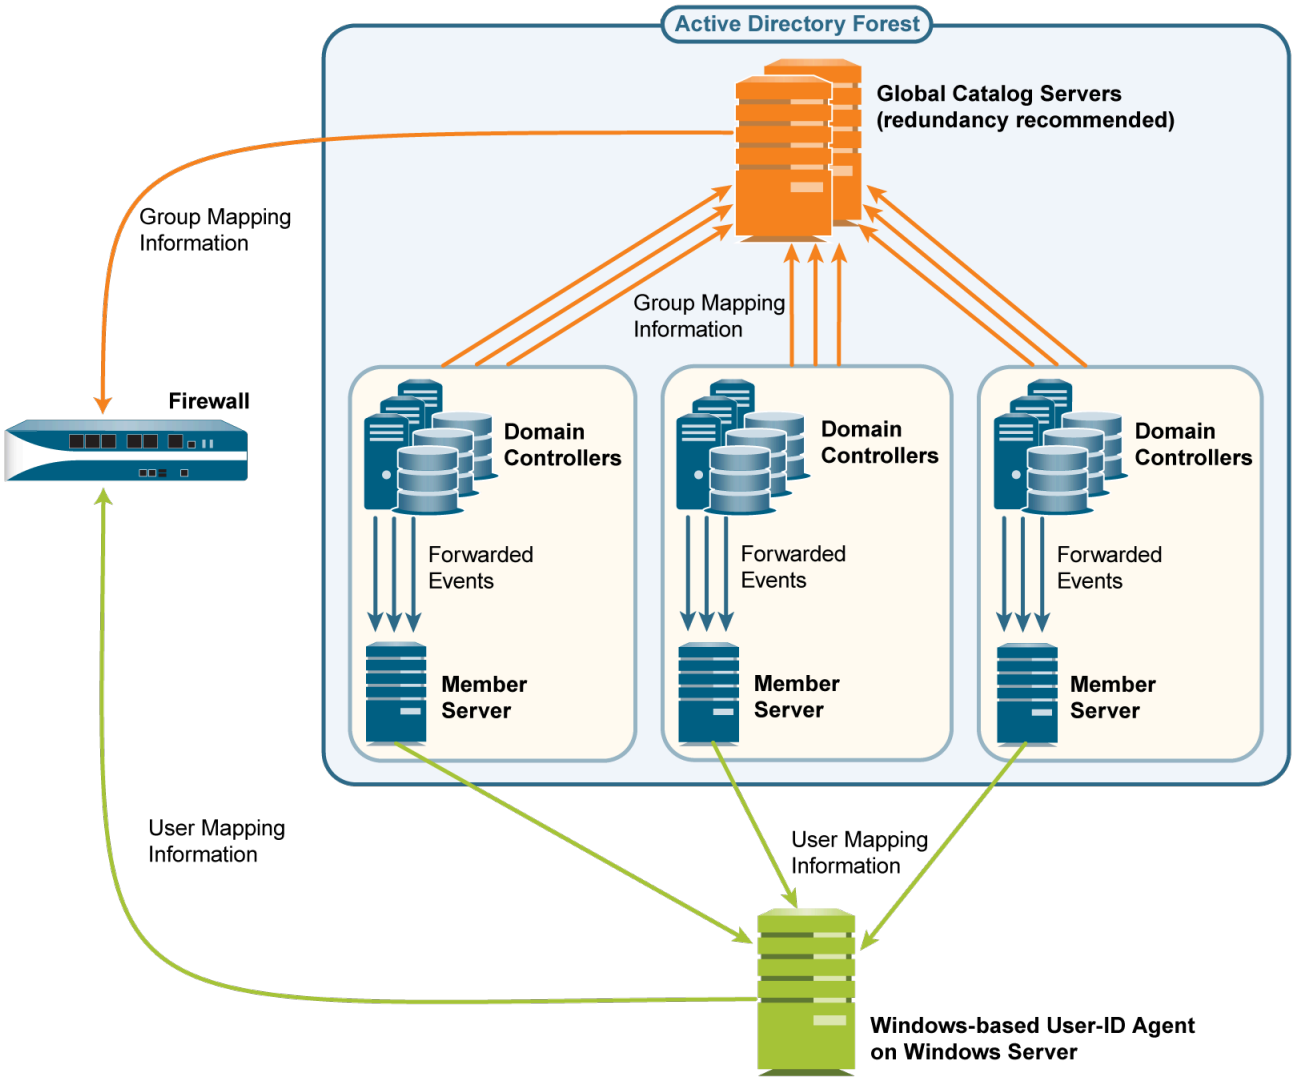
https://t.me/learningnets

configure Windows Log Forwarding, multiple domain controllers export their login events to a single domain member from which a User-ID agent collects the user mapping information.

You can configure Windows Log Forwarding for Windows Server versions 2012 and 2012 R2. Windows Log Forwarding is not available for non-Microsoft servers.

To collect group mapping information in a large-scale network, you can configure the firewall to query a Global Catalog server that receives account information from the domain controllers.

The following figure illustrates user mapping and group mapping for a large-scale network in which the firewall uses a Windows-based User-ID agent. See Plan a Large-Scale User-ID Deployment to determine if this deployment suits your network.



Plan a Large-Scale User-ID Deployment

When deciding whether to use Windows Log Forwarding and Global Catalog servers for your User-ID implementation, consult your system administrator to determine:

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

- Bandwidth required for domain controllers to forward login events to member servers. The bandwidth is a multiple of the login rate (number of logins per minute) of the domain controllers and the byte size of each login event.

  Domain controllers won't forward their entire security logs, they forward only the events that the user mapping process requires per login: four events for Windows Server 2012 and MS Exchange.
- Whether the following network elements support the required bandwidth:
  - **Domain controllers**—Must support the processing load associated with forwarding the events.
  - **Member Servers**—Must support the processing load associated with receiving the events.
  - **Connections**—The geographic distribution (local or remote) of the domain controllers, member servers, and Global Catalog servers is a factor. Generally, a remote distribution supports less bandwidth.

Configure Windows Log Forwarding

To configure Windows Log Forwarding, you need administrative privileges for configuring group policies on Windows servers. Configure Windows Log Forwarding on all the *Windows Event Collectors*—the member servers that collect login events from domain controllers. The following is an overview of the tasks; consult your Windows Server documentation for the specific steps.

STEP 1 | On each Windows Event Collector, enable event collection, add the domain controllers as event sources, and configure the event collection query (subscription). The events you specify in the subscription vary by domain controller platform:

- **Windows Server 2012 (including R2) and 2016, or MS Exchange**—The event IDs for the required events are 4768 (Authentication Ticket Granted), 4769 (Service Ticket Granted), 4770 (Ticket Granted Renewed), and 4624 (Logon Success).

  *To forward events as quickly as possible, **Minimize Latency** when configuring the subscription.*

  User-ID agents monitor the Security log on Windows Event Collectors, not the default forwarded events location. To change the event logging path to the Security log, perform the following steps on each Windows Event Collector.

  1. Open the Event Viewer.
  2. Right-click the **Security** log and select **Properties**.
  3. Copy the **Log path** (default **`%SystemRoot%\System32\Winevt\Logs\security.evtx`**) and click **OK**.
  4. Right-click the **Forwarded Events** folder and select **Properties**.
  5. Replace the default **Log path** (**`%SystemRoot%\System32\Winevt\Logs\ForwardedEvents.evtx`**) by pasting the value from the **Security** log, and then click **OK**.

STEP 2 | Configure a group policy to enable Windows Remote Management (WinRM) on the domain controllers.

STEP 3 | Configure a group policy to enable Windows Event Forwarding on the domain controllers.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Configure User-ID for Numerous Mapping Information Sources

STEP 1 | Configure Windows Log Forwarding on the member servers that will collect login events.

Configure Windows Log Forwarding. This step requires administrative privileges for configuring group policies on Windows servers.

STEP 2 | Install the Windows-based User-ID agent.

Install the Windows-Based User-ID Agent on a Windows server that can access the member servers. Make sure the system that will host the User-ID agent is a member of the same domain as the servers it will monitor.

STEP 3 | Configure the User-ID agent to collect user mapping information from the member servers.

1. Start the Windows-based User-ID agent.
2. Select **User Identification** > **Discovery** and perform the following steps for each member server that will receive events from domain controllers:
   1. In the Servers section, click **Add** and enter a **Name** to identify the member server.
   2. In the **Server Address** field, enter the FQDN or IP address of the member server.
   3. For the **Server Type**, select **Microsoft Active Directory**.
   4. Click **OK** to save the server entry.
3. Configure the remaining User-ID agent settings (refer to Configure the Windows-Based User-ID Agent for User Mapping).
4. If the User-ID sources provide usernames in multiple formats, specify the format for the **Primary Username** when you Map Users to Groups.

   The primary username is the username that identifies the user on the firewall and represents the user in reports and logs, regardless of the format that the User-ID source provides.

STEP 4 | Configure an LDAP server profile to specify how the firewall connects to the Global Catalog servers (up to four) for group mapping information.



To improve availability, use at least two Global Catalog servers for redundancy.

You can collect group mapping information only for universal groups, not local domain groups (subdomains).

1. Select **Device** > **Server Profiles** > **LDAP**, click **Add**, and enter a **Name** for the profile.
2. In the Servers section, for each Global Catalog, click **Add** and enter the server **Name**, IP address (**LDAP Server**), and **Port**. For a plaintext or Start Transport Layer Security (Start TLS) connection, use **Port** 3268. For an LDAP over SSL connection, use **Port** 3269. If the connection will use Start TLS or LDAP over SSL, select the **Require SSL/TLS secured connection** check box.
3. In the **Base DN** field, enter the Distinguished Name (DN) of the point in the Global Catalog server where the firewall will start searching for group mapping information (for example, `DC=acbdomain,DC=com`).
4. For the **Type**, select **active-directory**.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 5 | Configure an LDAP server profile to specify how the firewall connects to the servers (up to four) that contain domain mapping information.

User-ID uses this information to map DNS domain names to NetBIOS domain names. This mapping ensures consistent domain/username references in policy rules.

To improve availability, use at least two servers for redundancy.

The steps are the same as for the LDAP server profile you created for Global Catalogs in the previous step, except for the following fields:

- **LDAP Server**—Enter the IP address of the domain controller that contains the domain mapping information.
- **Port**—For a plaintext or Start TLS connection, use **Port** 389. For an LDAP over SSL connection, use **Port** 636. If the connection will use Start TLS or LDAP over SSL, select the **Require SSL/TLS secured connection** check box.
- **Base DN**—Select the DN of the point in the domain controller where the firewall will start searching for domain mapping information. The value must start with the string: `cn=partitions,cn=configuration` (for example, `cn=partitions,cn=configuration,DC=acbdomain,DC=com`).

STEP 6 | Create a group mapping configuration for each LDAP server profile you created.

1. Select **Device** > **User Identification** > **Group Mapping Settings**.
2. Click **Add** and enter a **Name** to identify the group mapping configuration.
3. Select the LDAP **Server Profile** and ensure the **Enabled** check box is selected.

   *If the Global Catalog and domain mapping servers reference more groups than your security rules require, configure the **Group Include List** and/or **Custom Group** list to limit the groups for which User-ID performs mapping.*
4. Click **OK** and **Commit**.

Insert Username in HTTP Headers

When you configure a secondary enforcement appliance with your Palo Alto Networks firewall to enforce user-based policy, the secondary appliance may not have the IP address-to-username mapping from the firewall. Transmitting user information to downstream appliances may require deployment of additional appliances such as proxies or negatively impact the user's experience (for example, users having to log in multiple times). By sharing the user's identity in the HTTP headers, you can enforce user-based policy without negatively impacting the user's experience or deploying additional infrastructure.

When you configure this feature, apply the URL profile to your Security policy, and commit your changes, the firewall:

1. Populates the user and domain values with the format of the primary username in the group mapping for the source user.
2. Encodes this information using Base64.
3. Adds the Base64-encoded header to the payload.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

4. Routes the traffic to the downstream appliance.

If you want to include the username and domain only when the user accesses specific domains, configure a domain list and the firewall inserts the header only when a domain in the list matches the Host header of the HTTP request.

To share user information with downstream appliances, you must first enable User-ID and configure group mapping.

To include the username and domain in the header, the firewall requires the IP address-to-username mapping for the user. If the user is not mapped, the firewall inserts `unknown` in Base64 encoding for both the domain and username in the header.

To include the username and domain in headers for HTTPS traffic, you must first create a decryption profile to decrypt HTTPS traffic.

This feature supports forward-proxy decryption traffic.

STEP 1 | Create or edit a **URL Filtering Profile**.

*The firewall does not insert headers if the action for the URL filtering profile is **`block`** for the domain.*

STEP 2 | Create or edit an HTTP header insertion entry using predefined types.

You can define up to five headers for each profile.

STEP 3 | Select **Dynamic Fields** as the header **Type**.

STEP 4 | **Add** the **Domains** where you want insert headers. When the user accesses a domain in the list, the firewall inserts the specified header.

STEP 5 | **Add** a new **Header** or select **X-Authenticated-User** to edit it.

STEP 6 | Select a header **Value** format (either **`($domain)\($user)`** or **`WinNT://($domain)/($user)`**) or enter your own format using the **`($domain)`** and **`($user)`** dynamic tokens (for example, **`($user)@($domain)`** for UserPrincipalName).

*Do not use the same dynamic token (either **`($user)`** or **`($domain)`**) more than once per value.*

Each value can be up to 512 characters. The firewall populates the **`($user)`** and **`($domain)`** dynamic tokens using the primary username in the group mapping profile. For example:

- If the primary username is the sAMAccountName, the value for **`($user)`** is the sAMAccountName and the value for **`($domain)`** is the NetBios domain name.
- If the primary username is the UserPrincipalName, the **`($user)`** the user account name (prefix) and the **`($domain)`** is the Domain Name System (DNS) name.

STEP 7 | (Optional) Select **Log** to enable logging for the header insertion.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 8 | Apply the URL filtering profile to the security policy rule for HTTP or HTTPS traffic.

STEP 9 | Select **OK** twice to confirm the HTTP header configuration.

STEP 10 | **Commit** your changes.

STEP 11 | Verify the firewall includes the username and domain in the HTTP headers.

- Use the **`show user user-ids all`** command to verify the group mapping is correct.
- Use the **`show counter global name ctd_header_insert`** command to view the number of HTTP headers inserted by the firewall.
- If you configured logging in Step 7, check the logs for the inserted Base64 encoded payload (for example, **`corpexample\testuser`** would appear in the logs as **`Y29ycGV4YW1wbGVcdGVzdHVzZXI=`**).

Redistribute Data and Authentication Timestamps

In a large-scale network, instead of configuring all your firewalls to directly query the mapping information sources, you can streamline resource usage by configuring some firewalls to collect mapping information through redistribution.

You can redistribute user mapping information collected through any method except Terminal Server (TS) agents. You cannot redistribute Group Mapping *or* HIP match *information.*

If you use Panorama to manage firewalls and aggregate firewall logs, you can use Panorama to manage User-ID redistribution*. Leveraging Panorama is a simpler solution than creating extra connections between firewalls to redistribute User-ID information.*

If you Configure Authentication Policy, your firewalls must also redistribute the Authentication Timestamps that are generated when users authenticate to access applications and services. Firewalls use the timestamps to evaluate the timeouts for Authentication policy rules. The timeouts allow a user who successfully authenticates to later request services and applications without authenticating again within the timeout periods. Redistributing timestamps enables you to enforce consistent timeouts across all the firewalls in your network.

Firewalls share data and authentication timestamps as part of the same redistribution flow; you don't have to configure redistribution for each information type separately.

- Firewall Deployment for Data Redistribution
- Configure Data Redistribution

Firewall Deployment for Data Redistribution

In a large-scale network, instead of configuring all your firewalls to directly query the data sources, you can streamline resource usage by configuring some firewalls to collect data through redistribution. Data redistribution also provides granularity, allowing you to redistribute only the types of information you specify to only the devices you select. You can also filter the IP user mappings or IP tag mappings using subnets and ranges to ensure the firewalls collect only the mappings they need to enforce policy.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Data redistribution can be unidirectional (the agent provides data to the client) or bidirectional, where both the agent and the client can simultaneously send and receive data.

To redistribute the data, you can use the following architecture types:

- **Hub and spoke architecture for a single region:**

  To redistribute data between firewalls, use a hub and spoke architecture as a best practice. In this configuration, a hub firewall collects the data from sources such as Windows User-ID agents, Syslog Servers, Domain Controllers, or other firewalls. Configure the redistribution client firewalls to collect the data from the hub firewall.

  For example, a hub (consisting of a pair of VM-50s for resiliency) could connect to the User-ID sources for the user mappings. The hub would then be able to redistribute the user mappings when the client firewalls that use the user mappings to enforce policy connect to the hub to receive data.

- **Multi-Hub and spoke architecture for multiple regions:**

  If you have firewalls deployed in multiple regions and want to distribute the data to the firewalls in all of these regions so that you can enforce policy consistently regardless of where the user logs in, you can use a multi-hub and spoke architecture for multiple regions.

  Start by configuring a firewall in each region to collect data from the sources. This firewall acts as a local hub for redistribution. This firewall collects the data from all sources in that region so that it can redistribute it to the client firewalls. Next, configure the client firewalls to connect to the redistribution hubs for their region and all other regions so that the client firewalls have all data from all hubs.

  As a best practice, enable bidirectional redistribution within a region if the firewalls need to both send and receive data. For example, if a firewall is acting as a GlobalProtect gateway for remote users and as a branch firewall for local users, the firewall must send the user mappings it collects for remote users to the hub firewall as well as receive the user mappings of the local users from the hub firewall.

- **Hierarchical architecture:**

  To redistribute data, you can also use a hierarchical architecture. For example, to redistribute data such as User-ID information, organize the redistribution sequence in layers, where each layer has one or more firewalls. In the bottom layer, PAN-OS integrated User-ID agents running on firewalls and Windows-based User-ID agents running on Windows servers map IP addresses to usernames. Each higher layer has firewalls that receive the mapping information and authentication timestamps from up to 100 redistribution points in the layer beneath it. The top-layer firewalls aggregate the mappings and timestamps from all layers. This deployment provides the option to configure policies for all users in top-layer firewalls and region- or function-specific policies for a subset of users in the corresponding domains served by lower-layer firewalls.

  In this scenario, three layers of firewalls redistribute mappings and timestamps from local offices to regional offices and then to a global data center. The data center firewall that aggregates all the information shares it with other data center firewalls so that they can all enforce policy and generate reports for users across your entire network. Only the bottom layer firewalls use User-ID agents to query the directory servers.

  The information sources that the User-ID agents query do not count towards the maximum of ten *hops* in the sequence. However, Windows-based User-ID agents that forward mapping

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

information to firewalls do count. Also in this example, the top layer has two hops: the first to aggregate information in one data center firewall and the second to share the information with other data center firewalls.

Configure Data Redistribution

Before you configure data redistribution:

- ❑ Plan the redistribution architecture. Some factors to consider are:
  - Which firewalls will enforce policies for all data types and which firewalls will enforce region- or function-specific policies for a subset of data?
  - How many hops does the redistribution sequence require to aggregate all data? The maximum allowed number of hops for user mappings is ten and the maximum allowed number of hops for IP address-to-username mappings and IP address-to-tag mappings is one.
  - How can you minimize the number of firewalls that query the user mapping information sources? The fewer the number of querying firewalls, the lower the processing load is on both the firewalls and sources.
- ❑ Configure the data sources from which your redistribution agents obtain the data to redistribute to their clients:
  - user mappings from PAN-OS Integrated User-ID agents or Windows-based User-ID agents
  - IP address-to-tag mappings for dynamic address groups
  - username-to-tag mappings for dynamic user groups
  - GlobalProtect for HIP-based Policy Enforcement
  - data for device quarantine (Panorama only)
- ❑ Configure Authentication Policy.

Data redistribution consists of:

- The redistribution agent that provides information
- The redistribution client that receives information

Perform the following steps on the firewalls in the data redistribution sequence.

STEP 1 | On a redistribution client firewall, configure a firewall, Panorama, or Windows User-ID agent as a data redistribution agent.

1. Select **Device** > **Data Redistribution** > **Agents**.
2. **Add** a redistribution agent and enter a **Name**.
3. Confirm that the agent is **Enabled**.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 2 | Add the agent using its **Serial Number** or its **Host and Port**.

- To add an agent using a serial number, select the **Serial Number** of the firewall you want to use as a redistribution agent.
- To add an agent using its host and port information:
  1. Enter the information for the **Host**.
  2. Select whether the host is an **LDAP Proxy**.
  3. Enter the **Port** (default is 5007, range is 1—65535).
  4. (Multiple virtual systems only) Enter the **Collector Name** to identify which virtual system you want to use as a redistribution agent.
  5. (Multiple virtual systems only) Enter and confirm the **Collector Pre-Shared Key** for the virtual system you want to use as a redistribution agent.

STEP 3 | Select one or more **Data Type** for the agent to redistribute.

- **IP User Mappings**—IP address-to-username mappings for User-ID.
- **IP Tags**—IP address-to-tag mappings for dynamic address groups.
- **User Tags**—Username-to-tag mappings for dynamic user groups.
- **HIP**—Host information profile (HIP) data from GlobalProtect, which includes HIP objects and profiles.
- **Quarantine List**—Devices that GlobalProtect identifies as quarantined.

STEP 4 | (Multiple virtual systems only) Configure a virtual system as a collector that can redistribute data.

Skip this step if the firewall receives but does not redistribute data.



You can redistribute information among virtual systems on different firewalls or on the same firewall. In both cases, each virtual system counts as one hop in the redistribution sequence.

1. Select **Device** > **Data Redistribution** > **Collector Settings**.
2. Edit the **Data Redistribution Agent Setup**.
3. Enter a **Collector Name** and **Pre-Shared Key** to identify this firewall or virtual system as a User-ID agent.
4. Click **OK** to save your changes.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 5 | (Optional but recommended) Configure which networks you want to include in data redistribution and which networks you want to exclude from data redistribution.

You can include or exclude networks and subnetworks when redistributing either IP address-to-tag mappings or IP address-to-username mappings.

As a best practice, always specify which networks to include and exclude to ensure that the agent is only communicating with internal resources.

1. Select **Device** > **Data Redistribution** > **Include/Exclude Networks**.
2. **Add** an entry and enter a **Name**.
3. Confirm that the entry is **Enabled**.
4. Select whether you want to **Include** or **Exclude** the entry.
5. Enter the **Network Address** for the entry.
6. Click **OK**.

STEP 6 | Configure the service route that the firewall uses to query other firewalls for User-ID information.

Skip this step if the firewall only receives user mapping information from Windows-based User-ID agents or directly from the information sources (such as directory servers) instead of from other firewalls.

1. Select **Device** > **Setup** > **Services**.
2. (**Firewalls with multiple virtual systems only**) Select **Global** (for a firewall-wide service route) or **Virtual Systems** (for a virtual system-specific service route), and then configure the service route.
3. Click **Service Route Configuration**, select **Customize**, and select **IPv4** or **IPv6** based on your network protocols. Configure the service route for both protocols if your network uses both.
4. Select **UID Agent** and then select the **Source Interface** and **Source Address**.
5. Click **OK** twice to save the service route.

STEP 7 | Enable the firewall to respond when other firewalls query it for data to redistribute.

Skip this step if the firewall receives but does not redistribute data.

Configure an Interface Management Profile with the **User-ID** service enabled and assign the profile to a firewall interface.

https://t.me/learningnets

STEP 8 | (Optional but recommended) Use a custom certificate from your enterprise PKI to establish a unique chain of trust from the redistribution client to the redistribution agent.

1. On the redistribution client firewall, create a custom SSL certificate profile to use for outgoing connections.
2. Select **Device** > **Setup** > **Management** > **Secure Communication Settings**.
3. **Edit** the settings.
4. Select the **Customize Secure Server Communication** option.
5. Select the **Certificate Profile** you created in Substep 1.
6. Click **OK**.
7. **Customize Communication** for **Data Redistribution**.
8. **Commit** your changes.
9. Enter the following CLI command to confirm the certificate profile (`SSL config`) uses `Custom certificates`: **`show redistribution agent state <agent-name>`** (where *<agent-name>* is the name of the redistribution agent or User-ID agent.

STEP 9 | (Optional but recommended) Use a custom certificate from your enterprise PKI to establish a unique chain of trust from the redistribution agent to the redistribution client.

1. On the redistribution agent firewall, create a custom SSL/TLS service profile for the firewall to use for incoming connections.
2. Select **Device** > **Setup** > **Management** > **Secure Communication Settings**.
3. **Edit** the settings.
4. Select the **Customize Secure Server Communication** option.
5. Select the **SSL/TLS Service Profile** you created in Step 1.
6. Click **OK**.
7. **Commit** your changes.
8. Enter the following CLI command to confirm the certificate profile (`SSL config`) uses `Custom certificates`: **`show redistribution service status`**.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 10 | Verify the agents correctly redistribute data to the clients.

1. View the agent statistics (**Device** > **Data Redistribution** > **Agents**) and select **Status** to view a summary of the activity for the redistribution agent, such as the number of mappings that the client firewall has received.
2. Confirm that the **Connected** status is **yes**.
3. On the agent, access the CLI and enter the following CLI command to check the status of the redistribution: **`show redistribution service status`**.
4. On the agent, enter the following CLI command to view the redistribution clients: **`show redistribution service client all`**.
5. On the client, enter the following CLI command to check the status of the redistribution: **`show redistribution service client all`**.
6. Confirm the **Source Name** in the User-ID logs (**Monitor** > **Logs** > **User-ID**) to verify that the firewall receives the mappings from the redistribution agents.
7. On the client, view the IP-Tag log (**Monitor** > **Logs** > **IP-Tag**) to confirm that the client firewall receives data.
8. On the client, enter the following CLI command and verify that the source the firewall receives the mappings `From` is `REDIST`: **`show user ip-user-mapping all`**.

STEP 11 | (Optional) To troubleshoot data redistribution, enable the traceroute option.

When you enable the traceroute option, the firewall that receives the data appends its IP address to the `<route>` field, which is a list of all firewall IP addresses that the data has traversed. This option requires that all PAN-OS devices in the redistribution route use PAN-OS version 10.0. If a PAN-OS device in the redistribution route uses PAN-OS 9.1.x or earlier versions, the traceroute information terminates at that device.

1. On the redistribution agent where the source originates, enter the following CLI command: **`debug user-id test cp-login traceroute yes ip-address <ip-address> user <username>`** (where ***`<ip-address>`*** is the IP address of the IP address-to-username mapping you want to verify and ***`<username>`*** is the username of the IP address-to-username mapping you want to verify.
2. On a client of the firewall where you configured the traceroute, verify the firewall redistributes the data by entering the following CLI command: **`show user ip-user-mapping all`**.

   The firewall displays the timestamp for the creation of the mapping (`SeqNumber`) and whether the user has GlobalProtect (`GP User`).

```
admin > show user ip-user-mapping-mp ip 192.0.2.0

IP address:  192.0.2.0 (vsys1)
User:    jimdoe
From:   REDIST
Timeout:  889s
Created:  11s ago
Origin:   198.51.100.0
SeqNumber:  15895329682-67831262
GP User:  No
Local HIP:  No
Route Node 0: 198.51.100.0 (vsys1)
```

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

```
Route Node 1: 198.51.100.1 (vsys1)
```

Share User-ID Mappings Across Virtual Systems

To simplify User-ID™ source configuration when you have multiple virtual systems, configure the User-ID sources on a single virtual system to share IP address-to-username mappings and username-to-group mappings with all other virtual systems on the firewall.

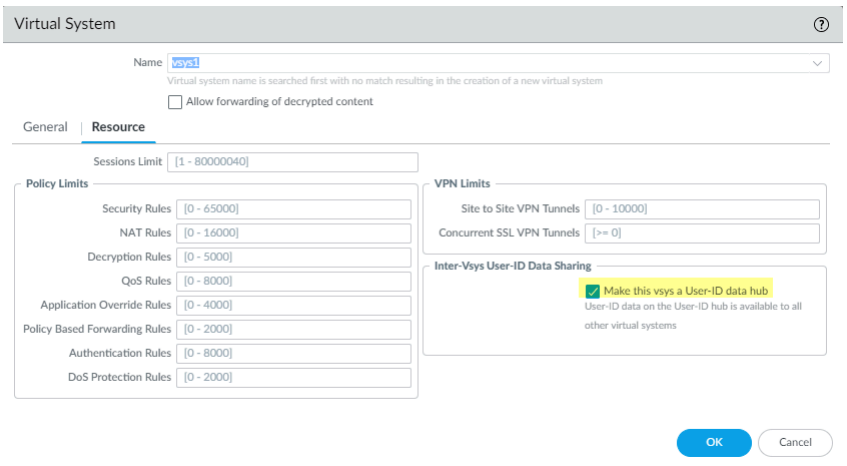
Configuring a single virtual system as a *User-ID hub* simplifies user mapping by eliminating the need to configure the sources on multiple virtual systems, especially if traffic will pass through multiple virtual systems based on the resources the user is trying to access (for example, in an academic networking environment where a student will be accessing different departments whose traffic is managed by different virtual systems).

To map the user or group, the firewall uses the mapping table on the local virtual system and applies the policy for that user or group. If the firewall does not find the mapping for a user or group on the virtual system where that user's traffic originated, the firewall queries the hub to fetch the IP address-to-username information for that user or group mapping information for that group. If the firewall locates the mapping on both the User-ID hub and the local virtual system, the firewall uses the mapping it learns locally. If the mapping on the local firewall differs from the mapping on the virtual system hub, the firewall uses the local mapping.

After you configure the User-ID hub, the virtual system can use the mapping table on the User-ID hub when it needs to identify a user for user-based policy enforcement or to display the username in a log or report but the source is not available locally. When you select a hub, the firewall retains the mappings on other virtual systems so we recommend consolidating the User-ID sources on the hub. However, if you don't want to share mappings from a specific source, you can configure an individual virtual system to perform user or group mapping.

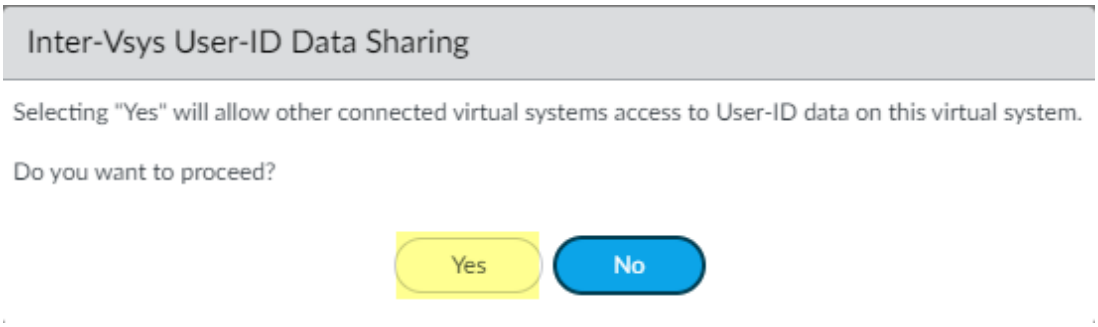
STEP 1 | Assign the virtual system as a User-ID hub.

1. Select **Device** > **Virtual Systems** and then select the virtual system where you consolidated your User-ID sources.
2. On the **Resource** tab, **Make this vsys a User-ID data hub** and click **Yes** to confirm. Then click **OK**.

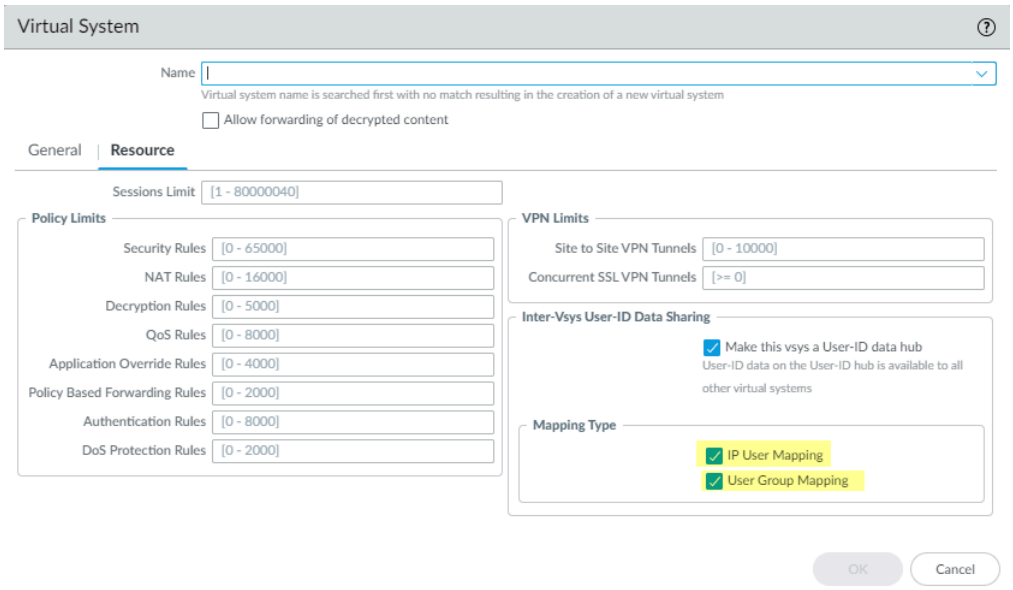


https://t.me/learningnets

STEP 2 | Click **Yes** to confirm.



STEP 3 | Select the **Mapping Type** that you want to share then click **OK**.



- **IP User Mapping**—Share IP address-to-username mapping information with other virtual systems.
- **User Group Mapping**—Share group mapping information with other virtual systems.

You must select at least one mapping type.

STEP 4 | Consolidate your User-ID sources and migrate them to the virtual system that you want to use as a User-ID hub.

This consolidates the User-ID configuration for operational simplicity. By configuring the hub to monitor servers and connect to agents that were previously monitored by other virtual systems, the hub collects the user mapping information instead of having each virtual system

https://t.me/learningnets

collect it independently. If you don't want to share mappings from specific virtual systems, configure those mappings on a virtual system that will not be used as the hub.

Use the same format for the Primary Username across virtual systems and firewalls.

1. Remove any sources that are unnecessary or outdated.
2. Identify all configurations for your Windows-based or integrated agents and any sources that send user mappings using the XML API and copy them to the virtual system you want to use as a User-ID hub.

   On the hub, you can configure any User-ID source that is currently configured on a virtual system. However, IP address-and-port-to-username mapping information from Terminal Server agents are not shared between the User-ID hub and the connected virtual systems.
3. Specify the subnetworks that User-ID should include in or exclude from mapping.
4. Define the **Ignore User List**.
5. On all other virtual systems, remove any sources that are on the User-ID hub.

STEP 5 | **Commit** the changes to enable the User-ID hub and begin collecting mappings for the consolidated sources.

STEP 6 | Confirm the User-ID hub is mapping the users and groups.

1. Use the **`show user ip-user-mapping all`** command to show the IP address-to-username mappings and which virtual system provides the mappings.
2. Use the **`show user user-id-agent statistics`** command to show which virtual system is serving as the User-ID hub.
3. Confirm the hub is sharing the group mappings by using the following CLI commands:
   - `show user group-mapping statistics`
   - `show user group-mapping state all`
   - `show user group list`
   - `show user group name` *<group-name>*

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

App-ID

To safely enable applications on your network, the Palo Alto Networks next-generation firewalls provide both an application and web perspective—App-ID and URL Filtering—to protect against a full spectrum of legal, regulatory, productivity, and resource utilization risks.

App-ID enables visibility into the applications on the network, so you can learn how they work and understand their behavioral characteristics and their relative risk. This application knowledge allows you to create and enforce security policy rules to enable, inspect, and shape desired applications and block unwanted applications. When you define policy rules to allow traffic, App-ID begins to classify traffic without any additional configuration.

New and modified App-IDs are released as part of Applications and Threat Content Updates—follow the Best Practices for Applications and Threats Content Updates to seamlessly keep your application and threat signatures up-to-date.

- App-ID Overview
- Streamlined App-ID Policy Rules
- App-ID and HTTP/2 Inspection
- Manage Custom or Unknown Applications
- Manage New and Modified App-IDs
- Use Application Objects in Policy
- Safely Enable Applications on Default Ports
- Applications with Implicit Support
- Security Policy Rule Optimization
- App-ID Cloud Engine
- SaaS App-ID Policy Recommendation
- Application Level Gateways
- Disable the SIP Application-level Gateway (ALG)
- Use HTTP Headers to Manage SaaS Application Access
- Maintain Custom Timeouts for Legacy Applications

https://t.me/learningnets

App-ID Overview

App-ID, a patented traffic classification system only available in Palo Alto Networks firewalls, determines what an application is irrespective of port, protocol, encryption (SSH or SSL) or any other evasive tactic used by the application. It applies multiple classification mechanisms —application signatures, application protocol decoding, and heuristics—to your network traffic stream to accurately identify applications.

Here's how App-ID identifies applications traversing your network:

- Traffic is matched against policy to check whether it is allowed on the network.
- Signatures are then applied to allowed traffic to identify the application based on unique application properties and related transaction characteristics. The signature also determines if the application is being used on its default port or it is using a non-standard port. If the traffic is allowed by policy, the traffic is then scanned for threats and further analyzed for identifying the application more granularly.
- If App-ID determines that encryption (SSL or SSH) is in use, and a Decryption policy rule is in place, the session is decrypted and application signatures are applied again on the decrypted flow.
- Decoders for known protocols are then used to apply additional context-based signatures to detect other applications that may be tunneling inside of the protocol (for example, Yahoo! Instant Messenger used across HTTP). Decoders validate that the traffic conforms to the protocol specification and provide support for NAT traversal and opening dynamic pinholes for applications such as SIP and FTP.
- For applications that are particularly evasive and cannot be identified through advanced signature and protocol analysis, heuristics or behavioral analysis may be used to determine the identity of the application.

When the application is identified, the policy check determines how to treat the application, for example—block, or allow and scan for threats, inspect for unauthorized file transfer and data patterns, or shape using QoS.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Streamlined App-ID Policy Rules

Safely enable a broad set of applications with common attributes using a single policy rule (for example, give your users broad access to web-based applications or safely enable all enterprise VoIP applications). Palo Alto Networks takes on the task of researching applications with common attributes and delivers this through tags in dynamic content updates. This:

- Minimizes errors and saves time.
- Helps you to create policies that automatically update to handle newly released applications.
- Simplifies the transition toward an App-ID based rule set using Policy Optimizer.

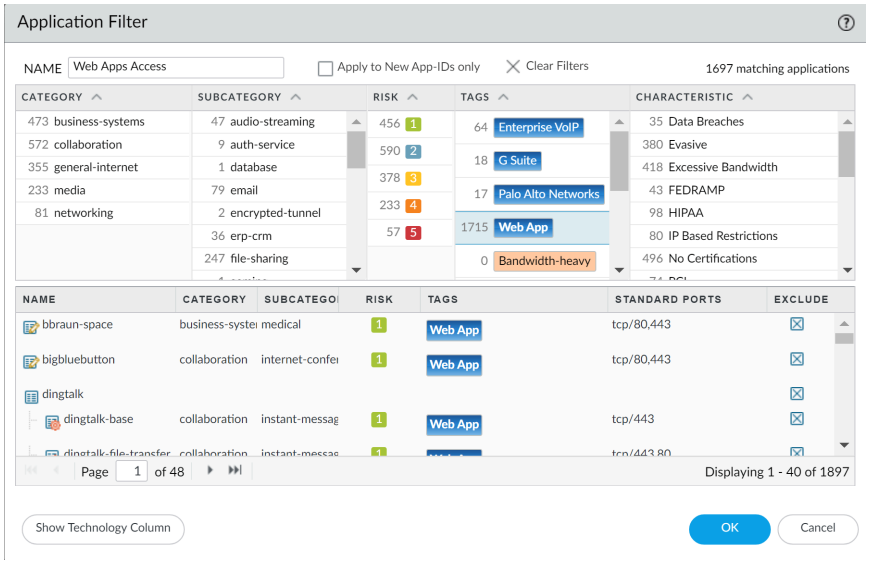
Your firewall can then use your tag-based application filter to dynamically enforce new and updated App-IDs without requiring you to review or update policy rules whenever new applications are added. If you choose to exclude applications from a specific tag, new content updates honor those exclusions. You can also use your own tags to define applications types based on your policy requirements.

- Create an Application Filter Using Tags
- Create an Application Filter Based on Custom Tags

Create an Application Filter Using Tags

STEP 1 | Create an application filter using one or more tags.

If you select more than one tag, applications must match both tags to be included in the filter.



STEP 2 | (Optional) Exclude tags from your filter by selecting the check box in the **Exclude** column.

STEP 3 | Create a security policy rule and **Add** your new application filter on the **Application** tab.

STEP 4 | **Commit** your changes.

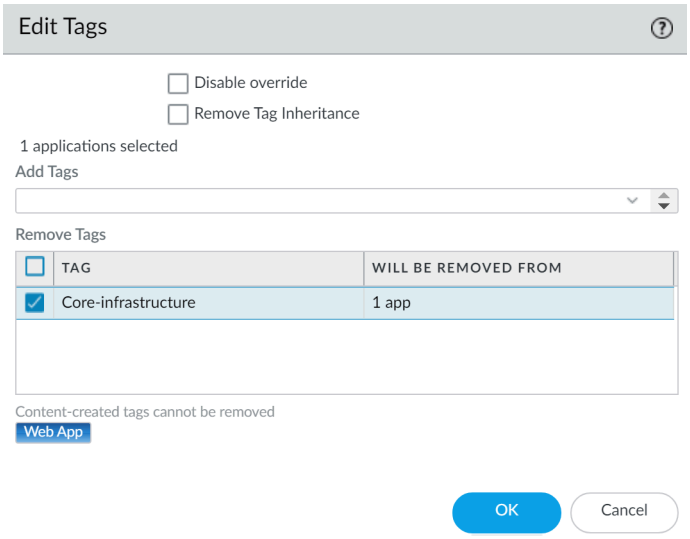
 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Create an Application Filter Based on Custom Tags

STEP 1 | Create a custom tag and apply to App-IDs.

1. (Optional) Remove tags from an application.
2. Filter or search for applications, then select the specific applications to remove tags.
3. **Edit Tags** and select the tags to remove.

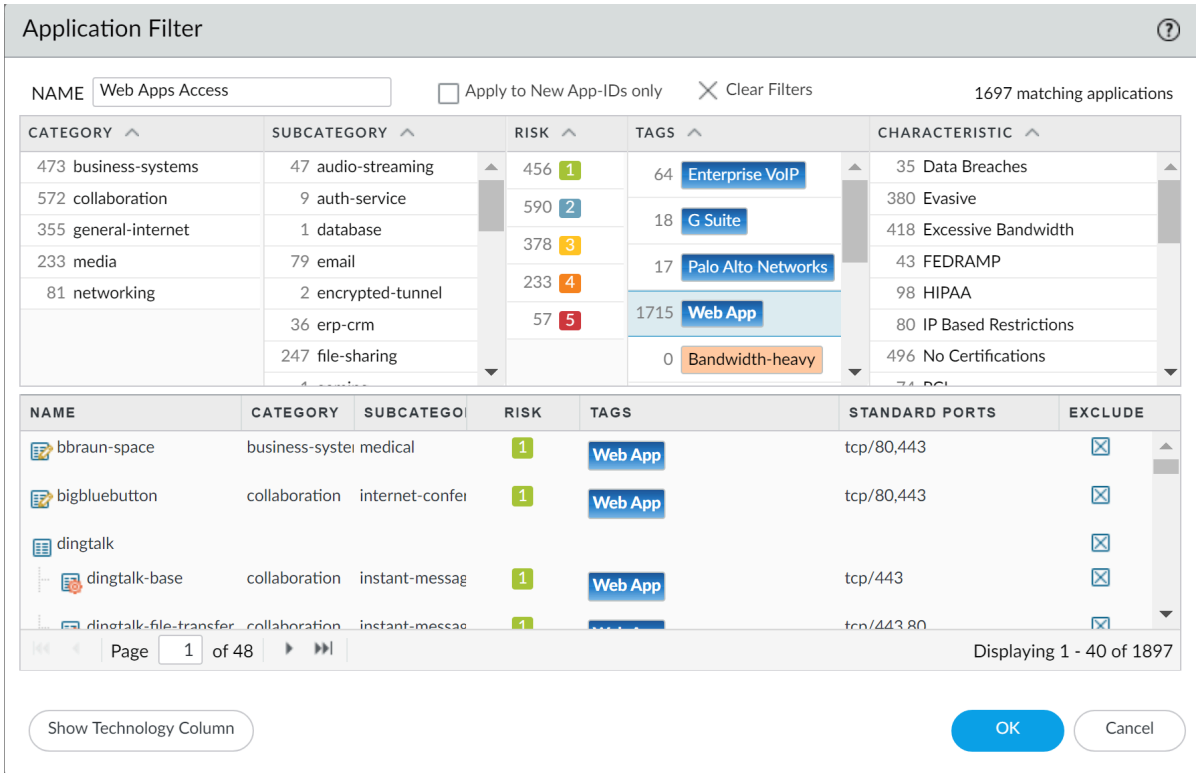


4. Click **OK**.

https://t.me/learningnets

STEP 2 | Create an application filter using one or more tags.

If you select more than one tag, applications must match both tags to be included in the filter.



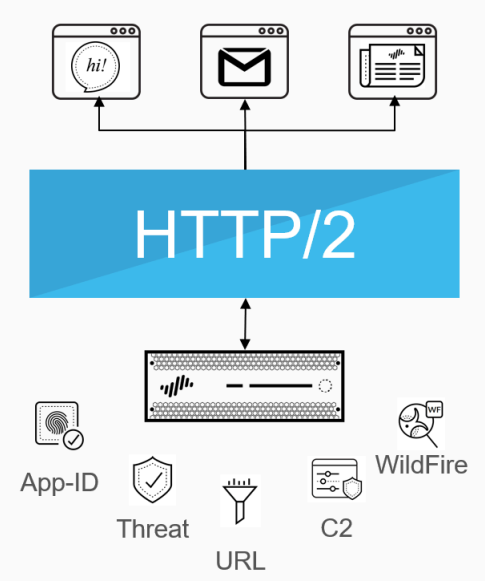
STEP 3 | Create a security policy rule and **Add** your new application filter on the **Application** tab.

STEP 4 | **Commit** your changes.

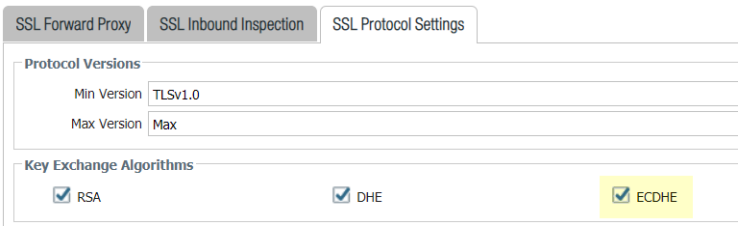
https://t.me/learningnets

App-ID and HTTP/2 Inspection

You can now safely enable applications running over HTTP/2, without any additional configuration on the firewall. As more websites continue to adopt HTTP/2, the firewall can enforce security policy and all threat detection and prevention capabilities on a stream-by-stream basis. This visibility into HTTP/2 traffic enables you to secure web servers that provide services over HTTP/2, and allow your users to benefit from the speed and resource efficiency gains that HTTP/2 provides.



The firewall processes and inspects HTTP/2 traffic by default when SSL decryption is enabled. For HTTP/2 inspection to work correctly, the firewall must be enabled to use ECDHE (elliptic curve Diffie-Hellman) as a key exchange algorithm for SSL sessions. ECDHE is enabled by default, but you can check to confirm that it's enabled by selecting **Objects** > **Decryption** > **Decryption Profile** > **SSL Decryption** > **SSL Protocol Settings**.



When the Decryption logs introduced in PAN-OS 10.1 are enabled, you must enable Tunnel Content Inspection *to obtain the App-ID for HTTP/2 traffic.*

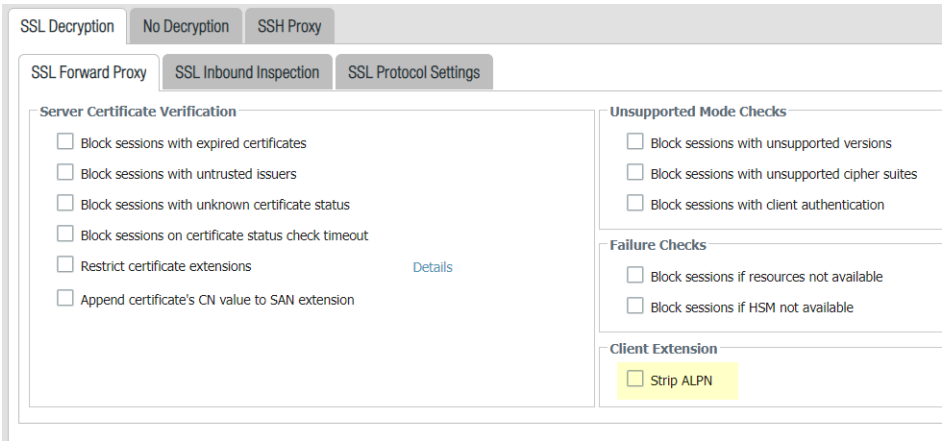
You can disable HTTP/2 inspection for targeted traffic, or globally:

- Disable HTTP/2 inspection for targeted traffic.

  You'll need to specify for the firewall to remove any value contained in the Application-Layer Protocol Negotiation (ALPN) TLS extension. ALPN is used to secure HTTP/2 connections—

https://t.me/learningnets

when there is no value specified for this TLS extension, the firewall either downgrades HTTP/2 traffic to HTTP/1.1 or classifies it as unknown TCP traffic.



1. Select **Objects** > **Decryption** > **Decryption Profile** > **SSL Decryption** > **SSL Forward Proxy** and then select **Strip ALPN**.
2. Attach the decryption profile to a decryption policy (**Policies** > **Decryption**) to turn off HTTP/2 inspection for traffic that matches the policy.
3. **Commit** your changes.

- Disable HTTP/2 inspection globally.

  Use the CLI command: `set deviceconfig setting http2 enable no` and **Commit** your changes. The firewall will classify HTTP/2 traffic as unknown TCP traffic.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Manage Custom or Unknown Applications

Palo Alto Networks provides weekly application updates to identify new App-ID signatures. By default, App-ID is always enabled on the firewall, and you don't need to enable a series of signatures to identify well-known applications. Typically, the only applications that are classified as unknown traffic—tcp, udp or non-syn-tcp—in the ACC and the traffic logs are commercially available applications that have not yet been added to App-ID, internal or custom applications on your network, or potential threats.

On occasion, the firewall may report an application as unknown for the following reasons:

- Incomplete data—A handshake took place, but no data packets were sent prior to the timeout.
- Insufficient data—A handshake took place followed by one or more data packets; however, not enough data packets were exchanged to identify the application.

The following choices are available to handle unknown applications:

- Create security policies to control unknown applications by unknown TCP, unknown UDP or by a combination of source zone, destination zone, and IP addresses.
- Request an App-ID from Palo Alto Networks—If you would like to inspect and control the applications that traverse your network, for any unknown traffic, you can record a packet capture. If the packet capture reveals that the application is a commercial application, you can submit this packet capture to Palo Alto Networks for App-ID development. If it is an internal application, you can create a custom App-ID and/or define an application override policy.
- Create a Custom Application with a signature and attach it to a security policy, or create a custom application and define an application override policy—A custom application allows you to customize the definition of the internal application—its characteristics, category and sub-category, risk, port, timeout—and exercise granular policy control in order to minimize the range of unidentified traffic on your network. Creating a custom application also allows you to correctly identify the application in the **ACC** and traffic logs and is useful in auditing/reporting on the applications on your network. For a custom application you can specify a signature and a pattern that uniquely identifies the application and attach it to a security policy that allows or denies the application.

  Alternatively, if you would like the firewall to process the custom application using fast path (Layer-4 inspection instead of using App-ID for Layer-7 inspection), you can reference the custom application in an application override policy rule. An application override with a custom application will prevent the session from being processed by the App-ID engine, which is a Layer-7 inspection. Instead it forces the firewall to handle the session as a regular stateful inspection firewall at Layer-4, and thereby saves application processing time.

  For example, if you build a custom application that triggers on a host header *www.mywebsite.com*, the packets are first identified as *web-browsing* and then are matched as your custom application (whose parent application is web-browsing). Because the parent application is web-browsing, the custom application is inspected at Layer-7 and scanned for content and vulnerabilities.

  If you define an application override, the firewall stops processing at Layer-4. The custom application name is assigned to the session to help identify it in the logs, and the traffic is not scanned for threats.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Manage New and Modified App-IDs

New and modified App-IDs are delivered to the firewall as part of Applications and Threats Content Updates. While new and modified App-IDs enable the firewall to enforce your security policy with ever-increasing precision, changes in security policy enforcement that can occur when a content update release is installed can impact application availability. For this reason, you will need to think about how to best deploy content updates so that you can get the latest threat prevention as it's made available, and adjust your security policy to best leverage new and modified App-IDs.

The following options enable you to assess the impact of new App-IDs on existing policy enforcement, disable (and enable) App-IDs, and seamlessly update policy rules to secure and enforce newly-identified applications:

- Workflow to Best Incorporate New and Modified App-IDs
- See the New and Modified App-IDs in a Content Release
- See How New and Modified App-IDs Impact Your Security Policy
- Ensure Critical New App-IDs are Allowed
- Monitor New App-IDs
- Disable and Enable App-IDs

You can also take advantage of the Streamlined App-ID Policy Rules that use application tags provided in the content updates.

Workflow to Best Incorporate New and Modified App-IDs

Refer to this master workflow to first set up Application and Threat content updates, and then to best incorporate new and modified App-IDs into your security policy. Everything you need to deploy content updates is referenced here.

STEP 1 | Align your business needs with an approach to deploying Application and Threat content updates.

Learn how Applications and Threat Content Updates work, and identify your organization as either mission-critical or security-first. Understanding which of these is most important to your business will help you to decide how to best deploy content updates and apply best practices to meet your business needs. You might find that you want to apply a mix of both approaches, perhaps depending on firewall deployment (data center or perimeter) or office location (remote or headquarters).

STEP 2 | Review and apply the Best Practices for Applications and Threats Content Updates based on your organization's network security and application availability requirements.

STEP 3 | Configure a security policy rule to always allow new App-IDs that might have network-wide impact, like authentication or software development applications.

The New App-ID characteristic matches to only the App-IDs introduced in the latest content release. When used in a security policy, this gives you a month's time to fine tune your security policy based on new App-IDs while ensuring constant availability for App-IDs that fall into critical categories (Ensure Critical New App-IDs are Allowed).

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 4 | Set the schedule to Deploy Application and Threat Content Updates; this includes the option to delay new App-ID installation until you've had time to make necessary security policy updates (using the **New App-ID Threshold**).

STEP 5 | After you've setup a content updates installation schedule, you'll want to regularly check in and See the New and Modified App-IDs in a Content Release.

STEP 6 | You can then See How New and Modified App-IDs Impact Your Security Policy, and make adjustments to your security policy as needed.

STEP 7 | Monitor New App-IDs to get a view into new App-ID activity on your network, so that you're best equipped to make the most effective security policy updates.

See the New and Modified App-IDs in a Content Release

For both downloaded and installed content updates, you can see a list of the new and modified App-IDs the update includes. Full application details are provided, and importantly, updates to applications with network-wide impact (for example, LDAP or IKE) are prominently flagged as a recommended for policy review. For modified App-IDs, application details also describe how coverage is either now expanded or more precise.

STEP 1 | Select **Device** > **Dynamic Updates** and select **Check Now** to refresh the list of available content updates.

STEP 2 | For either a downloaded or currently installed content release, click **Review Apps** link in the **Actions** column to view details on newly-identified and modified applications in that release:

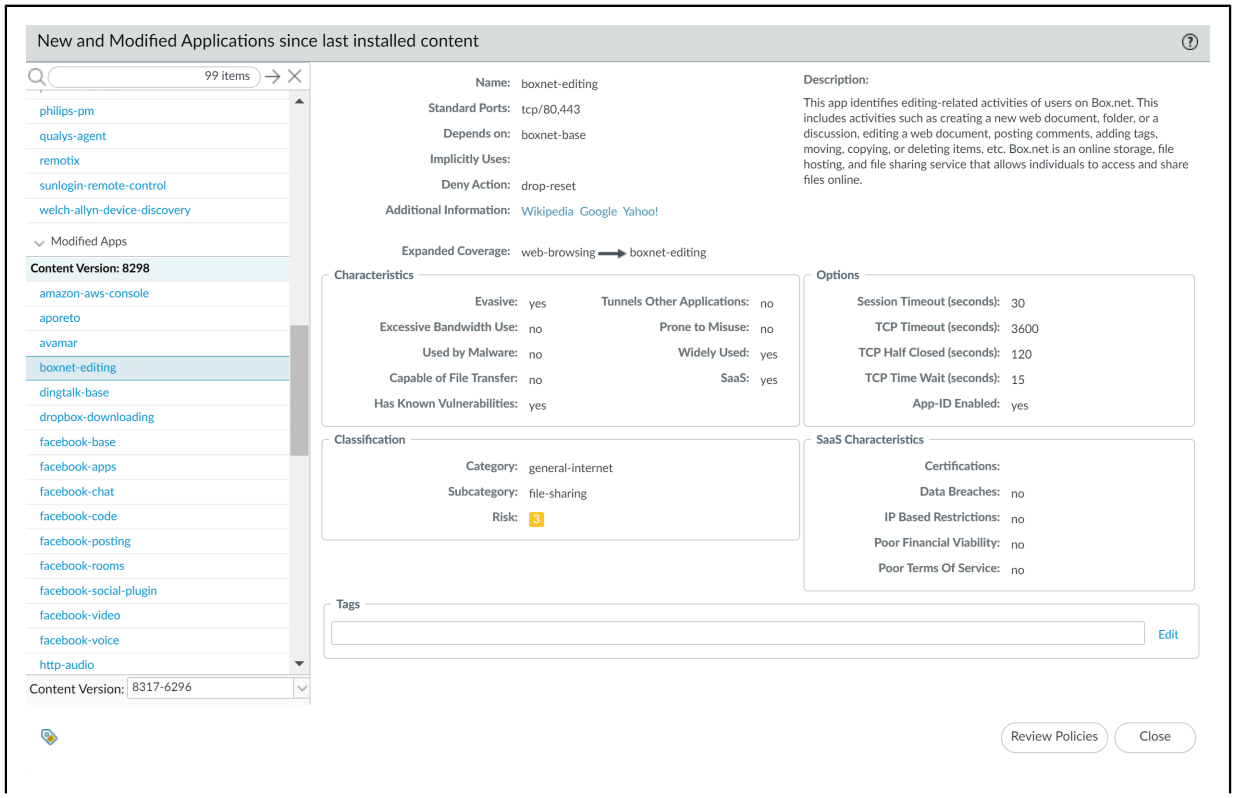
Applications and Threats **Last checked:** 2020/09/23 01:02:02 PDT **Schedule:** Every Wednesday at 01:02 (Download only)

8292-6181	panupv2-all-apps-8292-6181	Apps	Full	47 MB		2020/07/13 11:46:39 PDT	✓ previously		Revert	Release Notes
8317-6296	panupv2-all-apps-8317-6296	Apps	Full	48 MB		2020/09/08 17:55:10 PDT		✓	Review Policies Review Apps	Release Notes
8320-6309	panupv2-all-contents-8320-6309	Apps, Threats	Full	56 MB	192cfd8c2ff0058c188d0...	2020/09/14 18:13:54 PDT			Download	Release Notes
8320-6310	panupv2-all-contents-8320-6310	Apps, Threats	Full	57 MB	2436f79a8f02aeef37b82...	2020/09/15 10:19:15 PDT			Download	Release Notes
8321-6311	panupv2-all-contents-8321-6311	Apps, Threats	Full	56 MB	d3ac74a854c08527869cf...	2020/09/15 13:44:29 PDT			Download	Release Notes
8321-6312	panupv2-all-contents-8321-6312	Apps, Threats	Full	57 MB	a4275ee394b5d942c09e...	2020/09/15 14:26:20 PDT			Download	Release Notes

https://t.me/learningnets

STEP 3 | Review the App-IDs this content release introduces or modifies since the last content version.

New and modified App-IDs are listed separately. Full application details are provided for each, and App-IDs that Palo Alto Networks foresees as having network-wide impact are flagged as recommended for policy review.



New App-ID details that you can use to assess possible impact to policy enforcement include:

- **Depends on**—Lists the application signatures that this App-ID relies on to uniquely identify the application. If one of the application signatures listed in the **Depends On** field is disabled, the dependent App-ID is also disabled.
- **Previously Identified As**—Lists the App-IDs that matched to the application before the new App-ID was installed to uniquely identify the application.
- **App-ID Enabled**—All App-IDs display as enabled when a content release is downloaded, unless you choose to manually disable the App-ID signature before installing the content update.

For modified App-IDs, details include information on: **Expanded Coverage**, **Remove False Positive**, and application metadata changes. The Expanded Coverage and Remove False Positive fields both indicate how the application's coverage has changed (it's either more comprehensive or has been narrowed) and a clock icon indicates a metadata change, where certain application details are updated.

STEP 4 | Based on your findings, click **Review Policies** to see how the new and modified App-IDs impact security policy enforcement: See How New and Modified App-IDs Impact Your Security Policy.

https://t.me/learningnets

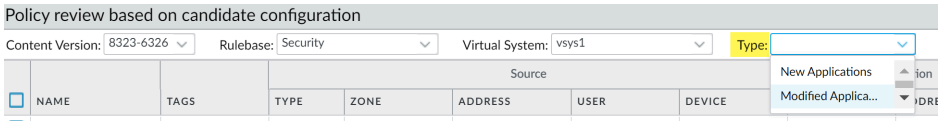
See How New and Modified App-IDs Impact Your Security Policy

Newly-categorized and modified App-IDs can change the way the firewall enforces traffic. Perform a content update policy review to see how new and modified App-IDs impact your security policy, and to easily make any necessary adjustments. You can perform a content update policy review for both downloaded and installed content.

STEP 1 | Select **Device** > **Dynamic Updates**.

STEP 2 | See the New and Modified App-IDs in a Content Release to learn more about each App-ID that a content release introduces or modifies.

STEP 3 | For a downloaded or currently installed content release, click **Review Policies** in the Action column. The **Policy review based on candidate configuration** dialog allows you to filter by **Content Version** and view either new or modified App-IDs introduced in a specific release (you can also filter the policy impact of new App-IDs according to **Rulebase**, **Virtual System**, and **Application**).



STEP 4 | Select an App-ID from the **Application** drop-down to view policy rules that currently enforce the application. The rules displayed are based on the App-IDs that match to the application before the new App-ID is installed (view application details to see the list of application signatures that an application was **Previously Identified As** before the new App-ID).

STEP 5 | Use the detail provided in the policy review to plan policy rule updates to take effect when the App-ID is installed, or if the content release version that included the App-ID is currently installed, the changes you make take effect immediately.

You can **Add app to selected policies** or **Remove app from selected policies**.

Ensure Critical New App-IDs are Allowed

New App-IDs can cause a change in policy enforcement for traffic that is newly-identified as belonging to a certain application. To mitigate any impact to security policy enforcement, you can use the **New App-ID** characteristic in a security policy rule so that the rule always enforces the most recently introduced App-IDs without requiring you to make configuration changes when new App-IDs are installed. The New App-ID characteristic always matches to only the new App-IDs in the most recently installed content releases. When a new content release is installed, the new App-ID characteristic automatically begins to match only to the new App-IDs in that content release version.

You can choose to enforce all new App-IDs, or target the security policy rule to enforce certain types of new App-IDs that might have network-wide or critical impact (for example, enforce only authentication or software development applications). Set the security policy rule to **Allow** to ensure that even if an App-ID release introduces expanded or more precise coverage for critical applications, the firewall continues to allow them.

New App-IDs are released monthly, so a policy rule that allows the latest App-IDs gives you a month's time (or, if the firewall is not installing content updates on a schedule, until the next time

 ©2022 Palo Alto Networks, Inc.

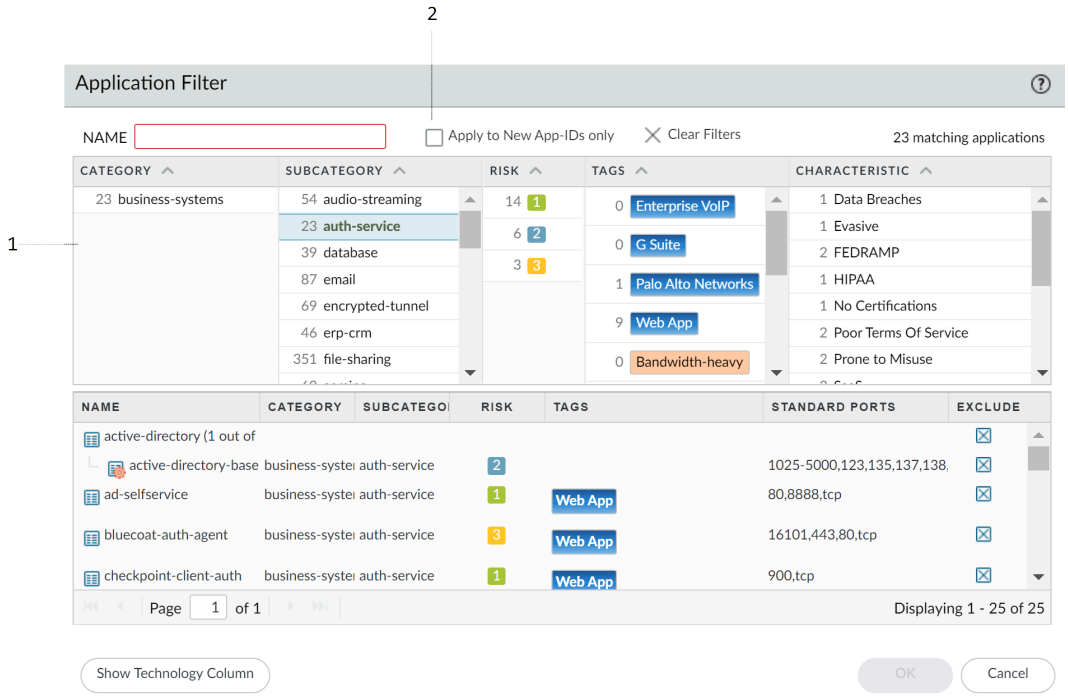
https://t.me/learningnets

you manually install content) to assess how newly-categorized applications might impact security policy enforcement and make any necessary adjustments.

STEP 1 | Select **Objects** > **Application Filters** and **Add** a new application filter.

STEP 2 | Define the types of new applications for which you want to ensure constant availability based on subcategory or characteristic. For example, select the category “auth-service” to ensure that any newly-installed applications that are known to perform or support authentication are allowed.

STEP 3 | Only after narrowing the types of new applications that you want to allow immediately upon installation, select **Apply to New App-IDs only**.



STEP 4 | Select **Policies** > **Security** and add or edit a security policy rule that is configured to allow matching traffic.

STEP 5 | Select **Application** and add the new **Application Filter** to the policy rule as match criteria.

STEP 6 | Click **OK** and **Commit** to save your changes.

STEP 7 | To continue to adjust your security policy to account for any changes to enforcement that new App-IDs introduce:

- Monitor New App-IDs—Monitor and get reports on new App-ID activity.
- See the New and Modified App-IDs in a Content Release—See how the newly-installed App-IDs impact your existing security policy rules.

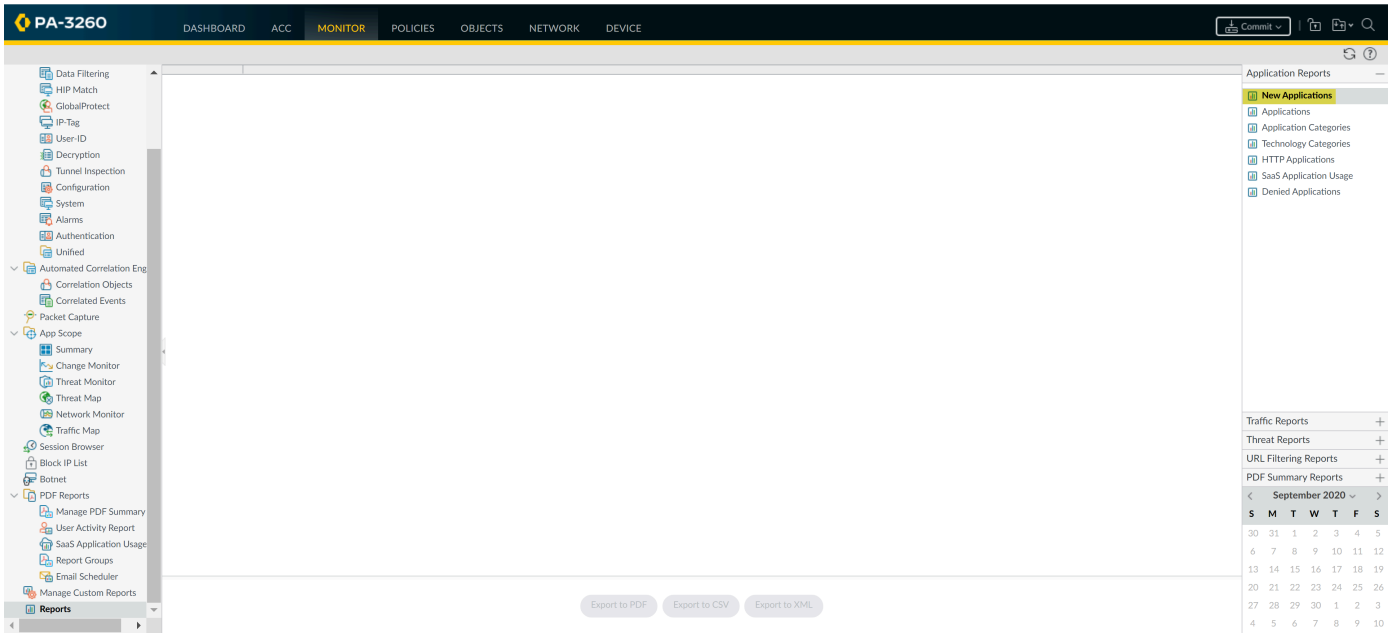
Monitor New App-IDs

The **New App-ID** characteristic enables you to monitor new applications on your network, so that you can better assess the security policy updates you might want to make. Use the New

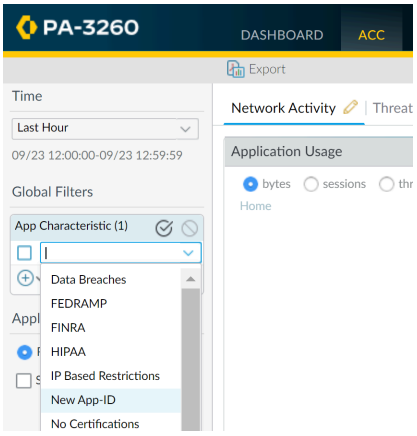
https://t.me/learningnets

App-ID characteristic on the ACC to get visibility into the new applications on your network, and to generate reports that detail newly-categorized application activity. What you learn can help you make the right decisions about how you to update your security policy to enforce the most recently-categorized App-IDs. Whether you're using it on the ACC or to generate reports (or to Ensure Critical New App-IDs are Allowed), the New App-ID characteristic always matches to only the new App-IDs in the most recently installed content releases. When a new content release is installed, the new App-ID characteristic automatically begins to match only to the new App-IDs in that content release version.

- Generate a report with details specifically regarding new applications (applications introduced only in the latest content release).



- Use the ACC to monitor new application activity: select **ACC** and under **Global Filters**, select **Application** > **Application Characteristics** > **New App-ID**.



https://t.me/learningnets

Disable and Enable App-IDs

You can disable all App-IDs introduced in a content release if you want to immediately benefit from the latest threat prevention, and plan to enable the App-IDs later, and you can disable App-IDs for specific applications.

Policy rules referencing App-IDs only match to and enforce traffic based on enabled App-IDs.

Certain App-IDs cannot be disabled and only allow a status of enabled. App-IDs that cannot be disabled include application signatures that are implicitly used by other App-IDs (such as unknown-tcp). Disabling a base App-ID could cause App-IDs which depend on the base App-ID to also be disabled. For example, disabling facebook-base will disable all other Facebook App-IDs.

- Disable all App-IDs in a content release or for scheduled content updates.

  While this option allows you to be protected against threats, by giving you the option to enable the App-ID at a later time, Palo Alto Networks recommends that instead of disabling App-IDs on a regular basis, you should instead configure a security policy rule to Temporarily Allow New App-IDs. This rule will always allow the new App-IDs introduced in only the latest content release. Because content updates that include new App-IDs are released only once a month, this gives you time to assess the new App-IDs and adjust your security policy to cover the new App-IDs if needed, all the while ensuring that availability for critical applications is not affected.

  - To disable all new App-IDs introduced in a content release, select **Device** > **Dynamic Updates** and **Install** an Application and Threats content release. When prompted, select **Disable new apps in content update**. Select the check box to disable apps and continue installing the content update.
  - On the **Device** > **Dynamic Updates** page, select **Schedule**. Choose to **Disable new apps in content update** for downloads and installations of content releases.

- Disable App-IDs for one application or multiple applications at a single time.

  - To quickly disable a single application or multiple applications at the same time, click **Objects** > **Applications**. Select one or more application check box and click **Disable**.
  - To review details for a single application, and then disable the App-ID for that application, select **Objects** > **Applications** and **Disable App-ID**. You can use this step to disable both pending App-IDs (where the content release including the App-ID is downloaded to the firewall but not installed) or installed App-IDs.

- Enable App-IDs.

  Enable App-IDs that you previously disabled by selecting **Objects** > **Applications**. Select one or more application check box and click **Enable** or open the details for a specific application and click **Enable App-ID**.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Use Application Objects in Policy

Use application objects to define how your security policy handles applications.

- Create an Application Group
- Create an Application Filter
- Create a Custom Application
- Resolve Application Dependencies

Create an Application Group

An application group is an object that contains applications that you want to treat similarly in policy. Application groups are useful for enabling access to applications that you explicitly sanction for use within your organization. Grouping sanctioned applications simplifies administration of your rulebases. Instead of having to update individual policy rules when there is a change in the applications you support, you can update only the affected application groups.

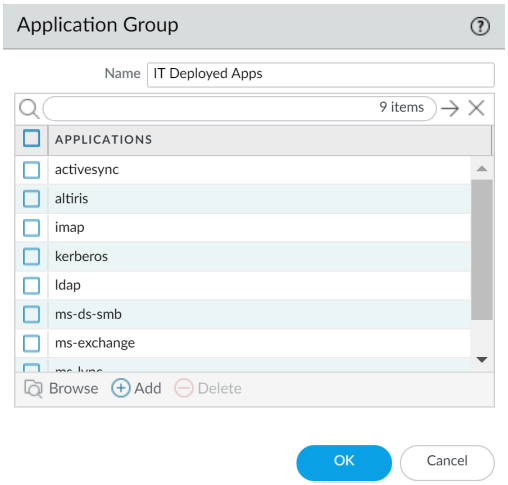
When deciding how to group applications, consider how you plan to enforce access to your sanctioned applications and create an application group that aligns with each of your policy goals. For example, you might have some applications that you will only allow your IT administrators to access, and other applications that you want to make available for any known user in your organization. In this case, you would create separate application groups for each of these policy goals. Although you generally want to enable access to applications on the default port only, you may want to group applications that are an exception to this and enforce access to those applications in a separate rule.

STEP 1 | Select **Objects** > **Application Groups**.

STEP 2 | **Add** a group and give it a descriptive **Name**.

STEP 3 | (Optional) Select **Shared** to create the object in a shared location for access as a shared object in Panorama or for use across all virtual systems in a multiple virtual system firewall.

STEP 4 | **Add** the applications you want in the group and then click **OK**.



 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 5 | **Commit** the configuration.

Create an Application Filter

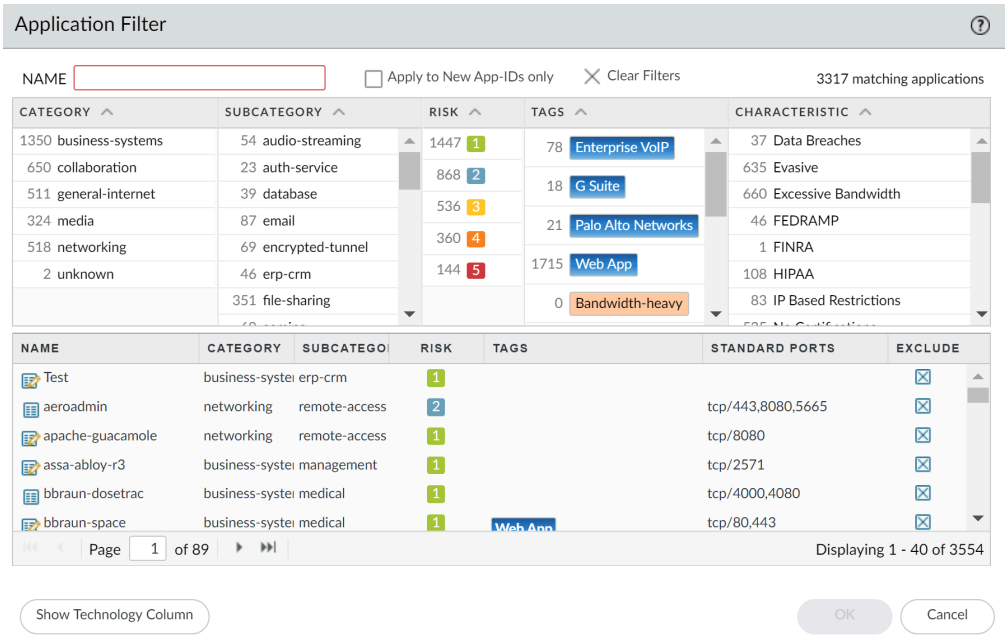
An application filter is an object that dynamically groups applications based on application attributes that you define, including category, subcategory, technology, risk factor, and characteristic. This is useful when you want to safely enable access to applications that you do not explicitly sanction, but that you want users to be able to access. For example, you may want to enable employees to choose their own office programs (such as Evernote, Google Docs, or Microsoft Office 365) for business use. To safely enable these types of applications, you could create an application filter that matches on the Category **business-systems** and the Subcategory **office-programs**. As new applications office programs emerge and new App-IDs get created, these new applications will automatically match the filter you defined; you will not have to make any additional changes to your policy rulebase to safely enable any application that matches the attributes you defined for the filter.

STEP 1 | Select **Objects** > **Application Filters**.

STEP 2 | **Add** a filter and give it a descriptive **Name**.

STEP 3 | (Optional) Select **Shared** to create the object in a shared location for access as a shared object in Panorama or for use across all virtual systems in a multiple virtual system firewall.

STEP 4 | Define the filter by selecting attribute values from the Category, Subcategory, Technology, Risk, and Characteristic sections. As you select values, notice that the list of matching applications at the bottom of the dialog narrows. When you have adjusted the filter attributes to match the types of applications you want to safely enable, click **OK**.



STEP 5 | **Commit** the configuration.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Create a Custom Application

To safely enable applications you must classify all traffic, across all ports, all the time. With App-ID, the only applications that are typically classified as unknown traffic—tcp, udp or non-syn-tcp—in the ACC and the Traffic logs are commercially available applications that have not yet been added to App-ID, internal or custom applications on your network, or potential threats.



If you are seeing unknown traffic for a commercial application that does not yet have an App-ID, you can submit a request for a new App-ID here: http://researchcenter.paloaltonetworks.com/submit-an-application/.

To ensure that your internal custom applications do not show up as unknown traffic, create a custom application. You can then exercise granular policy control over these applications in order to minimize the range of unidentified traffic on your network, thereby reducing the attack surface. Creating a custom application also allows you to correctly identify the application in the ACC and Traffic logs, which enables you to audit/report on the applications on your network.

To create a custom application, you must define the application attributes: its characteristics, category and sub-category, risk, port, timeout. In addition, you must define patterns or values that the firewall can use to match to the traffic flows themselves (the *signature*). Finally, you can attach the custom application to a security policy that allows or denies the application (or add it to an application group or match it to an application filter). You can also create custom applications to identify ephemeral applications with topical interest, such as ESPN3-Video for world cup soccer or March Madness.



In order to collect the right data to create a custom application signature, you'll need a good understanding of packet captures and how datagrams are formed. If the signature is created too broadly, you might inadvertently include other similar traffic; if it is defined too narrowly, the traffic will evade detection if it does not strictly match the pattern.

Custom applications are stored in a separate database on the firewall and this database is not impacted by the weekly App-ID updates.

The supported application protocol decoders that enable the firewall to detect applications that may be tunneling inside of the protocol include the following as of content release version 609: FTP, HTTP, IMAP, POP3, SMB, and SMTP.

The following is a basic example of how to create a custom application.

STEP 1 | Gather information about the application that you will be able to use to write custom signatures.

To do this, you must have an understanding of the application and how you want to control access to it. For example, you may want to limit what operations users can perform within the application (such as uploading, downloading, or live streaming). Or you may want to allow the application, but enforce QoS policing.

- Capture application packets so that you can find unique characteristics about the application on which to base your custom application signature. One way to do this is to run a protocol analyzer, such as Wireshark, on the client system to capture the packets between the client and the server. Perform different actions in the application, such as uploading and

 ©2022 Palo Alto Networks, Inc.

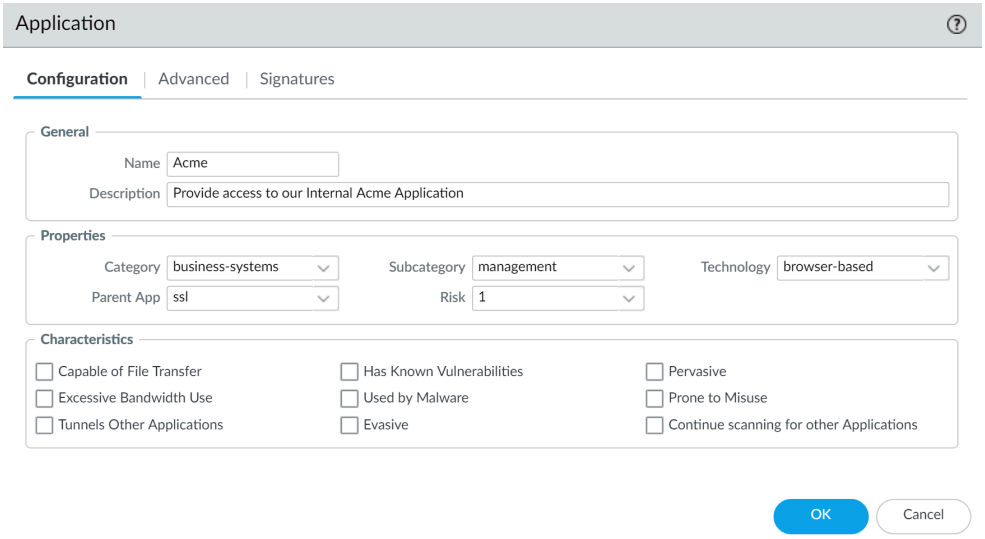
https://t.me/learningnets

downloading, so that you will be able to locate each type of session in the resulting packet captures (PCAPs).

- Because the firewall by default takes packet captures for all unknown traffic, if the firewall is between the client and the server you can view the packet capture for the unknown traffic directly from the Traffic log.
- Use the packet captures to find patterns or values in the packet *contexts* that you can use to create signatures that will uniquely match the application traffic. For example, look for string patterns in HTTP response or request headers, URI paths, or hostnames. For information on the different string contexts you can use to create application signatures and where you can find the corresponding values in the packet, refer to Creating Custom Threat Signatures.

STEP 2 | Add the custom application.

1. Select **Objects** > **Applications** and click **Add**.
2. On the **Configuration** tab, enter a **Name** and a **Description** for the custom application that will help other administrators understand why you created the application.
3. (Optional) Select **Shared** to create the object in a shared location for access as a shared object in Panorama or for use across all virtual systems in a multiple virtual system firewall.
4. Define the application Properties and Characteristics.



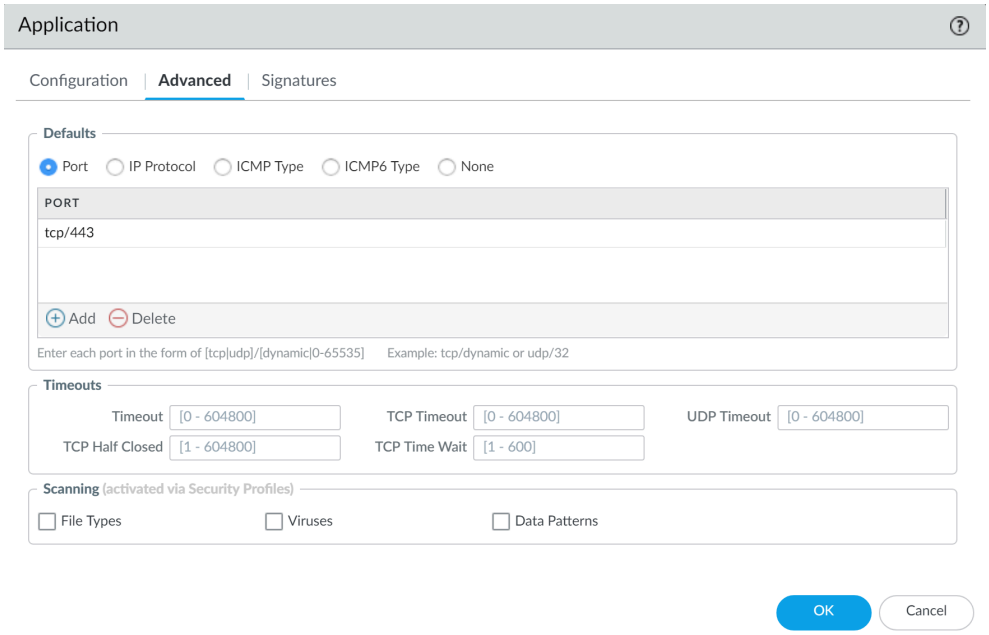
https://t.me/learningnets

STEP 3 | Define details about the application, such as the underlying protocol, the port number the application runs on, the timeout values, and any types of scanning you want to be able to perform on the traffic.

On the **Advanced** tab, define settings that will allow the firewall to identify the application protocol:

- Specify the default ports or protocol that the application uses.
- Specify the session timeout values. If you don't specify timeout values, the default timeout values will be used.
- Indicate any type of additional scanning you plan to perform on the application traffic.

For example, to create a custom TCP-based application that runs over SSL, but uses port 4443 (instead of the default port for SSL, 443), you would specify the port number. By adding the port number for a custom application, you can create policy rules that use the default port for the application rather than opening up additional ports on the firewall. This improves your security posture.



 ©2022 Palo Alto Networks, Inc.

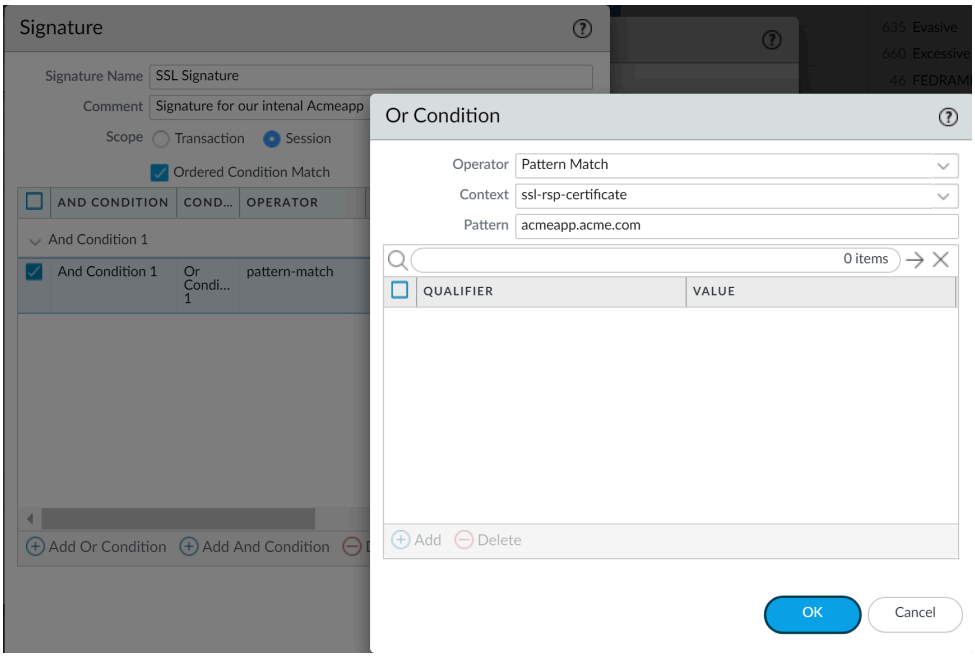
https://t.me/learningnets

STEP 4 | Define the criteria that the firewall will use to match the traffic to the new application.

You will use the information you gathered from the packet captures to specify unique string context values that the firewall can use to match patterns in the application traffic.

1. On the **Signatures** tab, click **Add** and define a **Signature Name** and optionally a **Comment** to provide information about how you intend to use this signature.
2. Specify the **Scope** of the signature: whether it matches to a full **Session** or a single **Transaction**.
3. Specify conditions to define signatures by clicking **Add And Condition** or **Add Or Condition**.
4. Select an **Operator** to define the type of match conditions you will use: **Pattern Match** or **Equal To**.
   - If you selected **Pattern Match**, select the **Context** and then use a regular expression to define the **Pattern** to match the selected context. Optionally, click **Add** to define a qualifier/value pair. The **Qualifier** list is specific to the **Context** you chose.
   - If you selected **Equal To**, select the **Context** and then use a regular expression to define the **Position** of the bytes in the packet header to use match the selected context. Choose from **first-4bytes** or **second-4bytes**. Define the 4-byte hex value for the **Mask** (for example, 0xffffff00) and **Value** (for example, 0xaabbccdd).

   For example, if you are creating a custom application for one of your internal applications, you could use the **ssl-rsp-certificate Context** to define a pattern match for the certificate response message of a SSL negotiation from the server and create a **Pattern** to match the commonName of the server in the message as shown here:



5. Repeat steps 4.c and 4.d for each matching condition.
6. If the order in which the firewall attempts to match the signature definitions is important, make sure the **Ordered Condition Match** check box is selected and then order the conditions so that they are evaluated in the appropriate order. Select a condition or a

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

group and click **Move Up** or **Move Down**. You cannot move conditions from one group to another.

7. Click **OK** to save the signature definition.

STEP 5 | Save the application.

1. Click **OK** to save the custom application definition.
2. Click **Commit**.

STEP 6 | Validate that traffic matches the custom application as expected.

1. Select **Policies** > **Security** and **Add** a security policy rule to allow the new application.
2. Run the application from a client system that is between the firewall and the application and then check the Traffic logs (**Monitor** > **Traffic**) to make sure that you see traffic matching the new application (and that it is being handled per your policy rule).

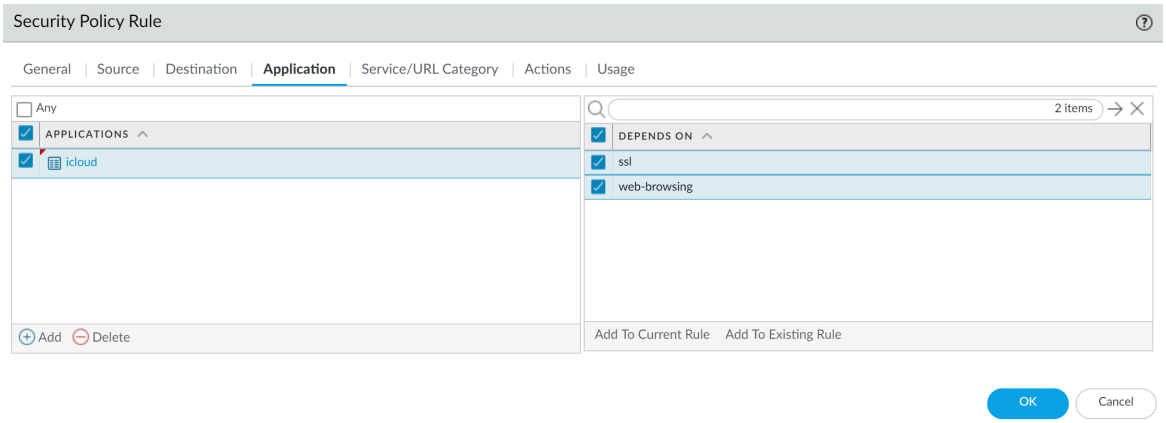
Resolve Application Dependencies

You can see application dependencies when you create a new Security policy rule and when performing Commits. When a policy does not include all application dependencies, you can directly access the associated Security policy rule to add the required applications.

STEP 1 | Create a security policy rule.

STEP 2 | Specify the application that the rule will allow or block.

1. In the **Applications** tab, **Add** the **Application** you want to safely enable. You can select multiple applications or you can use application groups or application filters.
2. View dependencies for selected applications and **Add To Current Rule** or **Add To Existing Rule**.



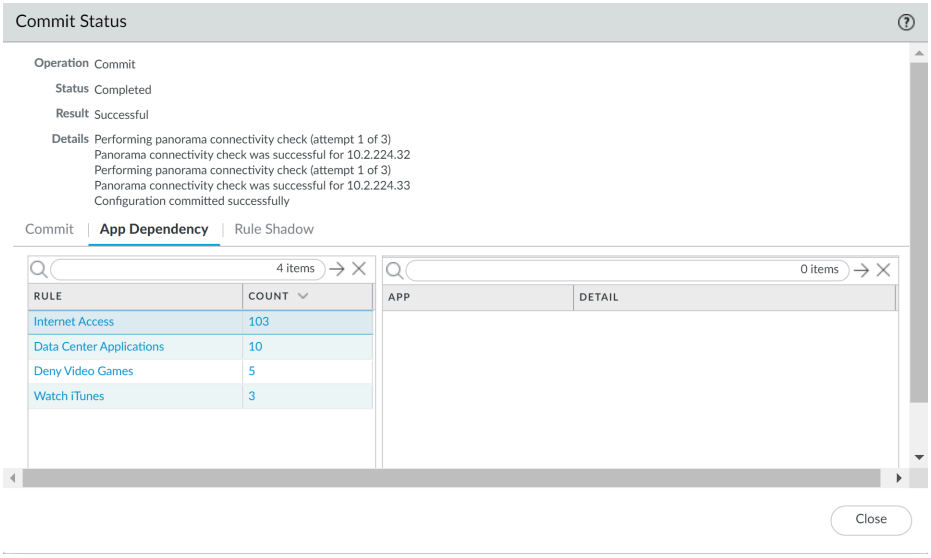
3. If adding to an existing rule, **Select Rule** and click **OK**.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 3 | Click **OK** and **Commit** your changes.

1. Review any Commit warnings in the **App Dependency** tab.



2. Select the **Count** to view the application dependencies not included.
3. Select the **Rule** name to open the policy and add the dependencies.

   Resolve any dependent applications or they'll continue to generate warnings on Commits.

4. Click **OK** and **Commit** your changes.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Safely Enable Applications on Default Ports

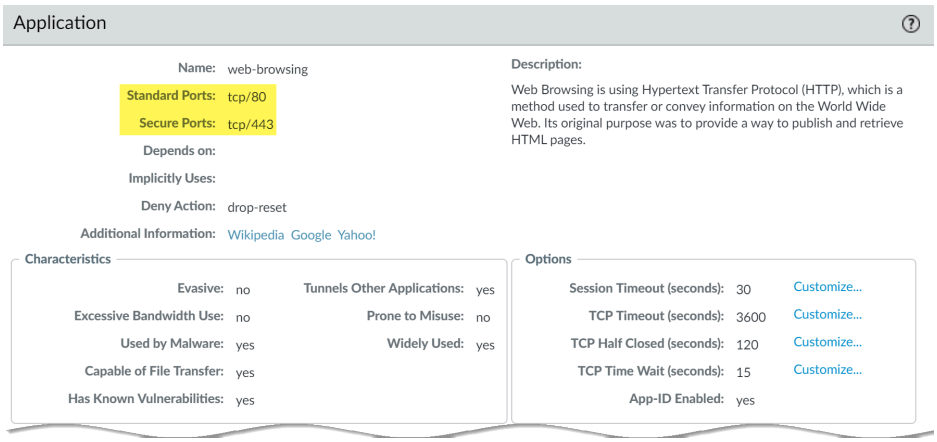
Applications running on unusual ports can indicate an attacker that is attempting to circumvent traditional port-based protections. Application-default is a feature of Palo Alto Networks firewalls that gives you an easy way to prevent this type of evasion and safely enable applications on their most commonly-used ports. Application-default is a best practice for application-based security policies—it reduces administrative overhead, and closes security gaps that port-based policy introduces:

- **Less overhead**—Write simple application-based security policy rules based on your business needs, instead of researching and maintaining application-to-port mappings. We've defined the default ports for all applications with an App-ID.
- **Stronger security**—Enabling applications to run only on their default ports is a security best practice. Application-default helps you to make sure that critical applications are available without compromising security if an application is behaving in an unexpected way.

  Additionally, the default ports an application uses can sometimes depend on whether the application is encrypted or cleartext. Port-based policy requires you to open all the default ports an application might use to account for encryption. Open ports introduce security gaps that an attacker can leverage to bypass your security policy. However, application-default differentiates between encrypted and clear-text application traffic. This means that it can enforce the default port for an application, regardless of whether it is encrypted or not.

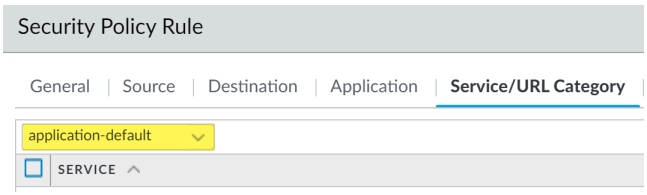
  For example, without application-default, you would need to open ports 80 and 443 to enable web-browsing traffic—you'd be allowing both cleartext and encrypted web-browsing traffic on both ports. With application-default turned on, the firewall strictly enforces cleartext web-browsing traffic only on port 80 and SSL-tunneled traffic only on port 443.

To see the ports that an application uses by default, you can visit Applipedia or select **Objects** > **Applications**. Application details include the application's standard port—the port it most commonly uses when in cleartext. For web-browsing, SMTP, FTP, LDAP, POP3, and IMAP details also include the application's secure port—the port the application uses when encrypted.



Select **Policy** > **Security** and add or a modify a rule to enforce applications only on their default port(s):

https://t.me/learningnets



*Using application-default as part of an application-based security policy and with SSL decryption is a best practice. Additionally, if you have existing security policy rules that control web-browsing traffic with the **Service** set to service-http and service-https, you should update those rules to use application-default instead.*

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Applications with Implicit Support

When creating a policy to allow specific applications, you must also be sure that you are allowing any other applications on which the application depends. In many cases, you do not have to explicitly allow access to the dependent applications for the traffic to flow because the firewall is able to determine the dependencies and allow them implicitly. This implicit support also applies to custom applications that are based on HTTP, SSL, MS-RPC, or RTSP. Applications for which the firewall cannot determine dependent applications on time will require that you explicitly allow the dependent applications when defining your policies. You can determine application dependencies from within your application-based security policy workflow using one of the following:

- Policy Optimizer
- Create an Application Filter Using Tags
- Create an Application Filter Based on Custom Tags
- Resolve Application Dependencies

Applipedia is also available if needed.

The following table lists the applications for which the firewall has implicit support (as of Content Update 595).

Application	Implicitly Supports
360-safeguard-update	http
apple-update	http
apt-get	http
as2	http
avg-update	http
avira-antivir-update	http, ssl
blokus	rtmp
bugzilla	http
clubcooee	http
corba	http
cubby	http, ssl
dropbox	ssl

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Application	Implicitly Supports
esignal	http
evernote	http, ssl
ezhelp	http
facebook	http, ssl
facebook-chat	jabber
facebook-social-plugin	http
fastviewer	http, ssl
forticlient-update	http
good-for-enterprise	http, ssl
google-cloud-print	http, ssl, jabber
google-desktop	http
google-talk	jabber
google-update	http
gotomypc-desktop-sharing	citrix-jedi
gotomypc-file-transfer	citrix-jedi
gotomypc-printing	citrix-jedi
hipchat	http
iheartradio	ssl, http, rtmp
infront	http
instagram	http, ssl
issuu	http, ssl
java-update	http
jepptech-updates	http

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Application	Implicitly Supports
kerberos	rpc
kik	http, ssl
lastpass	http, ssl
logmein	http, ssl
mcafee-update	http
megaupload	http
metatrader	http
mocha-rdp	t_120
mount	rpc
ms-frs	msrpc
ms-rdp	t_120
ms-scheduler	msrpc
ms-service-controller	msrpc
nfs	rpc
oovoo	http, ssl
paloalto-updates	ssl
panos-global-protect	http
panos-web-interface	http
pastebin	http
pastebin-posting	http
pinterest	http, ssl
portmapper	rpc
prezi	http, ssl

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Application	Implicitly Supports
rdp2tcp	t_120
renren-im	jabber
roboform	http, ssl
salesforce	http
stumbleupon	http
supremo	http
symantec-av-update	http
trendmicro	http
trillian	http, ssl
twitter	http
whatsapp	http, ssl
xm-radio	rtsp

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Security Policy Rule Optimization

Policy Optimizer provides a simple workflow to migrate your legacy Security policy rulebase to an App-ID based rulebase, which improves your security by reducing the attack surface and gaining visibility into applications so you can safely enable them. Policy Optimizer identifies port-based rules so you can convert them to application-based allow rules or add applications from a port-based rule to an existing application-based rule without compromising application availability. It also identifies over-provisioned App-ID based rules (App-ID rules configured with unused applications). Policy Optimizer helps you prioritize which port-based rules to migrate first, identify application-based rules that allow applications you don't use, and analyze rule usage characteristics such as hit count.

Converting port-based rules to application-based rules improves your security posture because you select the applications you want to allow and deny all other applications, so you eliminate unwanted and potentially malicious traffic from your network. Combined with restricting application traffic to its default ports (set the Service to **application-default**), converting to application-based rules also prevents evasive applications from running on non-standard ports.

You can use this feature on:

- Firewalls that run PAN-OS version 9.0 and have App-ID enabled.
- Panorama running PAN-OS version 9.0. You don't have to upgrade firewalls that Panorama manages to use the **Policy Optimizer** capabilities. However, to use the **Rule Usage** capabilities (Monitor Policy Rule Usage), managed firewalls must run PAN-OS 8.1 or later. If managed firewalls connect to Log Collectors, those Log Collectors must also run PAN-OS version 9.0. Managed PA-7000 Series firewalls that have a Log Processing Card (LPC) can also run PAN-OS 8.1 (or later).
- For Cortex Data Lake compatibility, Panorama running PAN-OS 10.0.3 or later with the Cloud Services plugin 2.0 Innovation or later installed.



PA-7000 Series Firewalls support two logging cards, the PA-7000 Series Firewall Log Processing Card (LPC) and the high-performance PA-7000 Series Firewall Log Forwarding Card (LFC). Unlike the LPC, the LFC does not have disks to store logs locally. Instead, the LFC forwards all logs to one or more external logging systems, such as Panorama or a syslog server. If you use the LFC, the application usage information for Policy Optimizer does not display on the firewall because traffic logs aren't stored locally. If you use the LPC, the traffic logs are stored locally on the firewall, so the application usage information for Policy Optimizer displays on the firewall.

Use this feature to:

- **Migrate port-based rules to application-based rules**—Instead of combing through traffic logs and manually mapping applications to port-based rules, use Policy Optimizer to identify port-based rules and list the applications that matched each rule, so you can select the applications you want to allow and safely enable them. Converting your legacy port-based rules to application-based allow rules supports your business applications and enables you to block any applications associated with malicious activity.

https://t.me/learningnets

- **Identify over-provisioned application-based rules**—Rules that are too broad allow applications you don't use on your network, which increases the attack surface and the risk of inadvertently allowing malicious traffic.

  Remove unused applications from Security policy rules to reduce the attack surface and keep the rulebase clean. Don't allow applications that nobody uses on your network.

- **Add App-ID Cloud Engine (ACE) applications to Security policy rules**—If you have a SaaS Security Inline subscription, you can use Policy Optimizer's New App Viewer to manage cloud-delivered App-IDs in Security policy. The ACE documentation describes how to use Policy Optimizer to gain visibility into and control cloud-delivered App-IDs.

  The Policy Optimizer examples in this section do not show the New App Viewer because they depict firewalls that do not have a SaaS Security Inline subscription.

To migrate a configuration from a legacy firewall to a Palo Alto Networks device, see Best Practices for Migrating to Application-Based Policy.

You can't sort Security policy rules in **Security** > **Policies** because sorting would change the rule order in the rulebase. However, under **Polices** > **Security** > **Policy Optimizer**, Policy Optimizer provides sorting options that don't affect the rule order, so you can sort rules to prioritize which rules to convert or clean up first. You can sort rules by the amount of traffic during the past 30 days, the number of applications seen on the rule, the number of days with no new applications, and the number of applications allowed (for over-provisioned rules).

You can use Policy Optimizer in other ways as well, including validating pre-production rules and troubleshooting existing rules. Note that Policy Optimizer honors only **Log at Session End** and ignores **Log at Session Start** to avoid counting transient applications on rules.

Due to resource constraints, VM-50 Lite virtual firewalls don't support Policy Optimizer.

- Policy Optimizer Concepts
- Migrate Port-Based to App-ID Based Security Policy Rules
- Rule Cloning Migration Use Case: Web Browsing and SSL Traffic
- Add Applications to an Existing Rule
- Identify Security Policy Rules with Unused Applications
- High Availability for Application Usage Statistics
- How to Disable Policy Optimizer

Policy Optimizer Concepts

Review the following topics to learn more about this feature's support:

- Sorting and Filtering Security Policy Rules
- Clear Application Usage Data

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Sorting and Filtering Security Policy Rules

You can filter Security policy rules to see the port-based rules, which are rules with no applications configured on them (**Policies** > **Security** > **Policy Optimizer** > **No App Specified**). You can also filter to see the rules that have applications configured on them, but traffic only matches some of the configured applications—the rule is over-provisioned and includes applications that aren't seen on the rule (**Policies** > **Security** > **Policy Optimizer** > **Unused Apps**). In addition, if you have a SaaS Security Inline license, you can use the New App Viewer to filter rules that have seen new App-ID Cloud Engine (ACE) applications (see the ACE documentation for how to do this). You can sort the filtered policy rules based on different types of statistics to help prioritize which rules to convert from port-based to application-based rules or to clean up first.

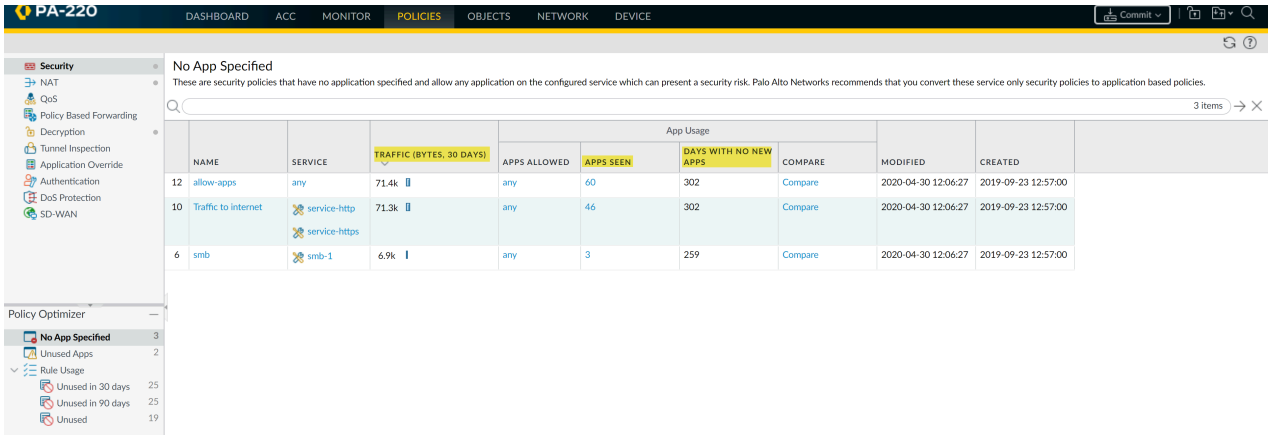


*You can't filter or sort rules in **Policies** > **Security** because that would change the order of the policy rules in the rulebase. Filtering and sorting **Policies** > **Security** > **Policy Optimizer** > **No App Specified**, **Policies** > **Security** > **Policy Optimizer** > **Unused Apps**, and **Policies** > **Security** > **Policy Optimizer** > **New App Viewer** (if you have a SaaS Inline Security subscription) does not change the order of the rules in the rulebase.*

You can click several column headers to sort rules based on application usage statistics. In addition, you can View Policy Rule Usage to help identify and remove unused rules to reduce security risks and keep your policy rule base organized. Rule usage tracking allows you to quickly validate new rule additions and rule changes and to monitor rule usage for operations and troubleshooting tasks.

©2022 Palo Alto Networks, Inc.

https://t.me/learningnets



- **Traffic (Bytes, 30 days)**—The amount of traffic seen on the rule over the last 30 days. The 30-day window places rules that *currently* match the most traffic at the top of the list by default (a longer time frame places more emphasis on older rules that would remain at the top of the list because they have large cumulative totals even though they may no longer see much traffic). Click to reverse the order.
- **Apps Seen**—Place the rules with the most or least applications seen at the top. The firewall never automatically purges the application data.

  *The firewall updates **Apps Seen** approximately every hour. However, if there is a large volume of application traffic or a large number of rules, it may take longer than an hour to update. After you add an application to a rule, wait at least an hour before running Traffic logs to see the application's log information.*
- **Days with No New Apps**—Place the rules with the most or least days since the last new application matched the rule at the top.
- (Unused Apps only) **Apps Allowed**—Place the rules with the most or least applications configured on the rule at the top.

Application usage statistics only count applications for rules that meet the following criteria:

https://t.me/learningnets

- The rule's Action must be **Allow**.
- The rule's Log Setting must be **Log at Session End** (this is the default Log Setting). Rules that **Log at Session Start** are ignored to prevent counting transient applications.
- Valid traffic must match the rule. For example, if the session ends before enough traffic passes through the firewall to identify the application, it is not counted. The following traffic types are not valid and therefore don't count for Policy Optimizer statistics:
  - Insufficient-data
  - Not-applicable
  - Non-syn-tcp
  - Incomplete

  You can filter the Traffic logs (**Monitor** > **Logs** > **Traffic**) to see traffic identified as one of these types. For example, to see all traffic identified as incomplete, use the filter **`(app eq incomplete)`**.

If these criteria aren't met, the application isn't counted for statistics such as **Apps Seen**, doesn't affect statistics such as **Days with No New Apps**, and doesn't appear in lists of applications.

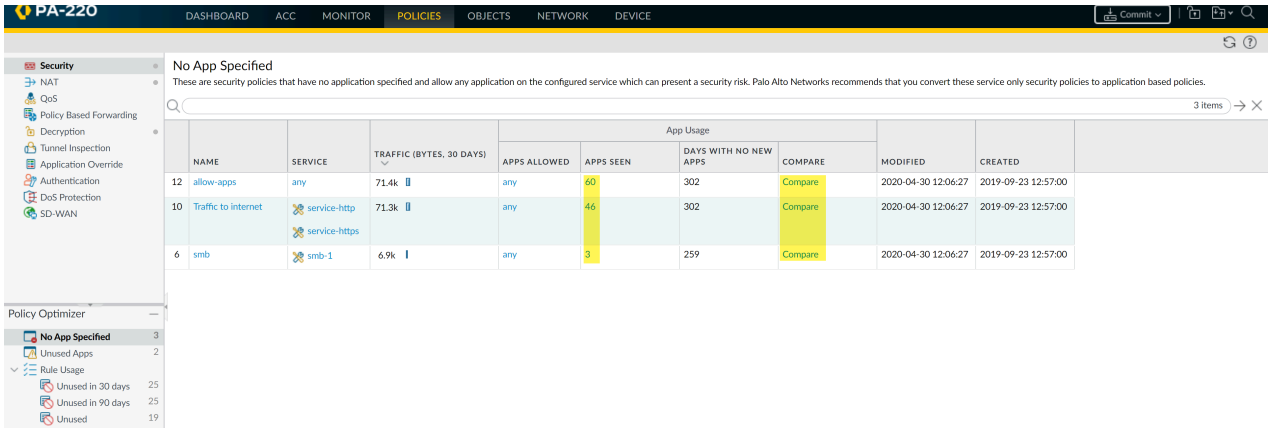
The firewall doesn't track application usage statistics for the interzone-default and intrazone-default Security policy rules.

If the UUID of a rule changes, the application usage statistics for that rule reset because the UUID change makes the firewall see the rule as a different (new) rule.

To see and sort the applications seen on a rule, in the rule's row, click **Compare** or click the number in **Apps Seen**.

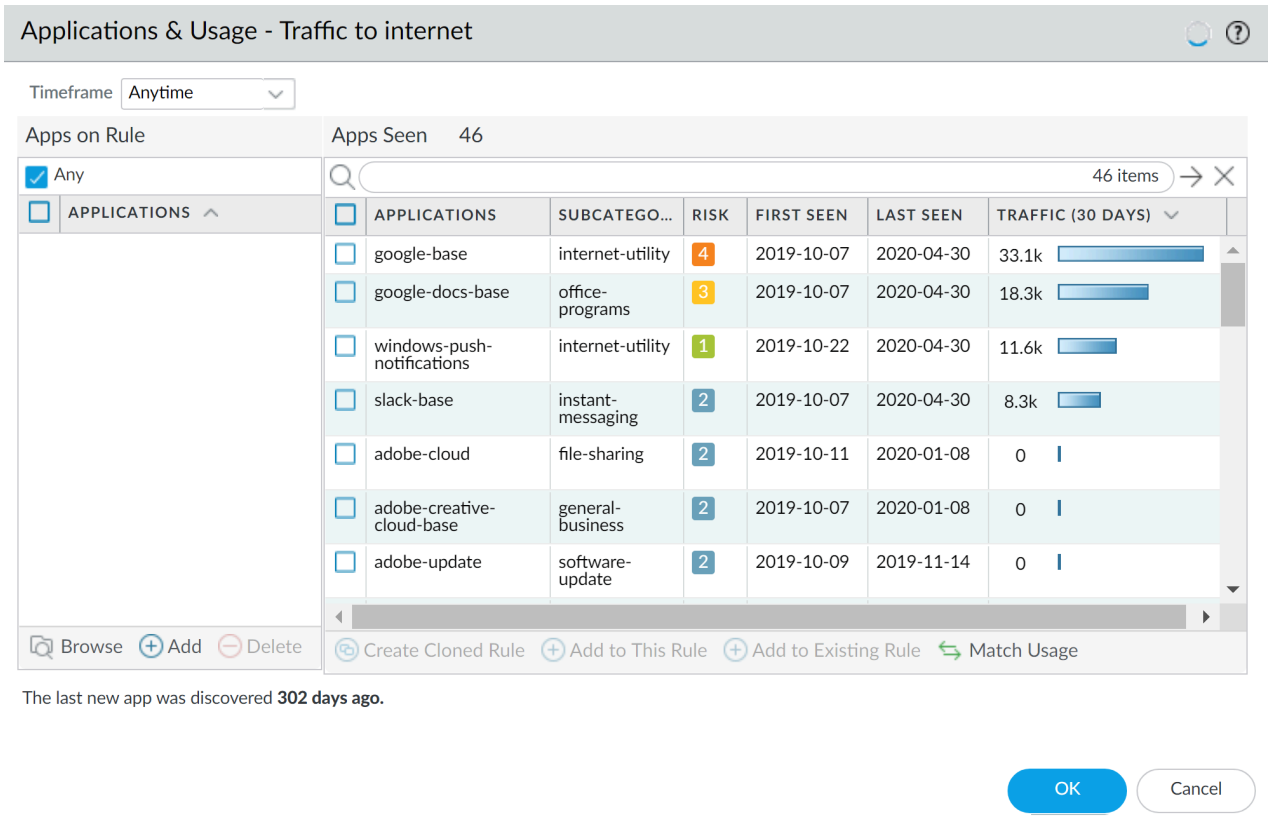
 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets



For the rules you see in **Policies** > **Security** > **Policy Optimizer** > **No App Specified** and **Policies** > **Security** > **Policy Optimizer** > **Unused Apps**, clicking **Compare** or the **Apps Seen** number brings up **Applications & Usage**, which gives you a view of the applications seen on the rule and the ability to sort them. **Applications & Usage** is also where you Migrate Port-Based to App-ID Based Security Policy Rules and remove unused applications from rules.

https://t.me/learningnets



You can sort the applications seen on the rule by all six of the **Apps Seen** statistics (**Apps Seen** is not updated in real time and takes an hour or longer to update, depending on the volume of traffic and number of rules).

- **Applications**—Alphabetical by application name. If you configure specific ports or port ranges for a rule's Service (the Service cannot be **any**), and there are standard (application default) ports for the application, and the configured ports don't match the application-default ports, then a yellow, triangular warning icon appears next to the application.
- **Subcategory**—Alphabetical by application subcategory, derived from the application content metadata.
- **Risk**—According to the risk rating of the application.
- **First Seen**—The first day the application was seen on the rule. The time stamp resolution is by the day only (not hourly).
- **Last Seen**—The last day the application was seen on the rule. The time stamp resolution is by the day only (not hourly).
- **Traffic (30 days)**—Traffic in bytes that matched the rule over the last 30 days is the default sorting method.

https://t.me/learningnets

Set the **Timeframe** to display statistics for a particular time period—**Anytime**, the **Past 7 days**, the **Past 15 days**, or the **Past 30 days**.

Traffic (30 days) *always displays only the last 30 days of traffic in bytes. Changing the* ***Timeframe*** *does not change the duration of the* ***Traffic (30 days)*** *bytes measurement.*

Clicking the column header orders the display and clicking the same column again reverses the order. For example, click **Risk** to sort applications from low risk to high risk. Click **Risk** again to sort applications from high risk to low risk.

The firewall doesn't report application usage statistics in real time for Policy Optimizer, so it isn't a replacement for running reports.

- The firewall updates **Apps Allowed**, **Apps Seen**, and the applications listed in **Applications & Usage** approximately every hour, not in real time. If there is a large amount of traffic or a large number of rules, updates may take longer. After you add an application to a rule, wait at least an hour before running Traffic logs to see the application's log information.

  The firewall updates **Apps Seen** approximately every hour. However, if there is a large volume of application traffic or a large number of rules, it may take longer than an hour to update. After you add an application to a rule, wait at least an hour before running Traffic logs to see the application's log information.
- The firewall updates **Days with No New Apps** and also **First Seen** and **Last Seen** on **Applications & Usage** once per day, at midnight device time.
- For rules with large numbers of applications seen, it may take longer to process application usage statistics.
- For Security policy rulebases with large numbers of rules that have many applications, it may take longer to process application usage statistics.
- For firewalls managed by Panorama, application usage data is visible only for rules Panorama pushes to the firewalls, not for rules configured locally on individual firewalls.

Clear Application Usage Data

You can use a CLI command to clear application usage data for an individual Security policy rule and reset **Apps Seen** and other application usage data.

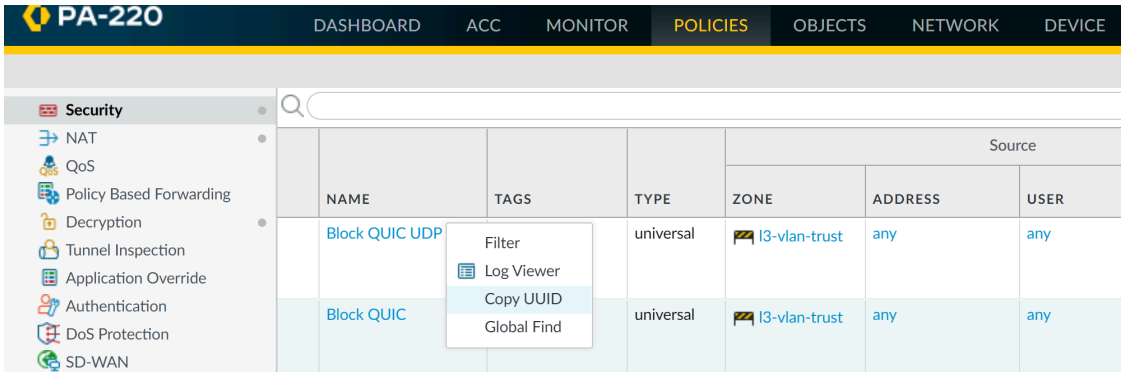
 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 1 | Find the UUID of the Security policy rule whose application usage data you want to clear.

There are two ways to find the UUID in the UI:

- In **Policies** > **Security**, copy the UUID from the **Rule UUID** column.
- In **Policies** > **Security**, select **Copy UUID** in the rule **Name** drop-down menu.



STEP 2 | Switch from the UI to the CLI.

Use the UUID you captured in the UI to clear the rule's application usage data:

```
admin@PA-VM>clear policy-app-usage-data ruleuuid <uuid-value>
```

Paste or type the rule's UUID as the value and execute the command to clear the rule's application usage data.

Migrate Port-Based to App-ID Based Security Policy Rules

When you transition from a legacy firewall to a Palo Alto Networks next-generation firewall, you inherit a large number of port-based rules that allow any application on the ports, which increases the attack surface because any application can use an open port. Policy Optimizer identifies all applications seen on any legacy port-based Security policy rule and provides an easy workflow for selecting the applications you want to allow on that rule. Migrate port-based rules to application-based rules to reduce the attack surface and safely enable applications on your network. Use Policy Optimizer to maintain the rulebase as you add new applications.

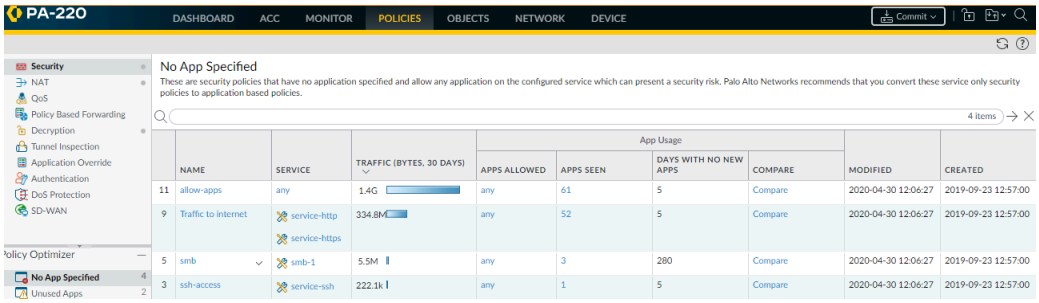
Migrate a few port-based rules at a time to application-based rules, in a prioritized manner. A gradual conversion is safer than migrating a large rulebase at one time and makes it easier to ensure that the new application-based rules control the necessary applications. Use ***Policy Optimizer*** *to prioritize which rules to convert first.*

To migrate a configuration from a legacy firewall to a Palo Alto Networks device, see Best Practices for Migrating to Application-Based Policy.

https://t.me/learningnets

STEP 1 | Identify port-based rules.

Port-based rules have no configured (allowed) applications. **Policies** > **Security** > **Policy Optimizer** > **No App Specified** displays all port-based rules (**Apps Allowed** is **any**).



STEP 2 | Prioritize which port-based rules to convert first.

Policies > **Security** > **Policy Optimizer** > **No App Specified** enables you to sort rules without affecting their order in the rulebase and provides other information that helps you prioritize rules for conversion based on your business goals and risk tolerance.

- **Traffic (Bytes, 30 days)**—(Click to sort.) Rules that *currently* match the most traffic are at the top of the list. This is the default sorting order.
- **Apps Seen**—(Click to sort.) A large number of legitimate applications matching a port-based rule may indicate you should replace it with multiple application-based rules that tightly define the applications, users, and sources and destinations. For example, if a port-based rule controls traffic for multiple applications for different user groups on different sets of devices, create separate rules that pair applications with their legitimate users and devices to reduce the attack surface and increase visibility. (Clicking the **Apps Seen** number or **Compare** shows you the applications that have matched the rule.)

  *The firewall updates **Apps Seen** approximately every hour. However, if there is a large volume of application traffic or a large number of rules, it may take longer than an hour to update. After you add an application to a rule, wait at least an hour before running Traffic logs to see the application's log information.*

- **Days with No New Apps**—(Click to sort.) When the applications seen on a port-based rule stabilize, you can be more confident the rule is mature, conversion won't accidentally exclude legitimate applications, and no more new applications will match the rule. The **Created** and **Modified** dates help you evaluate a rule's stability because older rules that have not been modified recently may also be more stable.
- **Hit Count**—Displays rules with the most matches over a selected time frame. You can exclude rules for which you reset the hit counter and specify the exclusion time period in days. Excluding rules with recently reset hit counters prevents misconceptions about rules that show fewer hits than you expect because you didn't know the counter was reset.

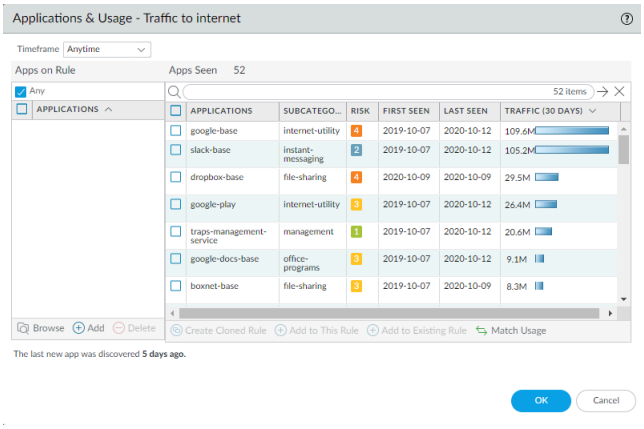
  *You can also use **Hit Count** to View Policy Rule Usage and help identify and remove unused rules to reduce security risks and keep your rulebase organized.*

STEP 3 | Review the **Apps Seen** on port-based rules, starting with the highest priority rules.

On **No Apps Specified**, click **Compare** or the number in **Apps Seen** to open **Applications & Usage**, which lists applications that matched a port-based rule over a specified **Timeframe**,

https://t.me/learningnets

with each application's **Risk**, the date it was **First Seen**, the date it was **Last Seen**, and the amount of traffic over the last 30 days.



You can check **Applications seen** on port-based rules over the past 7, 15, or 30 days, or over the rule's lifetime (**Anytime**). For migrating rules, **Anytime** provides the most complete assessment of applications that matched the rule.

You can search and filter the **Apps Seen**, but keep in mind that it takes an hour or more to update **Apps Seen**. You can also order the **Apps Seen** by clicking the column headers. For example, you can click **Traffic (30 days)** to bring the applications with the most recent traffic to the top of the list, or click **Subcategory** to organize the applications by subcategory.

*The granularity of measurement for **First Seen** and **Last Seen** data is one day, so on the day you define a rule, the dates in these two columns are the same. On the second day the firewall sees traffic on an application, you'll see a difference in the dates.*

https://t.me/learningnets

STEP 4 | Clone or add applications to the rule to specify the applications you want to allow on the rule.

On **Applications & Usage**, convert a port-based rule to an application-based rule in either of two ways:

- **Clone the rule**—Preserves the original port-based rule and places the cloned application-based rule directly above it in the rulebase.
- **Add Applications to the Rule**—Replaces the original port-based rule with the new application-based rule and deletes the original rule.

> *If you have existing application-based rules and you want to migrate applications to them from port-based rules, you can Add Applications to an Existing Rule instead of cloning a new rule or converting the port-based rule by adding applications to it.*

*Some applications appear on the network at intervals, for example, for quarterly or yearly events. These applications may not display on the **Applications & Usage** screen if the history isn't long enough to capture their latest activity.*

*When you clone a rule or add applications to a rule, nothing else about the original rule changes. The original rule's configuration remains the same except for the applications you added to the rule. For example, if the original rule's Service allowed **Any** application or specified a particular service, you need to change the Service to **Application-Default** to restrict the allowed applications to their default ports on the new rule.*

Cloning is the safest way to migrate rules, especially when **Applications & Usage** shows more than a few well-known applications matching the rule (Rule Cloning Migration Use Case: Web Browsing and SSL Traffic provides an example of this). Cloning preserves the original port-based rule and places it below the cloned application-based rule, which eliminates the risk of losing application availability because traffic that doesn't match the cloned rule flows through to the port-based rule. When traffic from legitimate applications hasn't hit the port-based rule for a reasonable period of time, you can remove it to complete that rule's migration.

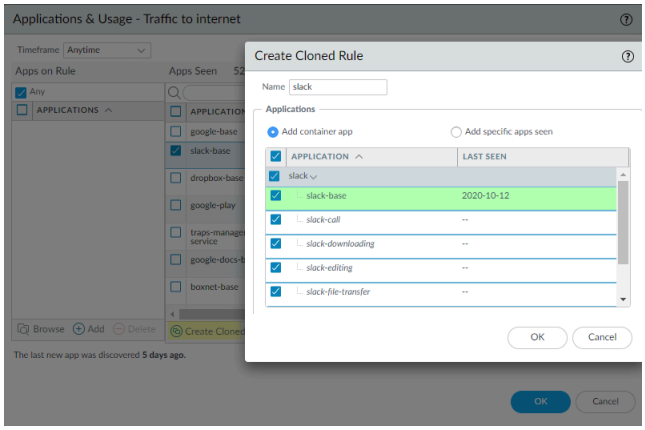
To **clone** a port-based rule:

1. In **Apps Seen**, click the check box next to each application you want in the cloned rule. Keep in mind that it takes an hour or more to update **Apps Seen**.
2. Click **Create Cloned Rule**. In the **Create Cloned Rule** dialog, **Name** the cloned rule ("slack" in this example) and add other applications in the same container and application

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

dependencies, if required. For example, to clone a rule by selecting the slack-base application:



The green text is the selected application to clone. The container application (**slack**) is in the gray row. The applications listed in *italics* are applications that have not been seen on the rule but are in the same container as the selected application. Individual applications that have been seen on the rule are in normal font. All the applications are included in the cloned rule by default (**Add Container App**, which adds all the applications in the container, is selected by default) to help prevent the rule from breaking in the future.

3. If you want to allow all of the applications in the container, leave **Add container app** selected. This also "future proofs" the rule because when an application is added to the container app, it's automatically added to the rule.

   If you want to constrain access to some of the individual applications in the container, uncheck the box next to each individual application you don't want users to access. This also unchecks the container app, so if you want to allow new applications in the container later, you have to add those applications individually.

   If you uncheck the container app, all the apps are unchecked and you manually select the apps you want to include in the cloned rule.
4. If application dependencies are listed in a box below the Applications (there are none in this example), leave them checked. The applications you selected need those application dependencies to run. Common dependencies include **ssl** and **web-browsing**.
5. Click **OK** to add the new application-based rule directly above the port-based rule in the rulebase.
6. **Commit** the configuration.

When you clone a rule and **Commit** the configuration, the applications you select for the cloned rule are removed from the original port-based rule's **Apps Seen** list. For example, if a port-based rule has 16 **Apps Seen** and you select two individual applications and one dependent application for the cloned rule, after cloning, the port-based rule shows 13 **Apps Seen** because the three selected applications have been removed from the port-based rule (16-3 = 13). The cloned rule shows the three added applications in **Apps on Rule**.

Creating a cloned rule with a container app works a bit differently. For example, a port-based rule has 16 **Apps Seen** and you select one individual application and a container app for the cloned rule. The container app has five individual applications and has one dependent application. After cloning, the cloned rule shows seven **Apps on Rule**—the

https://t.me/learningnets

individual application, the five individual applications in the container app, and the dependent application for the container app. However, in the original port-based rule, **Apps Seen** shows 13 applications because only the individual application, the container app, and the container app's dependent application are removed from the port-based rule.

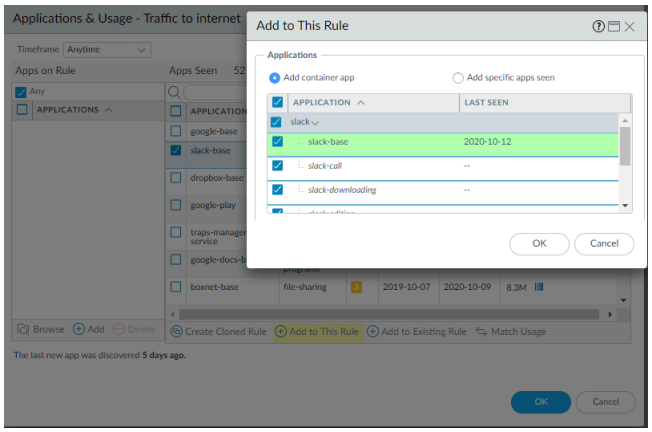
In contrast to cloning, adding applications to a port-based rule replaces the rule with the resulting application-based rule. Adding applications to a rule is simpler than cloning, but riskier because you may inadvertently miss applications that should be on the rule, and the original port-based rule is no longer in the rulebase to catch accidental omissions. However, adding applications to port-based rules that apply to only a few well-known applications migrates the rule quickly to an application-based rule. For example, for a port-based rule that only controls traffic to TCP port 22, the only legitimate application is SSH, so it's safe to add applications to the rule.

*Adding applications using the traditional Security policy rule's **Application** tab does not change **Apps Seen** or **Apps on Rule**. To preserve accurate application usage information, when replacing port-based rules with application-based rules, add applications using **Add to This Rule** or **Match Usage** (or create a cloned rule or add applications to an existing application-based rule instead) in **Apps Seen**.*

There are three ways to replace a port-based rule with an application-based rule by adding applications (**Add to This Rule** and **Match Usage** in **Apps Seen** and **Add** in **Apps on Rule**):

- **Add to This Rule** applications from **Apps Seen** (applications that matched the rule). Keep in mind that it takes an hour or more to update **Apps Seen**.

  1. Select applications from **Apps Seen** on the rule.
  2. Click **Add to This Rule**. In the **Add to This Rule** dialog, add other applications in the same container app and application dependencies, if required. For example, to add slack-base to a rule:



  Similar to the **Create Cloned Rule** dialog, the green text in **Add to This Rule** is the selected application to add to the rule. The container app (**slack**) is in the gray row. The applications listed in *italics* are applications that have not been seen on the rule but are in the same container as the selected application. Individual applications that have been seen on the rule are in normal font. All the applications are included in the cloned rule by

https://t.me/learningnets

default (**Add Container App**, which adds all the applications in the container, is selected by default) to help prevent the rule from breaking in the future.

3. If you want to allow all of the applications in the container, leave **Add container app** selected. This also "future proofs" the rule because when an application is added to the container app, it's automatically added to the rule.

   If you want to constrain access to some of the individual applications in the container, uncheck the box next to each individual application you don't want users to access. This also unchecks the container app, so if you want to allow new applications in the container later, you have to add those applications individually.

   If you uncheck the container app, all the apps are unchecked and you manually select the apps you want to include in the cloned rule.
4. If application dependencies are listed in a box below the Applications (there are none in this example), leave them checked. The applications you selected need those application dependencies to run.
5. Click **OK** to replace the port-based rule with the new application-based rule.

When you **Add to This Rule** and **Commit** the configuration, the applications you didn't add are removed from **Apps Seen** because the new application-based rule no longer allows them. For example, if a rule has 16 **Apps Seen** and you **Add to This Rule** three applications, the resulting new rule shows only those three added applications in **Apps Seen**.

Add to This Rule with a container app works a bit differently. For example, a port-based rule has 16 **Apps Seen** and you select one individual application and a container app to add to the new rule. The container app has five individual applications and has one dependent application. After adding the applications to the rule, the new rule shows seven **Apps on Rule**—the individual application, the five individual applications in the container app, and the dependent application for the container app. However, **Apps Seen** shows 13 applications because the individual application, the container app, and the container app's dependent application are removed from that list.

- Add all of the **Apps Seen** on the rule to the rule at one time with one click (**Match Usage**).

  *Port-based rules allow any application, so **Apps Seen** may include unneeded or unsafe applications. Use **Match Usage** to convert a rule only when the rule has seen a small number of well-known applications with legitimate business purposes. A good example is TCP port 22, which should only allow SSH traffic, so if SSH is the only application seen on a port-based rule that opens port 22, you can safely **Match Usage**.*

  1. In **Apps Seen**, click **Match Usage**. Keep in mind that it takes an hour or more to update **Apps Seen**. All the applications in **Apps Seen** are copied to **Apps on Rule**.
  2. Click **OK** to create the application-based rule and replace the port-based rule.
- If you know the applications you want on the rule, you can **Add** applications manually in **Apps on Rule**. However, this method is equivalent to using the traditional Security policy rule **Application** tab and does not change **Apps Seen** or **Apps on Rule**. To preserve accurate

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

application usage information, convert rules using **Add to This Rule**, **Create Cloned Rule**, or **Match Usage** in **Apps Seen**.

1. In **Apps on Rule**, **Add** (or **Browse**) and select applications to add to the rule. This is equivalent to adding applications on the **Application** tab.
2. Click **OK** to add the applications to the rule and replace the port-based rule with the new application-based rule.

   *Because this method is equivalent to adding applications using the **Application** tab, the dialog to add application dependencies doesn't pop up.*

STEP 5 | For each application-based rule, set the **Service** to **application-default**.

If business needs require you to allow applications (for example, internal custom applications) on non-standard ports between particular clients and servers, restrict the exception to only the required application, sources, and destinations. Consider rewriting custom applications so they use the application default port.

STEP 6 | **Commit** the configuration.

STEP 7 | Monitor the rules.

- **Cloned rules**—Monitor the original port-based rule to ensure the application-based rule matches the desired traffic. If applications you want to allow match the port-based rule, add them to the application-based rule or clone another application-based rule for them. When only applications that you don't want on your network match the port-based rule for a reasonable period of time, the cloned rule is robust (it catches all the application traffic you want to control) and you can safely remove it.
- **Rules with Added Applications**—Because you convert only port-based rules that have a few well-known applications directly to application-based rules, in most cases the rule is solid from the start. Monitor the converted rule to see if the expected traffic matches the rule—if there's less traffic than expected, the rule may not allow all of the necessary applications. If there's more traffic than expected, the rule may allow unwanted traffic. Listen to user feedback—if users can't access applications they need for business purposes, the rule (or another rule) may be too tight.

Rule Cloning Migration Use Case: Web Browsing and SSL Traffic

A port-based rule that allows web access on TCP ports 80 (HTTP web-browsing) and 443 (HTTPS SSL) provides no control over which applications use those open ports. There are many web applications, so a general rule that allows web traffic allows thousands of applications, many of which you don't want on your network.

This use case shows how to migrate a port-based policy that allows all web applications to an application-based policy that allows only the applications you want, so you can safely enable the applications you choose to allow. For rules that see a lot of applications, cloning the original port-based rule is safer than adding applications to the rule because adding replaces the port-based rule, so if you inadvertently forget to add a critical application, you affect application availability. And if you **Match Usage**, which also replaces the port-based rule, you allow all of the applications the rule has seen, which could be dangerous, especially with web browsing traffic.

https://t.me/learningnets

Cloning the rule retains the original port-based rule and places the cloned rule directly above the port-based rule in the rulebase, so you can monitor the rules. Cloning also allows you to split rules that see a lot of different applications—such as a port-based web traffic rule—into multiple application-based rules so you can treat different groups of applications differently. When you're sure you're allowing all the applications you need to allow in the cloned rule (or rules), you can remove the port-based rule.

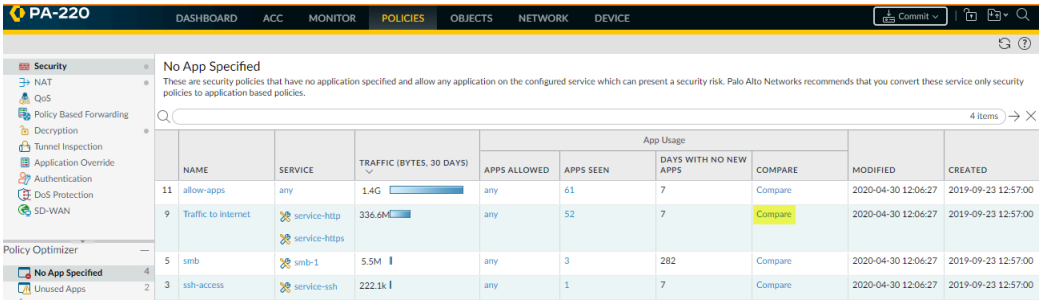
This example clones a port-based web traffic rule to create an application-based rule for web-based file sharing traffic (a subset of the application traffic seen on the port-based rule).

This example does not apply to using the New App Viewer to clone App-ID Cloud Engine (ACE) applications (see the ACE documentation for examples of how to do this); ACE requires a SaaS Security Inline license.

STEP 1 | Navigate to **Policies** > **Security** > **Policy Optimizer** > **No App Specified** to view the port-based rules.

STEP 2 | Click **Compare** for the rule you want to migrate.

In this example, the port-based rule that allows web access is named Traffic to internet.

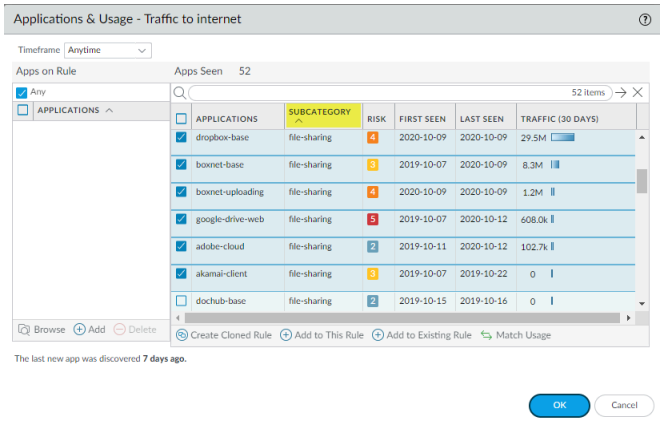


https://t.me/learningnets

STEP 3 | Use the sorting options to review and select the applications you want to allow from **Apps Seen**.

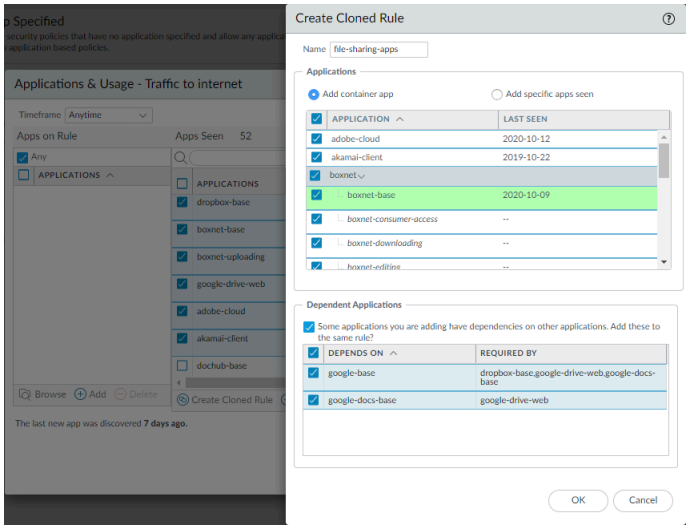
*The number of **Apps Seen** is updated approximately every hour, so if you don't see as many applications as you expect, check again after about an hour. Depending on the firewall's load, it may take longer than one hour for these fields to update.*

For example, click **Subcategory** to sort the applications, scroll to the file-sharing subcategory, and then select the applications you want to allow. Alternatively, you can filter (search) for file-sharing applications.



STEP 4 | Click **Create Cloned Rule** and **Name** the cloned rule (file-sharing-apps in this example).

Create Cloned Rule shows the selected applications shaded green, the container apps shaded gray, individual applications in the container that haven't been seen on the rule in *italics*, and individual applications that have been seen on the rule in normal text font. Scrolling through **Applications** shows all the container apps and their individual applications.



Create Cloned Rule also shows the dependent applications for the selected applications. In this example, some of the selected applications require (**Required By**) the google-base and google-docs-base applications to run.

https://t.me/learningnets

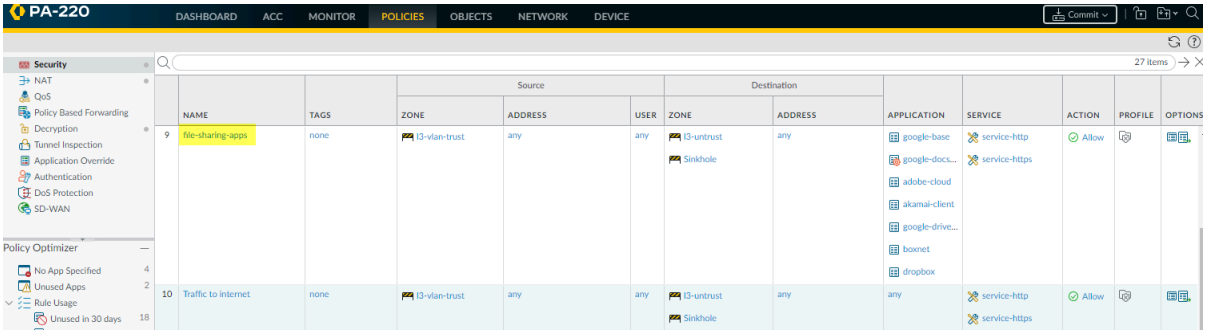
STEP 5 | Select the applications you want in the cloned rule.

For applications you don't want to include, uncheck the corresponding box, which also unchecks the container app. If you don't include the container app, then when new apps are added to the container, they won't automatically be added to the rule.

If you uncheck the container app, all the individual applications in the container are unchecked and you must select the apps you want to add manually.

STEP 6 | Click **OK** to create the cloned rule.

STEP 7 | In **Policies** > **Security**, the cloned rule (file-sharing-apps) is inserted in the rulebase above the original port-based rule (Traffic to internet).



STEP 8 | Click the rule name to edit the cloned rule, which inherits the properties of the original port-based rule.

STEP 9 | On the **Service/URL Category** tab, delete service-http and service-https from **Service**.

This changes the **Service** to **application-default**, which prevents applications from using non-standard ports and further reduces the attack surface.

If business needs require you to allow applications (for example, internal custom applications) on non-standard ports between particular clients and servers, restrict the exception to only the required application, sources, and destinations. Consider rewriting custom applications so they use the application default port.

STEP 10 | On the **Source**, **User**, and **Destination** tabs, tighten the rule to apply to only the right users in only the right locations (zones, subnets).

For example, you may decide to limit web file sharing activity to only the user groups that have business reasons to share files across the web.

STEP 11 | Click **OK**.

STEP 12 | **Commit** the configuration.

STEP 13 | Repeat the process for other application categories in the port-based web access rule until your application-based rules allow only the applications you want to allow on your network.

When traffic you want to allow stops hitting the original port-based rule for a sufficient amount of time to be confident that the port-based rule is no longer needed, you can remove the port-based rule from the rulebase.

https://t.me/learningnets

Add Applications to an Existing Rule

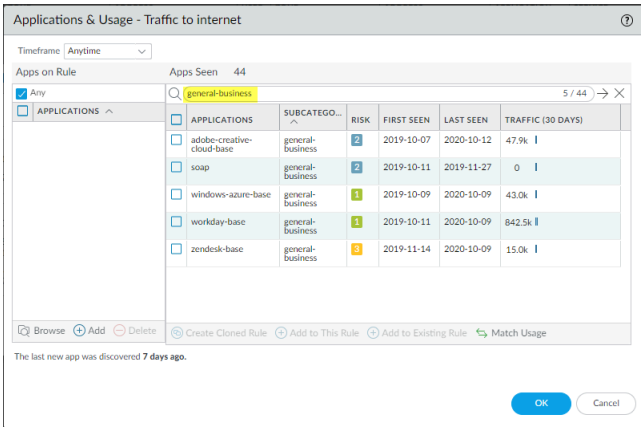
In some cases, you may want to add applications learned (seen) on a port-based rule to a rule that already exists. For example, an administrator may create a cloned application-based rule for general business web applications from a port-based rule that allows internet access (a port 80/443 rule). Later, the administrator notices that the port-based internet access rule has seen more general business applications and wants to add some or all of them to the cloned application-based rule (cloning another application-based rule for the same type of application would create an unnecessary rule and complicate the rulebase).

This example assumes that an application-based Security policy rule to control general business traffic already exists or was cloned from a port-based internet access rule, similarly to the Rule Cloning Migration Use Case: Web Browsing and SSL Traffic. In that example, we cloned an application-based rule from the port-based internet access rule and changed the new rule's Service to application-default to prevent web-based applications from using non-standard ports.

In addition to adding applications to an existing application-based rule, you can add applications to an existing port-based rule. This converts the port-based rule to an application-based rule for the applications you add to the rule. If you do this, go to the rule and change the Service to application-default to prevent the applications from using non-standard ports (also, the Service configured on the rule may not match the application).

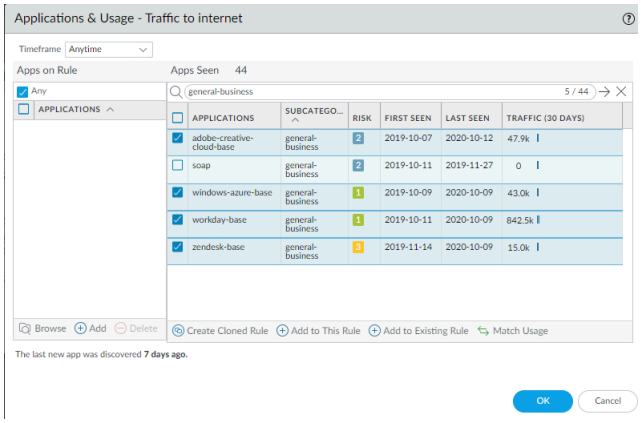
This example does not apply to using the New App Viewer to add App-ID Cloud Engine (ACE) applications to an existing rule (see the ACE documentation for examples of how to do this); ACE requires a SaaS Security Inline license.

STEP 1 | You check the port-based internet access rule and discover that the rule has seen general business applications and that you need to allow some of them for business purposes.

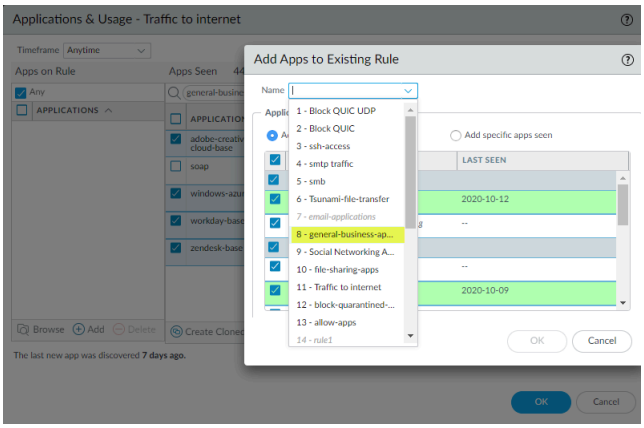


https://t.me/learningnets

STEP 2 | Select the general business apps you want to add to the existing rule.



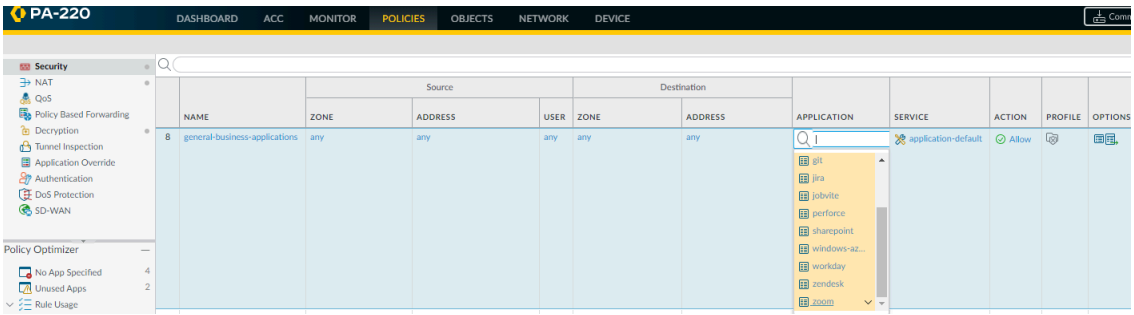
STEP 3 | Click **Add to Existing Rule** and select the **Name** of the rule to which you want to add the applications, in this example, **general-business-applications**.



STEP 4 | Click **OK** in **Add Apps to Existing Rule**to add the selected applications to the **general-business-applications** rule.

STEP 5 | Click **OK** in **Applications & Usage**.

STEP 6 | The updated rule now controls the original applications on the rule and the applications you just added.



https://t.me/learningnets

Identify Security Policy Rules with Unused Applications

If you have application-based Security policy rules that allow a large number of applications, you can remove unused applications (applications never seen on the rules) to tighten those rules so that they only allow applications actually seen in traffic that matches the rule. Identifying and removing unused applications from Security policy rules is a best practice that strengthens your security posture by reducing the attack surface.

STEP 1 | Identify Security policy rules that have unused applications.

Policies > **Security** > **Policy Optimizer** > **Unused Apps** displays all application-based rules that are configured with applications that have not matched (been seen on) the rule. This means that these rules allow applications that you may not use in your network (or that another rule shadows the rule, so traffic that you expect to match the rule matches an earlier rule in the rulebase).



*The number of **Apps Allowed** and **Apps Seen** are updated approximately every hour, so if you configure applications on a rule and don't see as many **Apps Allowed** as you expect, check again after about an hour. Depending on the firewall's load, it may take longer than one hour for these fields to update.*

STEP 2 | Prioritize which rules with unused applications to modify first.

Policies > **Security** > **Policy Optimizer** > **Unused Apps** enables you to sort rules without affecting their order in the rulebase and provides other information that helps you prioritize rules to clean up based on your business goals and risk tolerance.

- The difference between **Apps Allowed** (the number of applications on the allow list) and **Apps Seen** (the number of allowed applications actually seen on the rule) shows how many applications are configured on each rule but not actually seen on the rule, which indicates to what extent the rule is over-provisioned. Click **Apps Allowed** to sort by the number of applications allowed in a rule and click **Apps Seen** to sort by the number of applications actually seen on a rule.
- **Days with No New Apps** (click to sort) shows you the number of days since the last time a new application hit the rule. This indicates how likely it is that the rule is mature and won't see any applications that haven't already been seen. The longer the **Days with No New Apps**, the less likely that new applications will hit the rule and the more likely that you know all the applications the rule allows.
- **Created** and **Modified** dates also help determine whether a rule has matured enough to understand whether applications not seen on the rule may be seen at a later date or if the rule has seen all the applications expected to hit the rule. The longer the time since a rule was **Modified**, the more likely the rule is mature. (If **Created** and **Modified** are the same, the rule hasn't been modified.)
- **Hit Count**—Displays rules with the most matches over a selected time frame. You can exclude rules for which you reset the hit counter and specify the exclusion time period in

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

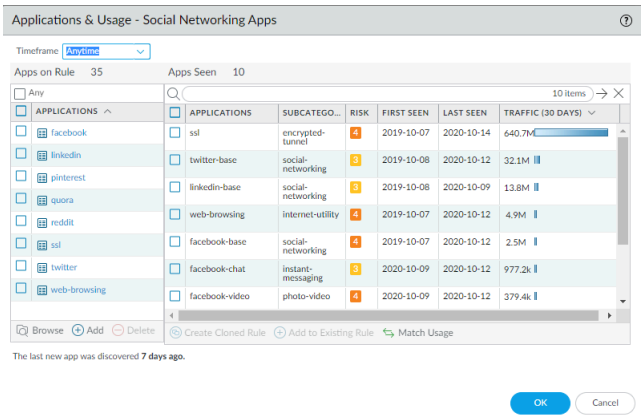
days. Excluding rules with recently reset hit counters prevents misconceptions about rules that show fewer hits than you expect because you didn't know the counter was reset.

*You can also use **Hit Count** to View Policy Rule Usage.*

You can also click **Traffic (Bytes, 30 days)** to sort by the amount of traffic a rule has seen over the last 30 days. Use this information to prioritize which rules to modify first. For example, you can prioritize rules with the largest difference between **Apps Allowed** and **Apps Seen** and that also have the most **Days with No New Apps**, because those rules have the greatest number of unused applications and are the most mature.

STEP 3 | Review the **Apps Seen** on the rule.

On **Unused Apps**, click **Compare** or the number in the **Apps Seen** column to open **Applications & Usage**, which shows the applications configured on the rule (**Apps on Rule**) and the **Apps Seen** on the rule.

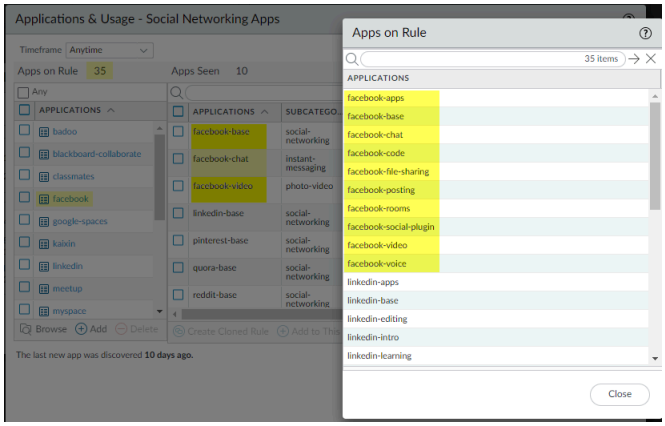


- The number next to **Apps Seen** (10 in this example) is the number of applications that matched the rule. Keep in mind that it takes at least one hour for the firewall to update **Apps Seen**.
- The number next to **Apps on Rule** (35 in this example) is how many applications are configured on the rule, which is calculated by counting each application in a container app (but not the container app itself—if you configure a container app on the rule, the rule allows the container app's individual applications). Because the **Applications** list shows only the applications you configure manually on the rule, when you configure a container app on a rule, **Applications** only shows the container app, not all of the individual applications in the container (unless you also manually configure the individual applications on the rule). For this reason, the number of **Apps on Rule** may not be the same as the number of applications you see in the **Applications** list.
- Click the number next to **Apps on Rule** to see all of the individual applications on the rule.

  This example rule has 10 **Apps Seen** (applications that matched the rule) but allows 35 **Apps on Rule**. The **facebook** container app is configured on the rule and the rule sees traffic from the individual applications facebook-base, facebook-chat, and facebook-video (**Apps Seen**).

https://t.me/learningnets

When you click the **Apps on Rule** number, the **Apps on Rule** dialog displays the individual applications allowed, but not the container app itself.



You cannot add or delete applications from the pop-up dialog.

Compare the **Apps Seen** on the rule to the **Apps on Rule**. If an application on the rule isn't used (you don't see the application or you don't see applications in an allowed container in **Apps Seen**), consider removing the application from the rule to reduce the attack surface. Take into account periodically used applications, such as for quarterly or annual events, which may look unused if you don't examine a long enough time frame. **Timeframe** enables you to select the time frame for the **Apps Seen** on the rule. Select **Anytime** to see every application seen over the life of the rule. Depending on the **Created** or **Modified** date in the **No App Specified** dialog and the time between periodic events, the rule may not have been on the firewall long enough to see all periodically used applications.

STEP 4 | Remove unused applications from the rule.

Delete (or **Add**) applications in **Apps on Rule** to remove (or add) applications manually, or **Match Usage** to add the **Apps Seen** on the rule and delete applications for which no matching traffic has been seen on the rule with one click.

To remove applications from the rule manually, select applications from **Apps on Rule** and **Delete** them. Ensure that none of the applications are required for periodic events before you remove them from the rule. (You can also add or delete applications on the Security policy rule's **Application** tab.)

Match Usage moves the **Apps Seen** on the rule to **Apps on Rule** and removes all unused applications from the rule.

*You can clone rules from **Policies** > **Security** and from **No App Specified** to Migrate Port-Based to App-ID Based Security Policy Rules. You can't clone a rule starting from **Unused Apps**.*

STEP 5 | **Commit** the configuration.

https://t.me/learningnets

STEP 6 | Monitor updated rules and listen to user feedback to ensure that updated rules allow the applications you want to allow and don't inadvertently block periodically used applications.

> *The number of **Apps Allowed** and **Apps Seen** are updated approximately every hour. After you remove all of the unused applications from a rule, the rule remains listed in **Policies** > **Security** > **Policy Optimzer** > **Unused Apps** until the firewall updates the display. When the firewall updates the display and the number of **Apps Allowed** is the same as the number of **Apps Seen**, the rule no longer displays in the **Unused Apps** screen. However, depending on the firewall's load, it may take longer than one hour for these fields to update.*

High Availability for Application Usage Statistics

When you configure two firewalls as a High Availability (HA) pair, the application usage statistics are local to the firewall that generates the Traffic logs for the application. Where you can view application usage statistics also depends in part on the HA configuration:

- **Active/Passive**—The active device generates the application usage statistics. If a passive device has seen no user traffic, then only the active device displays the application usage statistics. If a passive device has seen traffic, then the passive device only displays the application usage statistics from the traffic that it has seen.

  On a failover, the application usage statistics are based only on the Traffic logs generated on the newly active device (the device that was passive before the failover).
- **Active/Active**—The device that owns a session generates the Traffic logs for that session, so the application usage statistics for a session are only available on the device that owns the session. If one active device owns a session, the other active device does not display that session's application usage statistics.

How to Disable Policy Optimizer

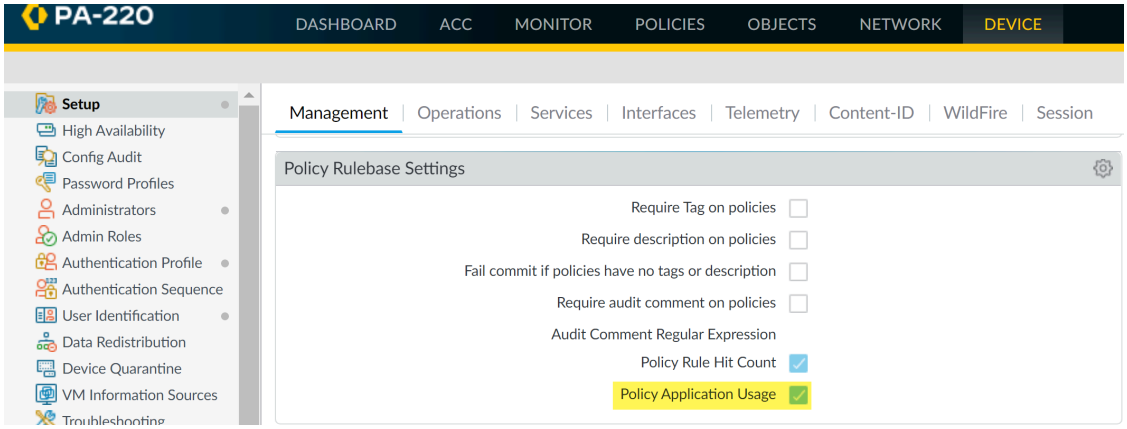
Policy Optimizer is enabled by default. Policy Optimizer provides many capabilities that make it easier to Migrate Port-Based to App-ID Based Security Policy Rules and to Identify Security Policy Rules with Unused Applications and remove the unused applications from the rules, but if you wish to disable the feature, you can.

STEP 1 | Navigate to **Device** > **Setup** > **Management** > **Policy Rulebase Settings**.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 2 | Select the **Policy Application Usage** check box to enable the feature and deselect the check box to disable the feature.



https://t.me/learningnets

App-ID Cloud Engine

The App-ID Cloud Engine (ACE) is a new service that enables the firewall or Panorama to download App-IDs from the cloud for applications that do not have specific predefined App-IDs from the Palo Alto Networks content team. These are the applications that the firewall identifies as ssl, web-browsing, unknown-tcp, or unknown-udp traffic. Use ACE App-IDs in Security policy rules to gain visibility into and control those applications and use Policy Optimizer to add and manage applications in Security policy. You cannot use ACE App-IDs in any other types of policy rules. ACE:

- Vastly increases the number of known App-IDs to identify and control applications. As ACE defines new App-IDs for applications, they become available on the firewall.
- Speeds up the availability and delivery of new App-IDs to the firewall.
- Speeds up and can automate the addition of applications to Security policy through the use of Application Filters in Security policy rules.
- Dramatically increases visibility into applications that previously were identified as ssl, web-browsing, unknown-tcp, or unknown-udp.

ACE requires a SaaS Security Inline subscription. Each appliance that uses ACE must have a valid device certificate installed.

All hardware platforms that support PAN-OS 10.1 or later support ACE and all appliances on which you want to use ACE require PAN-OS 10.1 or later. Panorama cannot push and commit ACE-based polices or objects to firewalls that don't have a SaaS Security Inline license installed or to firewalls that run an earlier version of PAN-OS than 10.1.

ACE is supported in the US, APAC, and EU GCP regions. The region is selected automatically based on your CDL region.

*Verify that the firewall uses the correct Content Cloud FQDN (**Device** > **Setup** > **Content-ID** > **Content Cloud Setting**) for your region and change the FQDN if necessary:*

- *US—**hawkeye.services-edge.paloaltonetworks.com***
- *EU—**eu.hawkeye.services-edge.paloaltonetworks.com***
- *APAC—**apac.hawkeye.services-edge.paloaltonetworks.com***

ACE data, including traffic payloads, is sent to the servers in the selected region. If you specify a Content Cloud FQDN that is outside of your region (for example, if you are in the EU region but you specify the APAC region FQDN), you may break your country's or your organization's privacy and legal regulations.

Predefined content-delivered App-ID delivers new applications once per month and you need to analyze the new App-IDs before you install them to understand changes that they may make to Security policy rules. The monthly cadence and need for analysis slows down the adoption of new App-IDs in policy. Although Palo Alto Networks will continue to provide new App-IDs via monthly content updates that you need to review, ACE improves the adoption of new App-IDs by providing on-demand App-IDs for applications initially identified as any of the following four types:

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

- **ssl**—Encrypted SSL traffic is by far the most common type of network traffic, with most experts claiming that it exceeds 90% of total traffic. If you don't or can't decrypt that traffic, the firewall often can only identify it as ssl instead of as the actual underlying application.
- **web-browsing**—The firewall can't specifically identify some unencrypted web-browsing traffic because there are so many applications that content-delivered App-ID can't keep up with the ever-increasing amount.
- **unknown-tcp** and **unknown-udp**—This traffic may be internal or custom applications or unknown external applications. It's important to identify that traffic by its specific App-ID so that you can make intelligent access decisions and construct appropriate Security policy rules to control and inspect the traffic.

ACE provides specific identification of these applications, which enables you to understand them and control them appropriately in policy.

ACE App-IDs do not identify other types of public applications and do not identify private and custom applications. The ACE App-ID catalog does not contain predefined, content-provided App-IDs. Content-provided App-IDs still arrive monthly in content updates.

When the firewall encounters ssl, web-browsing, unknown-tcp, or unknown-udp traffic, the firewall sends the payload to ACE for analysis. If there is a matching App-ID in the ACE database, ACE returns the App-ID to the requesting firewall. If ACE has no matching App-ID for the traffic, ACE sends the payload to the Machine Learning (ML) engine. The ML engine analyzes the payload and develops the new App-ID in conjunction with the human content team and drops traffic that isn't related to applications. When development finishes, the ML engine uploads new App-ID to the ACE database, and the requesting firewall (and any other firewalls) can download the App-ID and use it in Security policy.

Because it can take several minutes to retrieve an application from ACE for which it has an App-ID and longer if a new App-ID must be developed, cloud application detection is not inline on the firewall. The firewall does not wait for a verdict to process the application traffic. The firewall processes the traffic as ssl, web-browsing, unknown-tcp, or unknown-udp until it receives an App-ID from ACE and then continues to process the traffic in that way until you receive the new App-ID and use it in Security policy.

*If you downgrade a firewall or Panorama after ACE has been enabled and ACE cloud App-IDs are still in use in Security policy rules or Application Groups, the downgrade fails. The fail reason lists the objects that you need to remove from the configuration in order to downgrade. Remove those objects from the configuration and **Commit** the configuration, and then the downgrade will succeed.*

- Prepare to Deploy App-ID Cloud Engine
- Enable or Disable the App-ID Cloud Engine
- App-ID Cloud Engine Processing and Usage
- New App Viewer (Policy Optimizer)
- Add Apps to an Application Filter with Policy Optimizer
- Add Apps to an Application Group with Policy Optimizer
- Add Apps Directly to a Rule with Policy Optimizer

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

- Replace an RMA Firewall (ACE)
- Impact of License Expiration or Disabling ACE
- Commit Failure Due to Cloud Content Rollback
- Troubleshoot App-ID Cloud Engine

Prepare to Deploy App-ID Cloud Engine

There are several onboarding tasks to do before the firewall can use the App-ID Cloud Engine (ACE). You can deploy ACE on standalone firewalls or use Panorama to deploy ACE on managed firewalls.

Before a firewall can use ACE to provide specific App-IDs for traffic previously identified as ssl, web-browsing, unknown-tcp, and unknown-udp traffic, the PAN-OS administrator and the SaaS Security administrator must work together to:

- Install a valid device certificate on each appliance that will use ACE, including Panorama appliances that manage ACE firewalls. (PAN-OS administrator.)
- Activate SaaS Security Inline on each firewall that will use ACE. Panorama doesn't require a license. (SaaS Security administrator.)
- Configure a service route for communication between the firewall and ACE. (PAN-OS administrator.)
- Enable ACE on Panorama appliances which manage firewalls that will use ACE. (PAN-OS administrator.)

  On firewalls, ACE is enabled by default after activating SaaS Security Inline.

- Create Security policy rule that allows ACE traffic. (PAN-OS administrator.)
- Configure Log Forwarding from the firewall to the Cortex Data Lake (CDL). (PAN-OS administrator.)

At the appropriate step in the following procedure, the PAN-OS administrator should notify the SaaS Security administrator that the deployment is ready for SaaS Security Inline activation. After activating SaaS Security Inline, the SaaS Security Inline administrator should notify the PAN-OS administrator that the deployment is ready to complete on the PAN-OS devices. Communication between the administrators is essential to achieving a smooth deployment.

Requirements:

- Standalone firewalls, Panorama appliances, and managed firewalls must run PAN-OS 10.1 or later.
- All ACE firewalls must have purchased a SaaS Security Inline license. Panorama does not require a license to manage ACE firewalls or push ACE configurations to managed firewalls.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

- All ACE appliances must be able to connect to the US, APAC, or EU GCP region, depending on your location (the region is selected automatically based on your CDL region).

  Verify that the firewall uses the correct Content Cloud FQDN (**Device** > **Setup** > **Content-ID** > **Content Cloud Setting**) for your region and change the FQDN if necessary:

  - US—**hawkeye.services-edge.paloaltonetworks.com**
  - EU—**eu.hawkeye.services-edge.paloaltonetworks.com**
  - APAC—**apac.hawkeye.services-edge.paloaltonetworks.com**

  ACE data, including traffic payloads, is sent to the servers in the selected region. If you specify a Content Cloud FQDN that is outside of your region (for example, if you are in the EU region but you specify the APAC region FQDN), you may break your country's or your organization's privacy and legal regulations.

The PAN-OS administrator completes the first two steps of the procedure and then hands it off to the SaaS Security Inline administrator for activation (Step 3). After activation, the SaaS Security Inline administrator hands the rest of the procedure off to the PAN-OS administrator to complete on the PAN-OS devices.

STEP 1 | Bring the firewall and Panorama (if using) online. (PAN-OS administrator.)

STEP 2 | Install a Device Certifcate on individual firewalls so that they can use cloud services or use Panorama to Install the Device Certificate for Managed Firewalls. (PAN-OS administrator.)

Hand off the next step to the SaaS Security administrator.

STEP 3 | Activate SaaS Security Inline on every firewall that will use ACE. Activation enables ACE on the firewalls. (SaaS Security administrator.)

Panorama does not require a SaaS Security Inline license to manage firewalls that use ACE. Only managed firewalls need licenses, which you must retrieve manually as shown in the next step.

Hand off the rest of the steps to the PAN-OS administrator.

STEP 4 | Retrieve the SaaS Security Inline license on each firewall—Panorama doesn't need a license—and verify that it is activated. (PAN-OS administrator.)

The SaaS Security administrator's activation sets up the licenses for the firewall, so you don't have to go to the Customer Support Portal or obtain Auth Codes.

1. Go to **Device** > **Licenses** > **License Management** and select **Retrieve license keys from license server** to retrieve the license.
2. Check **Device** > **Licenses** to ensure that the SaaS Security Inline license is active.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 5 | Configure a data services (dataplane) service route so that the firewall can communicate with the App-ID Cloud Engine. (PAN-OS administrator.)

You can push this configuration to managed firewalls from Panorama. Both Panorama and the managed firewalls must run PAN-OS 10.1 or later.

By default, the firewall uses the management interface as the source interface for the data services service route, but it is recommended that you configure a dataplane interface that has connectivity to cloud services as the **Source Interface** and **Source Address** for data services, as shown later in this step.

The issue on firewalls is that if an explicit proxy is configured on the management interface and you use it for the data services service route, then the management interface can only connect to the Knowledge Cloud Service (KCS), which manages the cloud application and signatures. When an explicit proxy is configured on the management interface, it cannot connect to the Detection Cloud Service (DCS), which checks the application payload against existing ACE App-IDs and provides verdicts. KCS and DCS are services in the ACE cloud. If the management interface has an explicit proxy configured, you can't use it for the data services service route for ACE because it can't connect to all of the services. In this case, you must use a dataplane interface on the firewall to connect to the data services.

Panorama uses the management port by default to connect to the KCS and does not connect to the DCS.

To configure the service route on a data plane interface instead of using the default management interface:

1. Select **Device** > **Setup** > **Services** then in **Service Features**, select **Service Route Configuration**.
2. **Customize** a service route.
3. Select the **IPv4** protocol.
4. Click **Data Services** in the Service column to open the **Service Route Source** dialog box.
5. Select a **Source Interface** and **Source Address** (these cannot be the management interface).

   The source interface must have internet connectivity. The best practice is to use a dataplane interface that has connectivity to cloud services. See Configure Interfaces and Create an Address Object for more information about creating source interfaces and addresses.
6. Click **OK** to set the source interface and address.
7. Click **OK** to set the Service Route Configuration.
8. Select **Policies** > **Security** and add a Security policy rule that allows traffic from the source interface you specified earlier in this procedure to the FQDN addresses for the KCS and DCS services, which are **`kcs.ace.tpcloud.paloaltonetworks`** (KCS service for all regions) and **`hawkeye.services-edge.paloaltonetworks.com`** (US region DCS service), **`eu.hawkeye.services-edge.paloaltonetworks.com`** (EU

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

region DCS service), or **apac.hawkeye.services-edge.paloaltonetworks.com** (APAC region DCS service).

Also add and allow the following two FQDNs in a new or existing Security policy rule: **ocsp.paloaltonetworks.com** and **crl.paloaltonetworks.com** for certificate verification.

Finally, add or modify a Security policy rule to allow ACE traffic by allowing the following three applications: **paloalto-ace**, **paloalto-ace-kcs**, and **paloalto-dlp-service**.

STEP 6 | Make sure that hawkeye.services-edge.paloaltonetworks.com and kcs.ace.tpcloud.paloaltonetworks are reachable on firewalls and that kcs.ace.tpcloud.paloaltonetworks is reachable on Panorama devices. (PAN-OS administrator.)

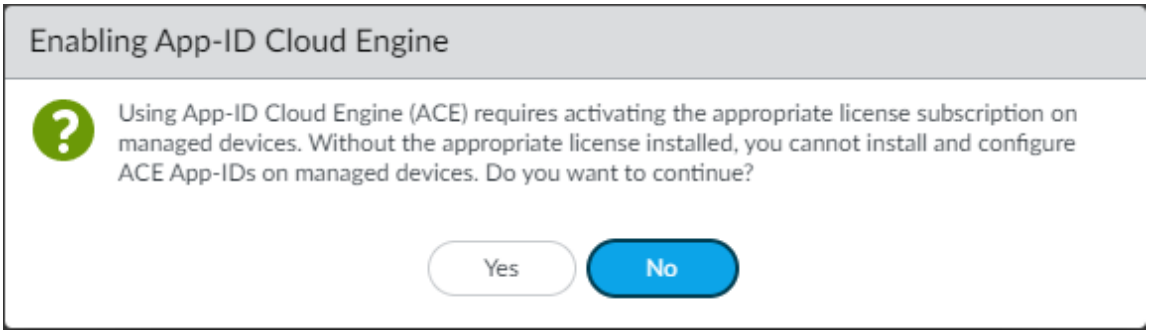
Run the operational command `admin@fw1> show cloud-appid connection-to-cloud`. The output informs you whether the connection is working and if the license is installed.

STEP 7 | (Panorama only) Enable ACE on any Panorama appliance that manages ACE-enabled firewalls. (PAN-OS administrator.)

ACE is disabled by default on Panorama.

If you push ACE configurations to managed groups that do not have ACE-enabled firewalls (some or all firewalls in the group do not have ACE enabled), the push fails.

1. Navigate to **Panorama** > **Setup** > **ACE** > **Settings**.
2. Click edit (⚙) and then de-select **Disable App-ID Cloud Engine.**
3. Click **OK**.
4. The **Enable App-ID Cloud Engine** dialog appears.



   Click **Yes** to enable ACE.
5. **Commit** the change.

STEP 8 | Wait for the App-ID catalog to download. (PAN-OS administrator.)

There are fewer than four thousand content-provided App-IDs. After you download the ACE catalog, you see many thousands more applications on the firewall and can confirm by checking **Objects** > **Applications** or by using the operational CLI command `show cloud-appid cloud-app-data application all` to see the new App-IDs.

https://t.me/learningnets

STEP 9 | (Panorama only) Push the desired configuration to the managed firewall(s). (PAN-OS administrator.)

STEP 10 | Configure Log Forwarding to Cortex Data Lake (CDL) and enable Log Forwarding with the correct Log Forwarding profile in Security policy rules. (PAN-OS administrator.)

A SaaS Security Inline connection to CDL is required for SaaS visibility and to support SaaS App-ID Policy Recommendation. At a minimum, you must forward Traffic logs and URL logs to CDL for SaaS Security Inline to work properly.

Enable or Disable the App-ID Cloud Engine

The App-ID Cloud Engine (ACE) is disabled by default on Panorama and enabled by default on firewalls when the SaaS Security Inline license is installed. You must enable ACE on Panorama appliances that manage ACE-enabled firewalls.

To enable or disable ACE:

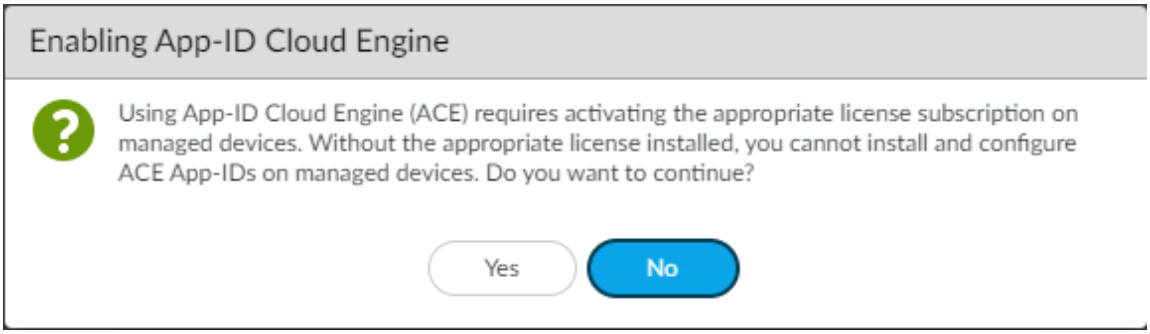
STEP 1 | Navigate to **Device** > **Setup** > **ACE** > **Settings** on the firewall or **Panorama** > **Setup** > **ACE** > **Settings** on Panorama.

STEP 2 | Click edit (⚙) and then either de-select **Disable App-ID Cloud Engine** to enable ACE or select **Disable App-ID Cloud Engine** to disable ACE.

ACE is disabled by default.

STEP 3 | Click **OK**.

STEP 4 | (Only if enabling ACE) If you are enabling ACE, the **Enable App-ID Cloud Engine** dialog appears.



If the firewall or Panorama-managed firewalls have the SaaS Security Inline license installed, click **Yes** to enable ACE.

STEP 5 | **Commit** the change.

App-ID Cloud Engine Processing and Usage

When the firewall downloads App-ID Cloud Engine (ACE) App-IDs, it's important to understand how the firewall handles those App-IDs and how the firewall handles ACE App-IDs when there are also predefined content-based App-IDs for the same applications. The Palo Alto Networks content team develops predefined content-based App-IDs and updates them with modified and new App-IDs through application content updates (a valid support contract is required for updates).

https://t.me/learningnets

ACE requires a SaaS Security Inline license. Firewalls that don't support ACE have only predefined content-based App-IDs. The ACE App-ID catalog doesn't contain content-based App-IDs.

You can only use ACE App-IDs in Security policy rules. You cannot use ACE App-IDs in any other type of policy rule.

- When the firewall first connects to the App-ID cloud engine, the firewall downloads a catalog of the available ACE App-IDs, and you can use those App-IDs in Security policy. It does not download the full signatures. The catalog enables you to use ACE App-IDs in Security policy even if the applications have never been seen on the firewall. ACE pushes catalog updates to firewalls regularly so that firewalls have access to the latest ACE App-IDs.

  If an application arrives at the firewall that is identified as ssl, web-browsing, unknown-tcp, or unknown udp and the firewall doesn't have its signature, the firewall sends the payload to ACE. If ACE has an App-ID for the traffic, ACE sends the full signatures back to the firewall. If the traffic doesn't match any ACE signatures, then ACE sends the payload to the Machine Learning (ML) engine. The ML engine analyzes the payload and develops the new App-ID in conjunction with the human content team and drops traffic that isn't related to applications. The ML engine sends the new App-ID to ACE and requesting firewalls can download it and use it in Security policy.

  Because it can take several minutes to retrieve an application from ACE for which it has an App-ID and longer if a new App-ID must be developed, cloud application detection is not inline on the firewall. The firewall does not wait for a verdict to process the application traffic. The firewall processes the traffic as ssl, web-browsing, unknown-tcp, or unknown-udp until it receives an App-ID from ACE and then continues to process the traffic in that way until you receive the new App-ID and use it in Security policy.
- When a firewall requests an App-ID from ACE, the firewall doesn't hold the traffic, it continues to process the traffic as usual until it receives an App-ID from ACE.
- The firewall handles cloud App-IDs downloaded from ACE differently than it handles content-delivered App-IDs. You don't have to examine how new ACE App-IDs affect Security policy before they are installed on the firewall because the firewall uses ACE App-IDs according to previously existing Security policy. Your existing Security policy rules control the new ACE App-IDs until you explicitly use ACE App-IDs in Security policy. For example:

  1. An application is identified only as "ssl" and you have a Security policy rule that allows SSL traffic, so the ssl rule allows that application.
  2. The firewall sees the ssl application and sends the payload to ACE.
  3. ACE identifies the actual application. If the application exists in the ACE database, then ACE sends that App-ID to the firewall. If it's a new application for which ACE does not have an App-ID, then ACE forwards the payload to the ML Engine. The firewall does not receive the App-ID until the ML Engine and the human content team assign an App-ID and send it to ACE.
  4. The rule that allows ssl traffic still allows the newly-identified application, even though its App-ID is no longer "ssl". (However, if you use the new ACE App-ID in Security policy, that policy controls the traffic. Similarly, traffic previously identified as web-browsing, unknown-

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

tcp, and unknown-udp continues to obey the Security policy rules that control those types of traffic until you use the ACE App-IDs in Security policy.)

In contrast to ACE App-IDs, if the App-ID was a predefined, content-provided App-ID, then the rule that allows ssl traffic would not longer match the application. The firewall would block it if no Security policy rule explicitly allows it.

The exception to this behavior is if another Security policy rule specifies the App-ID given to the traffic by ACE. The Security policy rule with the specific App-ID takes precedence over the rule with the less specific ssl App-ID. If the rule that specifies the actual App-ID is a block rule, the application is blocked even though there is a rule that allows ssl traffic. The rule with the more specific (granular) App-ID is the one the firewall acts on.

In this example, if you add the cloud App-ID for the application previously identified as "ssl" to an existing or cloned rule either directly or by using an Application Filter or an Application Group, that rule controls the application. The "ssl" rule no longer controls the application because the application is specifically identified in another rule.

If you don't explicitly add new ACE App-IDs to Security policy rules, the firewall continues to control them with the same rules that controlled those applications before they had ACE App-IDs and were identified as ssl, web-browsing, unknown-tcp, or unknown-udp traffic. For example, if the firewall sees an application identified as unknown-tcp and then receives an ACE App-ID for the traffic, but you don't use that ACE App-ID in a Security policy rule, then the firewall still controls that traffic using the rule that controls unknown-tcp traffic—if you block unknown-tcp traffic, then the traffic is blocked, and if you allow unknown-tcp traffic, the traffic is allowed.

- The firewall caches some information so that the firewall can check the cache and avoid repeatedly sending data to the cloud and requesting verdicts. If the firewall is waiting for a verdict from ACE, the firewall doesn't forward the same application data twice.
- A particular container app and its functional applications are either all cloud-based App-IDs or all content-based App-IDs. One App-ID delivery method defines a container app and all of its functional apps.
- If cloud-based, content-provided, and user-defined custom App-ID names overlap, the order of precedence is:

  1. **Custom App-IDs**—These App-IDs take precedence over all other App-IDs and if the firewall attempts to download an ACE application with the same App-ID, the commit fails because two applications on the same firewall cannot have the same App-ID.

     In this case, you can rename the custom application, or if the custom application is the same application as the ACE application, you can delete the custom application and use the ACE application.
  2. **Content-based, predefined App-IDs**—These App-IDs take precedence over ACE cloud App-ID definitions.
  3. **ACE cloud App-IDs**—Custom and content-based App-IDs take precedence over ACE App-ID definitions.

https://t.me/learningnets

- If an App-ID matches a container app, the firewall downloads the container app's App-ID and all of its functional apps. For example, if the firewall retrieves the facebook container app, it also retrieves facebook-base, facebook-chat, facebook-post, etc.
- When you take any of the following actions on an ACE App-ID, you affect how Security policy handles that ACE App-ID because the firewall will take action based on the specific ACE App-ID instead of based on the application's previous ssl, web-browsing, unknown-tcp, or unknown-udp App-ID:
  - Create Application Filters to automate adding ACE App-IDs to Security policy.

    Use Application Filters to automate adding ACE App-IDs to Security policy rules. When a new App-ID matches an Application Filter, the firewall automatically adds it to the filter. When you use that Application Filter in a Security policy rule, the rule controls the application traffic for the new App-IDs that were automatically added to the filter. In other words, Application Filters are your "Easy Button" for securing ACE App-IDs automatically to gain maximum application visibility and control with minimum effort.
  - Add the App-IDs to Application Groups.
  - Use Policy Optimizer to add the App-IDs to a cloned rule or to an existing rule, or to an existing Application Filter or Application Group. You can use Policy Optimizer to create new Application Filters and Application Groups directly from within the Policy Optimizer tool. Use Policy Optimizer's sorting and filtering tools to prioritize the rules to work on and to assess how many ACE App-IDs match those rules.
  - Add an ACE App-ID directly to a new or existing Security policy rule.

  When you add a cloud App-ID to a Security policy rule directly or by using an Application Filter or an Application Group, that rule controls the application. Until you take one of these actions to control cloud-delivered App-IDs, the firewall uses existing ssl, web-browsing, unknown-tcp, or unknown-udp Security policy rules to control ACE applications.
- When you create Application Filters, exclude ssl and web-browsing from the filters. Together, ssl and web-browsing match all browser-based cloud applications, so an Application Filter that includes ssl and web-browsing matches all browser-based cloud applications.
- Active/Passive High Availability:
  - The Active firewall syncs the ACE catalog to the passive firewall so that they have identical catalogs.
  - The Passive firewall does not initiate connections to ACE until it becomes the Active firewall.
- Active/Active High Availability: Each device fetches catalogs and signatures separately, so the catalogs and signatures are not synced. However, commits fail if the catalog is out-of-sync on peers and ACE App-IDs are referenced in Security policy rules. If the catalogs of peer HA firewalls are out-of-sync, wait a few minutes for the updates to reach the devices and become in-sync again.

https://t.me/learningnets

- A Panorama commit all/push failure to managed firewalls occurs if:
  - Managed firewalls do not have a valid SaaS Security Inline license so they do not have the ACE catalog. In this case, remove the ACE objects from the pushed configuration and try again.
  - The connection between a managed firewall and ACE goes down and the pushed configuration includes applications that are not in the ACE catalog on the firewall. In this case, check the connection to the ACE cloud and re-establish the connection if necessary so that the firewall can update its catalog.

    The operational CLI command `show cloud-appid connection-to-cloud` provides the cloud connection status and the ACE cloud server URL.
  - The ACE catalog on Panorama and the ACE catalog on managed firewalls is out-of-sync, which results in pushed configurations that include ACE apps that are not in the firewall's catalog. If the connection between the firewall and ACE is up, the outdated catalog will update in the next few minutes automatically and resolve the issue. (Wait five minutes and try again.)

    You can also use the CLI command `debug cloud-appid cloud-manual-pull check-cloud-app-data` to update the catalog manually.
- Some Security profiles such as the File Blocking, Antivirus, WildFire, and DLP profiles can specify applications as part of the profile. Only content-provided App-IDs are supported in Security profiles. ACE App-IDs are not supported in Security profiles. ACE App-IDs are intended for use in Security policy rules only.
- Because ACE App-IDs are supported only for Security policy, they are not supported in Application Override, Policy-Based Forwarding (PBF), QoS, or SD-WAN policy rules.

  You cannot see ACE App-IDs in Application Override or PBF rule configuration. However, ACE App-IDs are visible (able to be selected) in QoS and SD-WAN policy rule configuration and may be present in Application Groups or Application Filters applied to a rule. If you use ACE App-IDs in these rules, the policy doesn't control the application traffic and there is no effect on the application traffic—the rules do not apply to the ACE App-ID traffic even though ACE App-IDs were added to the rule.

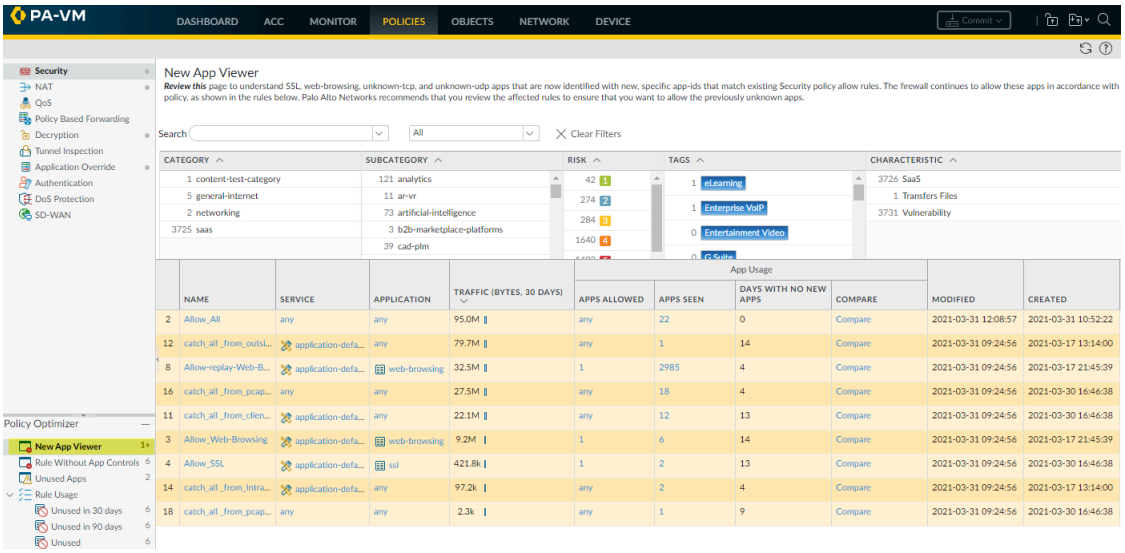
New App Viewer (Policy Optimizer)

The Policy Optimizer **New App Viewer** shows you the Security policy rules that match downloaded cloud App-IDs from ACE. You can use Policy Optimizer to manage the newly identified applications and add them to cloned rules or to existing rules. Select **Policies** > **Security** to expose the **New App Viewer** in the **Policy Optimizer** portion of the interface and then select **New App Viewer**.

The upper portion of the screen is similar to **Objects** > **Application Filters**. It works in a similar manner and filters the Security policy rules shown in the lower portion of the screen. You can filter the rules that allow applications by category, subcategory, etc. The only categories and subcategories available for filtering are the ones that match the new applications on the rules listed in the lower half of the screen, so you don't waste your time filtering for applications that aren't there.

https://t.me/learningnets

When you filter the rules, only the rules that include the filtered applications are shown in the lower portion of the screen. Rules that have not seen the apps in the filter are removed from the list. (You can see them all again by removing the filter.)



Click the number in the **Apps Seen** column to open the **Applications & Usage** dialog to change the way the firewall handles the cloud-based applications in Security policy. Add ACE App-IDs to Security policy rules using an Application Filter, an Application Group, Policy Optimizer, or by directly adding an ACE App-ID to a rule. Until you take one of these actions to control cloud-delivered App-IDs, the firewall continues to see the traffic as ssl, web-browsing, unknown-tcp, or unknown-udp traffic and uses existing ssl, web-browsing, unknown-tcp, or unknown-udp Security policy rules to control the applications.

Add Apps to an Application Filter with Policy Optimizer

Add App-IDs from the App-ID Cloud Engine (ACE, and/or content-provided App-IDs) to new or existing Application Filters to automate how you control cloud App-IDs in Security policy. When new ACE App-IDs match an Application Filter, the firewall adds them to the filter automatically. When you use the Application Filter in a Security policy rule, the rule automatically controls new ACE App-IDs as they arrive at the firewall and are added to the filter.

ACE provides App-IDs for applications that were previously identified as ssl, web-browsing, unknown-tcp, or unknown-udp.

Using Application Filters is a best practice because they:

- Improve your security posture. Application Filters automate adding new ACE App-IDs to Security policy rules that you design specifically to handle a particular type of application traffic, instead of matching the traffic to more general ssl, web-browsing, unknown-tcp, or unknown-udp rules.
- Save time. Firewall administrators can configure Application Filters to handle different types of traffic so that adding new ACE App-IDs to policy is automatic and requires no further effort by the administrator.

https://t.me/learningnets

When you create Application Filters, exclude ssl and web-browsing from the filters. Together, ssl and web-browsing match all browser-based cloud applications, so an Application Filter that includes ssl and web-browsing matches all browser-based cloud applications.

Use Policy Optimizer to add ACE App-IDs to Application Filters and to apply the filters to cloned or existing rules and control the ACE App-IDs in Security policy.

STEP 1 | Go to **Policies** > **Security** and then select **Policy Optimizer** > **New App Viewer**.

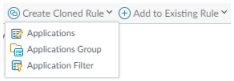
If the firewall has identified traffic with ACE App-IDs, a number displays next to **New App Viewer** in the left navigation window. The screen displays the Security policy rules that match cloud App-IDs.

STEP 2 | Click the number in **Apps Seen** for a Security policy rule to see the cloud-delivered applications that matched the rule in the **Applications & Usage** dialog.

STEP 3 | Select the applications that you want to add to an existing or new Application Filter.

You can sort and filter the applications in **Apps Seen** by subcategory, risk, amount of traffic seen over the last 30 days, or when the application was first or last seen.

STEP 4 | Select **Application Filter** from **Create Cloned Rule** or **Add to Existing Rule**, depending on how you want to handle the applications.



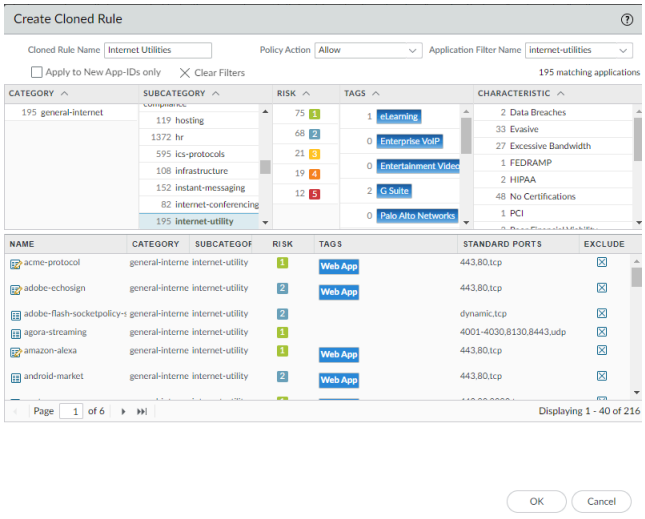
*The maximum number of applications you can clone using **Create Cloned Rule** is 1,000 applications. If there are more than 1,000 applications that you want to move to a different rule, use **Add to Existing Rule** instead. If you want to move the applications to a new rule, simply create the rule first (**Policies** > **Security**) and then use Policy Optimizer to add them to that rule.*

https://t.me/learningnets

STEP 5 | Select or create the Application Filter for the cloned or existing rule. Creating an Application Filter using Policy Optimizer is the almost exactly the same as using **Objects** > **Application Filters** to create an Application Filter—you use the same filtering tools and options.

Create Cloned Rule:

1. Type the **Cloned Rule Name** (the name for the cloned rule, which will appear in the Security policy rulebase immediately above the original rule).
2. Select the **Policy Action** (Allow or Deny).
3. Select the **Application Filter Name** from the menu or type the name of a new Application Filter.
4. Select whether the filter should **Apply to New App-IDs only** or if it should apply to all App-IDs.
5. Use the Category, Subcategory, Risk, Tags, and Characteristic values to filter the types of applications you want to add to the Application Filter. The firewall automatically adds new applications that meet the filter criteria to the Application Filter.

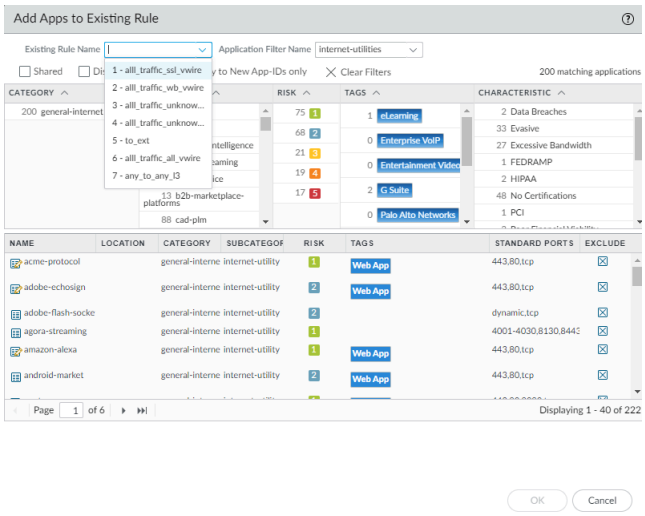


6. Click **OK** to add the applications to the new or existing Application Filter. The firewall includes the applications that you selected in Step 3 in the Application Filter.
7. **Commit** the changes.

Add to Existing Rule:

1. Select the **Existing Rule Name** to add the selected applications to an existing rule in an Application Filter.
2. Select the **Application Filter Name** from the menu or type the name of a new Application Filter.
3. Select whether the Application Filter is **Shared**, whether you want to **Disable override** of application characteristics for the filter, and whether the filter should **Apply to New App-IDs only** or if it should apply to all App-IDs.
4. Use the Category, Subcategory, Risk, Tags, and Characteristic values to filter the types of applications you want to add to the Application Filter. The firewall automatically adds new applications that meet the filter criteria to the Application Filter.

https://t.me/learningnets



5. Click **OK** to add the applications to the new or existing Application Filter. The firewall includes the applications that you selected in Step 3 in the Application Filter.
6. **Commit** the changes.

Add Apps to an Application Group with Policy Optimizer

Add App-IDs from the App-ID Cloud Engine (ACE, and/or content-provided App-IDs) to new or existing Application Groups and use the Application Groups in Security policy rules to control cloud App-IDs in Security policy.

ACE provides App-IDs for applications that were previously identified as ssl, web-browsing, unknown-tcp, or unknown-udp.

Use Policy Optimizer to add ACE App-IDs to Application Groups and to apply the groups to cloned or existing rules and control the ACE App-IDs in Security policy.

STEP 1 | Go to **Policies** > **Security** and then select **Policy Optimizer** > **New App Viewer**.

If the firewall or Panorama has downloaded ACE App-IDs, a number displays next to **New App Viewer** in the left navigation window. The screen displays the Security policy rules that match downloaded cloud App-IDs.

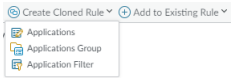
STEP 2 | Click the number in **Apps Seen** for a Security policy rule to see the cloud-delivered applications that matched the rule in the **Applications & Usage** dialog.

STEP 3 | Select the applications that you want to add to an existing or new Application Group.

You can sort and filter the applications in **Apps Seen** by subcategory, risk, amount of traffic seen over the last 30 days, or when the application was first or last seen.

https://t.me/learningnets

STEP 4 | Select **Application Group** from **Create Cloned Rule** or **Add to Existing Rule**, depending on how you want to handle the applications.



*The maximum number of applications you can clone using **Create Cloned Rule** is 1,000 applications. If there are more than 1,000 applications that you want to move to a different rule, use **Add to Existing Rule** instead. If you want to move the applications to a new rule, simply create the rule first (**Policies** > **Security**) and then use Policy Optimizer to add them to that rule.*

STEP 5 | Select or create the Application Group for the cloned or existing rule. Creating Application Groups using Policy Optimizer is similar to using **Objects** > **Application Groups** to create an Application Group.

Create Cloned Rule:

1. Type the **Cloned Rule Name** (the name for the cloned rule, which will appear in the Security policy rulebase immediately above the original rule).
2. Select the **Policy Action** (Allow or Deny).
3. In **Add to Application Group**, select the Application Group to which you want to add the applications that you selected in 3.
4. Select whether to **Add container app** (default) or only to **Add specific apps seen**.

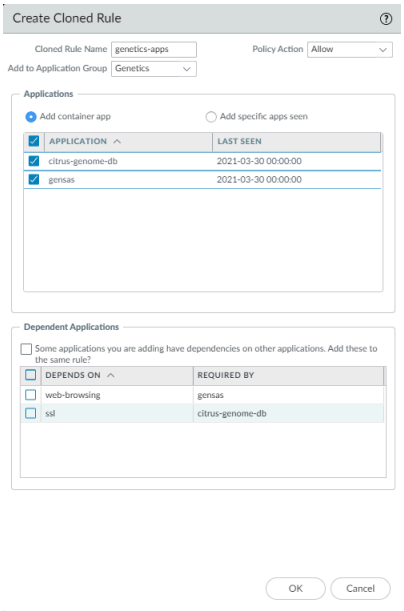
   When you add the container app, you also add all of the functional apps in that container, including functional apps that have not yet been seen on the firewall. For example, if you add the “facebook” container app, that also adds facebook-base, facebook-chat, facebook-posting, etc., and also any future applications added to the container. The container app and its functional apps are subject to the Security policy rule to which you add the Application Group. Selecting the container app essentially future-proofs and automates security for the container’s apps so that you don’t have to manually add new apps in that container to your Security policy.

   Adding only the specific apps seen means that only the applications that you selected are added to the Application Group. If new applications in the same container app arrive at the firewall, the Application Group doesn’t control them and you have to manually decide how to handle the new apps.
5. In some cases, the applications that you want to place in an Application Group require (depend on) other applications to function. In those cases, the **Create Cloned Rule** dialog box includes **Dependent Applications**, where you can select whether to add those

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

applications to the cloned rule. Add the dependent applications to the rule to ensure that the selected applications function properly.



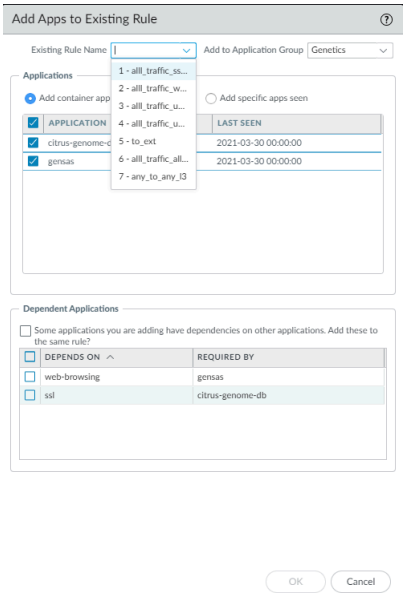
6. Click **OK** to add the applications to the new or existing Application Group.
7. **Commit** the changes.

Add Apps to Existing Rule:

1. Select the **Existing Rule Name** to add the selected applications to an existing rule in an Application Group.
2. Select the Application Group in **Add to Application Group** or type the name of a new Application Group.
3. As with cloning the rule, you can choose whether to **Add container app** or **Add specific apps seen**. Adding the container app adds all the functional apps in the container and any future apps added to that container. Adding only the specific apps only adds the specific selected apps.
4. As with cloning the rule, in some cases, the applications that you want to place in an Application Group require (depend on) other applications to function. In those cases, the **Add Apps to Existing Rule** dialog box includes **Dependent Applications**, where you can

https://t.me/learningnets

select whether to add those applications to the cloned rule. Add the dependent applications to the rule to ensure that the selected applications function properly.



5. Click **OK** to add the applications to the new or existing Application Group.
6. **Commit** the changes.

Add Apps Directly to a Rule with Policy Optimizer

You can add App-ID Cloud Engine (ACE, and/or content-provided App-IDs) App-IDs directly to a cloned or existing rule with Policy Optimizer. However, consider using Application Filters to automate adding ACE App-IDs to Security policy as they arrive at the firewall instead of adding them manually.

ACE provides App-IDs for applications that were previously identified as ssl, web-browsing, unknown-tcp, or unknown-udp.

STEP 1 | Go to **Policies** > **Security** and then select **Policy Optimizer** > **New App Viewer**.

If the firewall or Panorama has downloaded ACE App-IDs, a number displays next to **New App Viewer** in the left navigation window. The screen displays the Security policy rules that match downloaded cloud App-IDs.

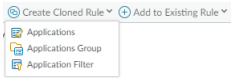
STEP 2 | Click the number in **Apps Seen** for a Security policy rule to see the cloud-delivered applications that matched the rule in the **Applications & Usage** dialog.

STEP 3 | Select the applications that you want to add to an existing or cloned Security policy rule.

You can sort and filter the applications in **Apps Seen** by subcategory, risk, amount of traffic seen over the last 30 days, or when the application was first or last seen.

https://t.me/learningnets

STEP 4 | Select **Applications** from **Create Cloned Rule** or **Add to Existing Rule**, depending on how you want to handle the applications.



*The maximum number of applications you can clone using **Create Cloned Rule** is 1,000 applications. If there are more than 1,000 applications that you want to move to a different rule, use **Add to Existing Rule** instead. If you want to move the applications to a new rule, simply create the rule first (**Policies** > **Security**) and then use Policy Optimizer to add them to that rule.*

STEP 5 | Add the selected applications to a cloned rule or to an existing rule.

Create Cloned Rule:

1. Type the **Name** (the name for the cloned rule, which will appear in the Security policy rulebase immediately above the original rule). The cloned rule has the same action (allow or deny) as the original rule.
2. Select whether to **Add container app** (default) or only to **Add specific apps seen**.

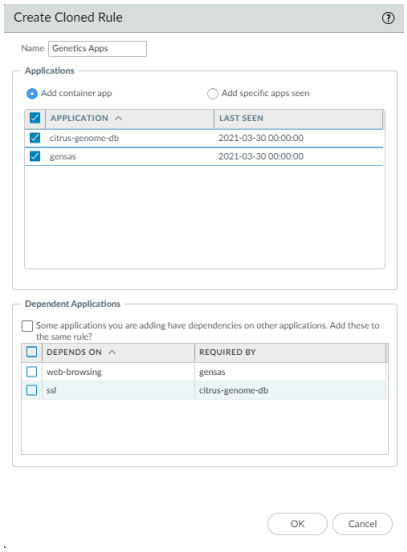
   When you add the container app, you also add all of the functional apps in that container, including functional apps that have not yet been seen on the firewall. For example, if you add the “facebook” container app, that also adds facebook-base, facebook-chat, facebook-posting, etc., and also any future applications added to the container. The container and its functional apps are subject to the Security policy rule that you are cloning. Selecting the container app essentially future-proofs and automates security for the container’s apps so that you don’t have to manually add new apps in that container to your Security policy.

   Adding only the specific apps seen means that only the applications that you selected are added to the cloned rule. If new applications in the same container app arrive at the firewall, the cloned rule doesn’t control them and you have to manually decide how to handle the new apps.
3. In some cases, the applications that you want to add to a rule require (depend on) other applications to function. In those cases, the **Create Cloned Rule** dialog box includes **Dependent Applications**, where you can select whether to add those applications to

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

the cloned rule. Add the dependent applications to the rule to ensure that the selected applications function properly.



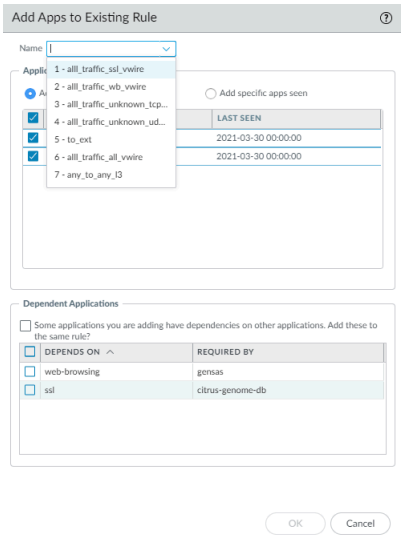
4. Click **OK** to add the applications to the cloned rule.
5. **Commit** the changes.

Add Apps to Existing Rule:

1. Select the **Name** of the existing rule to which you want to add the selected applications.
2. As with cloning the rule to add applications, you can choose whether to **Add container app** or **Add specific apps seen**. Adding the container app adds all the functional apps in the container and any future apps added to that container. Adding only the specific apps only adds the specific selected apps.
3. As with cloning the rule, in some cases, the applications that you want to add to a rule require (depend on) other applications to function. In those cases, the **Add Apps to Existing Rule** dialog box includes **Dependent Applications**, where you can select whether to add

https://t.me/learningnets

those applications to the cloned rule. Add the dependent applications to the rule to ensure that the selected applications function properly.



4. Click **OK** to add the applications to the existing rule.
5. **Commit** the changes.

Replace an RMA Firewall (ACE)

To restore the configuration on a managed firewall when there is a Return Merchandise Authorization (RMA), the procedure is to:

- Review Before Starting RMA Firewall Replacement.
- On Panorama, replace the serial number of the old firewall with the new firewall's serial number.
- In the firewall CLI, check to ensure that the firewall is online and connected to the Knowledge service so that the firewall can download the cloud application catalog:

  1. Access the firewall CLI.
  2. In Operational mode, check the cloud App-ID connection:

     ```
     admin@vm1> show cloud-appid connection-to-cloud
     ```

     If the firewall is connected to the cloud, the show command returns:

     ```
     ACE Cloud server: kcs.ace.tpcloud.paloaltonetworks.com:443Cloud connection: connected
     ```

     Information about the connection also displays. If the firewall is not connected to the cloud, check whether DNS services are functioning and check for any other network-related connectivity issues.
- With the firewall connected to the App-ID cloud, Restore the Firewall Configuration after Replacement.

https://t.me/learningnets

Impact of License Expiration or Disabling ACE

If you enable App-ID Cloud Engine (ACE) on a firewall, download ACE App-IDs to the firewall, and then use those App-IDs in objects such as Application Filters and in Security policy rules, then you need to understand what happens if the SaaS Security Inline license expires or if you disable ACE. Disabling ACE and the SaaS Security Inline license expiring both affect downloaded ACE App-IDs, the catalog of ACE App-IDs, Security policy rules that control ACE App-IDs, and objects that include ACE App-IDs. The effect is the same unless otherwise noted:

- ACE App-IDs remain on the firewall, but the firewall stops enforcing ACE App-IDs in Security policy.

  Security policy rules that control ACE App-IDs no longer control ACE App-IDs even though they are visible in the rule. Traffic that was controlled by ssl, web-browsing, unknown-tcp, or unknown-udp rules before ACE was enabled on the firewall is controlled by those rules again until you update and activate the SaaS Security Inline license and/or re-enable ACE or change those rules.
- Enforcement of Security policy rules based on ACE App-IDs stops within 4-6 hours of the license expiring (based on a timer that periodically checks license status).

  Enforcement of Security policy rules based on ACE App-IDs stops immediately after you commit the disabling ACE on the firewall.

  Disabling ACE stops enforcing Security policy rules based on ACE App-IDs as soon as you commit the change even if the SaaS Security Inline license is still valid and active.
- The catalog of ACE App-IDs remains on the firewall and on Panorama but the cloud engine no longer updates the catalog.
- The connection from the firewall to ACE no longer functions. If you re-enable ACE or renew the SaaS Security Inline license, it may take some time to download all of the catalog updates.
- If the SaaS Security Inline license expires, the ACE service stops working within 4-6 hours.

  Panorama doesn't require a SaaS Security Inline license, so there is no license to expire on Panorama. However, when the license expires on managed firewalls, configuration pushes to those firewalls from Panorama fail if they contain ACE configurations in Security policy or in Application Groups.
- Objects such as Application Filters and Application Groups are not changed, but any ACE App-IDs that you placed in those objects are no longer enforced even though the ACE App-IDs are still visible.
- If you are using SaaS Policy Recommendation, the firewall can no longer pull SaaS policy recommendations, so the SaaS administrator cannot push new policy recommendations to the firewall. Policy recommendations that were downloaded before license expiration remain in the configuration but they are not enforced (same behavior as Security policies configured with ACE App-IDs when the license expires or ACE is disabled).

Commit Failure Due to Cloud Content Rollback

Although it is extremely unlikely, it is possible that ACE App-IDs may need to be rolled back (reverted) because of bad metadata or issues with applications. If ACE must revert App-IDs and

 ©2022 Palo Alto Networks, Inc.

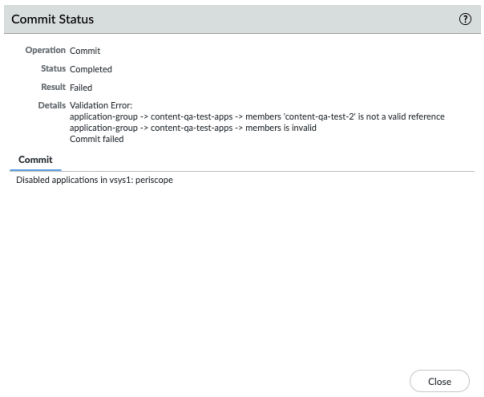
https://t.me/learningnets

you used those App-IDs in a Security policy rule (directly or in an Application Group), commit actions fail until those applications are removed from Security policy rules and from objects.

If it becomes necessary to roll back App-IDs, ACE reverts all of the most recently delivered cloud-based App-IDs, signatures, metadata, categories, subcategories, and tags from the ACE catalog. Removing the App-IDs from the catalog removes them from the firewall, which is why the commit action fails when the App-IDs are used in Security policy.

If you did not use the applications that ACE had to roll back in Security policy, there is no impact to the configuration and commit actions succeed.

When you attempt to commit a configuration after an ACE content rollback, the commit failure message lists the applications that ACE reverted, as in this example **Validation Error**:



To fix the issue, you must remove the listed applications from Security policy rules, regardless of whether they were added directly to a rule or were added using an Application Group. If the application is used in an Application Group, remove it from the Application Group.

In this example, `content-qa-test-2` is the reverted application, which is referenced in the Application Group `content-qa-test-apps`. After you remove `content-qa-test-2` from the Application Group, commit actions succeed.

Troubleshoot App-ID Cloud Engine

This topic provides general troubleshooting information for the App-ID Cloud Engine (ACE).

- To check if an appliance has a valid SaaS Security Inline license, run the operational CLI command `show cloud-appid connection-to-cloud`. If there is an issue, the command returns the message:

  ```
  ACE Error: License check failed. Check if SaaS license is installed
  and activeCloud connection: failed
  ```

  In addition, the output shows the time of the last successful connection, for example: `Last successful gRPC connection: 2021-05-20 16:00:00 -0800 PDT`

  If the license is installed and the connection to ACE is good, then the command returns the URL for the ACE cloud server connection and the status `Cloud connection: connected`, along with connection statistics and the status of the device certificate, including the certificate validity dates.

https://t.me/learningnets

- Panorama commit all/push to managed firewalls fails. Check if any of the following conditions exist and repair them:
  - Do managed firewalls have a valid SaaS Security Inline license? If not, then they do not have the ACE catalog and the commit all/push operation fails. Depending on whether you want to managed firewalls to handle ACE App-IDs, either remove the ACE objects from the pushed configuration and try again or install valid SaaS Security Inline licenses on the managed firewalls, wait for the catalog to download.

    *There are fewer than four thousand content-provided App-IDs. After you download the ACE catalog, you see many thousands more applications on the firewall and can confirm by checking **Objects** > **Applications** or by using the operational CLI command `show cloud-appid cloud-app-data application all` to see the new App-IDs.*
  - Has the connection between a managed firewall and ACE has gone down? Check the connection to the ACE cloud and restore the connection if necessary.

    The operational CLI command `show cloud-appid connection-to-cloud` provides the cloud connection status and the ACE cloud server URL.
  - The ACE catalog on Panorama and the ACE catalog on managed firewalls is out-of-sync, which results in pushed configurations that include ACE apps that are not in the firewall's catalog. If the connection between the firewall and ACE is up, the outdated catalog will update in the next few minutes automatically and resolve the issue. (Wait five minutes and try again.)

    You can also run the operational CLI command `debug cloud-appid cloud-manual-pull check-cloud-app-data` to update the catalog manually.
  - Are the firewalls all running PAN-OS 10.1 or later? (Pushing configurations that reference ACE applications and objects to firewalls running earlier versions than PAN-10 10.1 is not allowed.)
- In an HA pair (active/active or active/passive) that has an ACE configuration, if you run the operational command `show session all` or `show session id <id>`, the output for ACE applications may show the global App-ID number instead of the application name. The firewall only shows the application name if its data plane has the cloud application data. If not, then the firewall shows the global App-ID number for the application instead.
- To reset the connection to ACE (the gRPC connection), run the operational CLI command `debug cloud-appid reset connection-to-cloud`.
- View the ACE applications downloaded to the appliance with the operational CLI command `show cloud-appid cloud-app-data application`. You can view all downloaded apps or individual apps by App-ID or application name.
- View pending requests for ACE App-IDs with the operational CLI command `show cloud-appid signature-dp pending-request`. The output includes how many times the firewall sent the request to ACE (`tries`). After eleven tries, the send operation times out.
- The operational CLI command `show cloud-appid` has more useful options:

```
admin@PAN-ACE-VM-1> show cloud-appid ?
```


https://t.me/learningnets

```
> app-objects-in-policy          Show application-filter/application-
groups referred in policy
> app-to-filtergroup-mapping     Show application to matched filter
 and groups
> application                    Show Application info for UI
> application-filter             Show cloud apps in application-
filters
> application-group              Show cloud apps in application-
groups
> cloud-app-data                 Show cloud application, container
 and metadata
> connection-to-cloud            Show gRPC connection status to cloud
 application server
> ha-info                        Show statistics of cloud application
 high availability
> overlap-appid                  Show duplicated applications in
 predefined content
> signature-dp                   Show cloud signatures and
 applications used on DP
> task                           Show task on management-plane
> transaction                    Show cloud application transaction
> version                        Show Cloud-AppID version
```

- To view the global counters for ACE, run the operational CLI command `show counter global filter value all category cad` (`cad` stands for "cloud app-identification).
- To view statistics for bytes and packets received and sent to/from shared memory and to/from the security client for services such as ACE, DLP, and IoT, run the operational command `show ctd-agent statistics`.
- If you notice a discrepancy between the number of applications that match an Application Filter when you look in the user interface versus when you look in the CLI, it's because of the way the firewall counts matching applications in the user interfaces versus in the CLI:
  - When you look at an Application Filter in **Objects** > **Application Filters**, the firewall displays all of the matching applications in the ACE catalog, regardless of whether the firewall has actually seen those applications and downloaded their App-IDs, and the number count includes all of those applications.
  - When you look at an Application Filter in the CLI with the `show cloud-appid application-filter` operational command, the firewall only displays the number of matching applications for which the firewall has downloaded ACE App-IDs.

  For this reason, the user interface may show more matching applications than the CLI for the same Application Filter.

  The same thing applies to Application Groups when you look at them in the user interface versus the CLI.
- ACE App-IDs are supported for Security policy only. ACE App-IDs are not supported for any other policy type.

  However, when you configure QoS or SD-WAN policy, ACE App-IDs are visible (able to be selected) and may be present in Application Groups or Application Filters applied to the rule, but adding them to QoS or SD-WAN policy has no effect on the application traffic. (The QoS and SD-WAN policies don't control the application traffic.)

https://t.me/learningnets

SaaS App-ID Policy Recommendation

The rapid proliferation of SaaS applications makes it difficult to assign all of them specific App-IDs, gain visibility into those applications, and control them. Security policy rules that allow ssl, web-browsing, or “any” application may allow unsanctioned SaaS applications that can introduce security risks to your network. To gain visibility into those applications and control them on the firewall, SaaS Security administrators can recommend Security policy rules with specific SaaS App-IDs provided by the App-ID Cloud Engine (ACE) to PAN-OS firewall administrators. PAN-OS administrators can import those rules on firewall’s that have a SaaS Security Inline subscription.

SaaS Policy Recommendation requires a SaaS Security Inline subscription. Each appliance that uses the SaaS Policy Recommendation Engine needs to generate and install a valid device certificate or use Panorama to generate and install a valid device certificate.

A SaaS Security Inline connection to Cortex Data Lake (CDL) is required for SaaS visibility. Configure Log Forwarding to CDL and enable Log Forwarding with the correct Log Forwarding profile in Security policy rules. At a minimum, you must forward Traffic logs and URL logs to CDL for SaaS Security Inline to work properly.

All hardware platforms that support PAN-OS 10.1 or later support SaaS Policy Recommendation and all appliances on which you want to use SaaS Policy Recommendation require PAN-OS 10.1 or later. Panorama cannot push and commit SaaS Policy Recommendations to firewalls that don’t have a SaaS Security Inline license installed or to firewalls that run an earlier version of PAN-OS than 10.1.

- The *SaaS Security Administrator’s Guide* describes the SaaS Security administrator’s procedure for creating Security policy rule recommendations and then pushing them to the firewall.
- The *PAN-OS Administrator’s Guide* describes how the PAN-OS administrator imports and manages policy recommendations from the SaaS Security administrator.

The SaaS Security administrator creates the new rule, adds applications, users, and groups to the rule, and sets the rule action. The rule action can be allow or block; no other actions are permitted for pushed rules. The SaaS Security administrator then pushes the rule to the appropriate appliances and the rule appears in the firewall interface (**Device** > **Policy Recommendation** > **SaaS**).

The PAN-OS administrator evaluates the recommended rule and decides whether to implement it on the firewall. If the PAN-OS administrator chooses to implement the rule, the administrator imports it on the firewall and selects where to place the policy rule in the firewall rulebase. When a PAN-OS administrator imports a policy recommendation, the firewall creates the required HIP profiles, tags, and Application Groups automatically (the PAN-OS administrator doesn’t have to do it manually).

If the SaaS Security administrator pushes Security profiles with the policy recommendation and those profiles don’t exist on the firewall, the firewall import fails. If the profiles already exist on the firewall, the import succeeds.

If the SaaS Security administrator updates a policy rule recommendation, the PAN-OS administrator sees the update and imports it into the firewall. If the SaaS Security administrator

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

deletes a policy rule recommendation, the PAN-OS administrator sees the action and deletes the rule from the firewall Security policy rulebase.

If the SaaS Security Inline license expires, the firewall no longer pulls SaaS policy recommendations, so you see no new recommendations. However, Security policy rules that you already imported continue to work.

If you disable ACE, the firewall no longer receives new cloud application signatures and App-IDs and the firewall cannot import SaaS policy recommendations based on new ACE App-IDs.

The ACE deployment process (connecting to the cloud, installing device certificates, activating the license on the SaaS Security Portal and pushing it to Panorama and firewalls, etc.) also sets up SaaS Policy Recommendation.

Update all appliances to the latest Threat content updates.

User interface additions for this new feature include:

- **Device** > **Policy Recommendation** > **SaaS** displays policy recommendations from SaaS administrators and enables firewall administrators to import, update, remove, and control recommended SaaS policies. The page display includes Application Groups configured by the SaaS administrator for the policy.
- Role-based interface access (**Device** > **Admin Roles**) has a new option on the **Web UI** tab for SaaS policy recommendation permissions: **Device** > **Policy Recommendation** > **SaaS**.
- SaaS policy recommendations are automatically tagged **SaaSSecurityRecommended**, which is displayed in the **Tags** column in the interface.

You can import and update SaaS policy recommendations pushed by SaaS administrators and remove SaaS policy recommendations that the SaaS administrator has deleted.

- Import SaaS Policy Recommendation
- Import Updated SaaS Policy Recommendation
- Remove Deleted SaaS Policy Recommendation

Import SaaS Policy Recommendation

When a SaaS Security administrator pushes Security policy rule recommendations to a PAN-OS firewall, the PAN-OS firewall administrator can import those rules on the firewall to gain visibility into and control of the applications in the policy recommendation.

See the *SaaS Security Administrator's Guide* for the SaaS administrator's policy recommendation and push procedures. This procedure shows PAN-OS administrators how to import policy recommendations.

If the SaaS Security administrator pushes Security profiles with the policy recommendation and those profiles don't exist on the firewall, the firewall import fails. If the profiles already exist on the firewall, the import succeeds.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 1 | **Device** > **Policy Recommendation** > **SaaS** on the firewall and **Panorama** > **Policy Recommendation** > **SaaS** on Panorama show all of the SaaS policy recommendations pushed from the SaaS administrator. Push policy recommendations from Panorama to managed firewalls.

STEP 2 | Refresh ( ) **Device** > **Policy Recommendation** > **SaaS** (or **Panorama** > **Policy Recommendation** > **SaaS**) to ensure that the SaaS policy recommendations are up-to-date.

Any time you push policy recommendations from Panorama to managed firewalls, refresh ( ) the page on the firewalls to ensure that the recommendations are up-to-date.

Newly pushed policy recommendations appear at the top of the screen. **Active Recommendations** shows the value **active** and **New Updates Available** shows the value **Yes**.

STEP 3 | Select a new policy recommendation.

You import one policy recommendation at a time. The **Applications** column shows an Application Group for each policy recommendation. Click the name of the group to see the applications in that group.

The **Device** column shows the source device that the SaaS administrator configured for the rule. The term "SaaS" precedes the source device. The source device can be:

- MCD—Managed Compliant Device
- MNCD—Managed Non-compliant Device
- UMCD—Unmanaged Compliant Device
- UMNCD—Unmanaged Non-compliant Device

For example, **SaaS - MCD** indicates a managed, compliant source device.

STEP 4 | **Import Policy Rule**.

In the **Import Policy Rule** dialog:

- **Name**—Name the imported rule using a name that describes the rule's intent.

  If you specify a rule name that already exists in the Security policy rulebase, the imported rule overwrites the existing rule.

- **After Rule**—Select the rule after which to place the imported SaaS rule. Think about the firewall's rulebase and how the new rule may affect existing rules. If you do not select a rule (**No Rule Selection**), then the rule is placed at the top of the Security policy rulebase. In some cases, that's not where you want to place the rule. For example, you may want some particular block rules to always be at the top of the rulebase, such as blocking QUIC

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

protocol. Be aware of the intent of the imported rule and be careful not to shadow existing rules.

The **Description** comes from the description entered when the SaaS administrator created the rule. You can change it or leave it as-is.

The import process automatically creates an Application Group for the applications in the policy recommendation. The name of the Application Group is derived from the Name that the SaaS Security administrator gave to the rule. The firewall also automatically creates any HIP profiles and tags that the SaaS administrator applied to the rule.

STEP 5 | Click **OK** to import the rule and add it to the Security policy rulebase in the position selected in **After Rule**.

STEP 6 | When you see the status message "You've successfully updated your Security policy rules", click **OK**.

The **Location** column now shows the rule's location (vsys) on the firewall, which corresponds to the vsys to which the SaaS administrator pushed the rule.

STEP 7 | Confirm that the imported policy rule is in the Security policy rulebase (**Security** > **Policies**) at the specified location and that the firewall created the associated objects.

For example, check the Security policy rule for:

- The rule's **Source Device** is populated and shows the source device for the rule on the **Source** tab.
- The Application Group populates the rule's **Application** tab.
- Associated profiles are attached to the rule (**Actions** tab).

Also check that:

- **Objects** > **Applications Group** shows the imported Application Group.
- **Objects** > **GlobalProtect** > **HIP Objects** and **Objects** > **GlobalProtect** > **HIP Profiles** show the HIP information pushed from the SaaS Security administrator with the rule.

Import Updated SaaS Policy Recommendation

When a SaaS Security administrator pushes Security policy rule recommendations to a PAN-OS firewall (or Panorama), the PAN-OS administrator can import those rules to gain visibility into and control of the applications in the policy recommendation. However, if the SaaS administrator updates the rule, for example by adding or removing applications, the rule also needs to be updated on the firewall.

If the SaaS Security administrator pushes new or updated Application Groups, HIP profiles, or tags, the firewall automatically creates or updates those objects. If the SaaS Security administrator pushes Security profiles with the policy recommendation update and those profiles don't exist on the firewall, the firewall import fails. If the profiles already exist on the firewall, the import succeeds.

https://t.me/learningnets

STEP 1 | Refresh ( ) **Device** > **Policy Recommendation** > **SaaS** (or **Panorama** > **Policy Recommendation** > **SaaS**) to ensure that you see all of the latest SaaS policy recommendations that the SaaS administrator pushed to the firewall.

STEP 2 | Check **New Updates Available**.

If the value in the **New Updates Available** column is **No**, then there are no updates to the rule. If the value is **Yes**, then the SaaS administrator has pushed an update to the rule to the firewall. In addition, **Active Recommendations** shows the value **active**.

STEP 3 | Click the Application Group name in the **Applications** column to see the updated list of applications that the rule controls.

STEP 4 | Select a policy recommendation to update.

You update only one policy recommendation at a time.

STEP 5 | Click **Import Policy Rule** to import the policy (if there are no updates to the rule, this option is grayed out and you can't select it).

The **Import Policy Rule** dialog appears. The **Name** is already populated and cannot be changed because the rule has already been imported. **After Rule** also cannot be changed in the dialog, but if you want to change the rule's location in the Security policy rulebase, you can do that on **Policies** > **Security** in the same way that you change the position of any Security policy rule. You can change the **Description** or leave it as-is.

STEP 6 | Click **OK**.

STEP 7 | Click **Yes** in **Confirm Change** to import the updated rule (or click **No** if you don't want to import the changed rule).

The firewall automatically makes any changes to the Application Group, HIP profiles, and tags associated with the rule.

Remove Deleted SaaS Policy Recommendation

When a SaaS Security administrator pushes Security policy rule recommendations to a PAN-OS appliance, the PAN-OS administrator can import those rules to gain visibility into and control of the applications in the policy recommendation. However, if the SaaS Security administrator deletes the rule, you should also delete that rule from the PAN-OS appliance.

When a SaaS Security administrator deletes a rule, the **Active Recommendation** column shows the value **removed** (for valid rules, the value is **active**).

STEP 1 | Select a rule that the SaaS Security administrator **removed** (you can select only one rule to remove at a time).

*The **Import Policy Rule** option is grayed out because the rule can no longer be imported.*

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 2 | Click **Remove Recommendation Mapping**.

This removes local mapping of the Security policy rule on the firewall. For example, mappings to locations, users, and the rule are deleted. The **Remove Recommendation Mapping** dialog box shows you the location of the rule so that you know from where the rule is removed.

STEP 3 | Click **OK**.

STEP 4 | In the **Confirm Change** dialog, click **Yes** to remove the rule from the policy recommendation database.



This action only removes the rule from the policy recommendation rule list. It does NOT remove the rule from the Security policy rulebase. You must manually remove the rule from the rulebase.

STEP 5 | A **Status** dialog appears to confirm that the policy recommendation mapping has been removed, but you still need to remove the rule from the Security policy rulebase.

STEP 6 | Go to **Policies** > **Security** and delete the rule from the Security policy rulebase.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Application Level Gateways

The Palo Alto Networks firewall does not classify traffic by port and protocol; instead it identifies the application based on its unique properties and transaction characteristics using the App-ID technology. Some applications, however, require the firewall to dynamically open *pinholes* to establish the connection, determine the parameters for the session and negotiate the ports that will be used for the transfer of data; these applications use the application-layer payload to communicate the dynamic TCP or UDP ports on which the application opens data connections. For such applications, the firewall serves as an Application Level Gateway (ALG), and it opens a pinhole for a limited time and for exclusively transferring data or control traffic. The firewall also performs a NAT rewrite of the payload when necessary.

- *H.323 (H.225 and H.248) ALG is not supported in gatekeeper routed mode.*
- *When the firewall serves as an ALG for the Session Initiation Protocol (SIP), by default it performs NAT on the payload and opens dynamic pinholes for media ports. In some cases, depending on the SIP applications in use in your environment, the SIP endpoints have NAT intelligence embedded in their clients. In such cases, you might need to disable the SIP ALG functionality to prevent the firewall from modifying the signaling sessions. When SIP ALG is disabled, if App-ID determines that a session is SIP, the payload is not translated and dynamic pinholes are not opened. See* Disable the SIP Application-level Gateway (ALG).

When you use Dynamic IP and Port (DIPP) NAT, the Palo Alto Networks firewall ALG decoder needs a combination of IP and Port (Sent-by Address and Sent-by Port) under SIP headers (Contact and Via fields) to be able to translate the mentioned headers and open predict sessions based on them.

The following table lists IPv4, NAT, IPv6, NPTv6 and NAT64 ALGs and indicates with a check mark whether the ALG supports each protocol (such as SIP).

App-ID	IPv4	NAT	IPv6	NPTv6	NAT64
SIP	✓	✓	✓	–	–
SCCP	✓	✓	✓	–	–
MGCP	✓	✓	–	–	–
FTP	✓	✓	✓	✓	–
RTSP	✓	✓	✓	✓	–
MySQL	✓	✓	–	–	–

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

App-ID	IPv4	NAT	IPv6	NPTv6	NAT64
Oracle/ SQLNet/ TNS	✓	✓	✓	✓	—
RPC	✓	✓	—	—	—
RSH	✓	✓	—	—	—
UNIStim	✓	✓	—	—	—
H.225	✓	✓	—	—	—
H.248	✓	✓	—	—	—

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Disable the SIP Application-level Gateway (ALG)

The Palo Alto Networks firewall uses the Session Initiation Protocol (SIP) application-level gateway (ALG) to open dynamic pinholes in the firewall where NAT is enabled. However, some applications—such as VoIP—have NAT intelligence embedded in the client application. In these cases, the SIP ALG on the firewall can interfere with the signaling sessions and cause the client application to stop working.

One solution to this problem is to define an Application Override Policy for SIP, but using this approach disables the App-ID and threat detection functionality. A better approach is to disable the SIP ALG, which does not disable App-ID or threat detection.

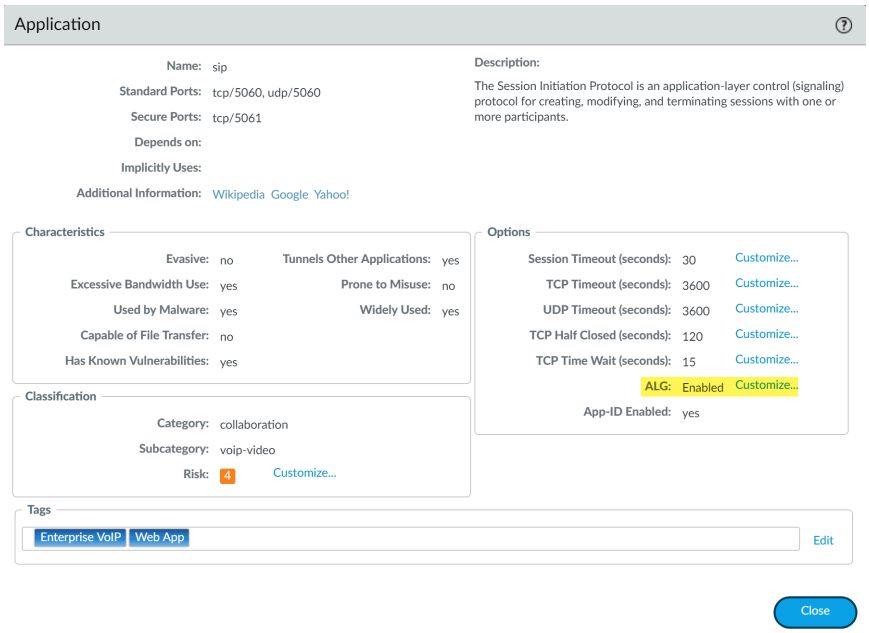
The following procedure describes how to disable the SIP ALG.

STEP 1 | Select **Objects** > **Applications**.

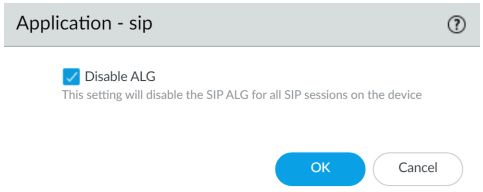
STEP 2 | Select the **sip** application.

You can type **sip** in the **Search** box to help find the sip application.

STEP 3 | Select **Customize...** for **ALG** in the Options section of the Application dialog box.



STEP 4 | Select the **Disable ALG** check box in the Application - sip dialog box and click **OK**.



STEP 5 | **Close** the Application dialog box and **Commit** the change.

https://t.me/learningnets

Use HTTP Headers to Manage SaaS Application Access

Unsanctioned usage of SaaS applications can be a way for your users to transmit sensitive information outside of your network, usually by accessing a consumer version of an application. However, if you need to allow access to the enterprise version of these applications for specific individuals or organizations, then you can't block the SaaS application entirely.

You can use custom HTTP headers to disallow SaaS consumer accounts while allowing a specific enterprise account. Many SaaS applications allow or disallow access to applications based on information contained in specific HTTP headers. You can Create HTTP Header Insertion Entries using Predefined Types to manage access to popular SaaS applications, such as Google G Suite and Microsoft Office 365. Palo Alto Networks® uses content updates to maintain predefined rule sets specific to these applications, as well as to add new predefined rule sets.

You can also Create Custom HTTP Header Insertion Entries if you want to manage access to a SaaS application—that uses HTTP headers to limit service access—for which Palo Alto Networks has not provided a predefined set of rules.

Be aware that commercial SaaS applications always use SSL so decryption is necessary to perform HTTP header insertion. You can configure the firewall to decrypt traffic using SSL Forward Proxy decryption if traffic is not already decrypted by an upstream firewall.

You don't need a URL Filtering license to use this feature.

To understand how to use HTTP headers to manage SaaS applications, see the following:

- Understand SaaS Custom Headers
- Domains used by the Predefined SaaS Application Types
- Create HTTP Header Insertion Entries using Predefined Types
- Create Custom HTTP Header Insertion Entries

Understand SaaS Custom Headers

Before you begin, make sure you understand the custom HTTP headers you will use with the SaaS application you are managing. You need to understand what you can accomplish with these headers and the information you need to specify to accomplish your goals.

Be aware that SaaS applications that use custom headers do not always use them to control access to types of accounts. For example, Palo Alto Networks® provides predefined support for YouTube custom headers that determine whether network users can access restricted content.

You should also read the documentation for the SaaS application to which you want to control access so that you understand what headers you need to use for that application.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

The following limits apply to HTTP header insertion:

- *Header name character length: 100.*
- *Header value character length: 512.*

Be aware that some SaaS applications might define custom header names, or assign values to their custom headers, that exceed these limits. These situations should be rare, but if a SaaS application does exceed one or both of these character length limits, then your next-generation firewall can not successfully manage access to that SaaS application.

The following table lists the headers that you can use for the SaaS applications for which Palo Alto Networks provides predefined support; each header also includes a link to more information specific to that header.

Application	Headers	For More Information
Dropbox	`X-Dropbox-allowed-Team-Ids`	www.dropbox.com/help/business/network-control You can allow access to sanctioned Enterprise Dropbox accounts. This header's value is the business account's team ID, which you can obtain from the network control section of the Dropbox admin console. You must also enable this functionality from the same location. For details on managing this header, as well as how to enable your Dropbox clients so that you can decrypt their traffic, contact your Dropbox account representative.
Google G Suite	`X-GooGApps-Allowed-Domains`	support.google.com/a/answer/1668854?hl=en You can allow access to specific Google accounts from your domain. The values that you give to this header are your domain and subdomains. To successfully insert headers for Google applications, you must also:

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Application	Headers	For More Information
		1. Create an SSL decryption profile that includes the following categories and URLs: • business-and-economy • computer-and-internet-info • content-delivery-networks • internet-communications-and-telephony • low-risk • online-storage-and-backup • search-engine • web-based-email • drive.google.com • *.google.com • *.googleusercontent.com • *.gstatic.com 2. HTTP header insertion is not currently supported for HTTP/2. To insert headers, downgrade HTTP/2 connections to HTTP/1.1 using the **Strip ALPN** feature in the appropriate decryption profile. For more information, see App-ID and HTTP/2 Inspection. 3. Create rules to block the Quick UDP Internet Connections (QUIC) App-ID and place them at the top of your security policy because the firewall does not support header insertion for this protocol. When you do, the app reverts to using HTTP/2 over TLS, which the firewall handles in the previous step.
Microsoft Office 365	`Restrict-Access-To-Tenants` `Restrict-Access-Context`	docs.microsoft.com/en-us/azure/active-directory/active-directory-tenant-restrictions You provide `Restrict-Access-To-Tenants` with a list of tenants you want to allow your users to access. You can use any domain that is registered with a tenant to identify the tenant in this list. You provide `Restrict-Access-Context` with the directory ID that is setting the tenant restriction. You can find your directory ID in the Azure portal. Sign in as an administrator,

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Application	Headers	For More Information
		select **Azure Active Directory**, then select **Properties**.
YouTube	`YouTube-Restrict`	support.google.com/a/answer/6214622?hl=en You provide this header with information on the type of videos you want your users to be able to view. You can specify either a **Strict** or **Moderate** setting. See support.google.com/a/answer/6212415 for details on these different settings.

Domains used by the Predefined SaaS Application Types

SaaS applications use HTTPS so, to insert custom headers into this traffic, custom headers must be decrypted. If you use the forward-proxy decryption available on the firewall to decrypt custom headers, you must identify the specific HTTPS traffic you want to decrypt by identifying the domains associated with the traffic. The following table identifies the relevant domains for each of the SaaS applications for which Palo Alto Networks® has provided predefined rules.

Application	Domains
Dropbox	`*.dropbox.com`
G Suite	`*.google.com` `gmail.com`

Microsoft Office 365	`login.microsoftonline.com`
	`login.microsoft.com`
	`login.windows.net`

YouTube	`www.youtube.com`
	`m.youtube.com`
	`youtubei.googleapis.com`
	`youtube.googleapis.com`
	`www.youtube-nocookie.com`

©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Create HTTP Header Insertion Entries using Predefined Types

STEP 1 | If there are no upstream devices already decrypting HTTPS traffic, configure Decryption using Configure SSL Forward Proxy.

> *If you are configuring SSL decryption for Dropbox, then you must also configure your Dropbox clients to allow SSL traffic. These procedures are specific and private to Dropbox — to obtain these procedures, contact your Dropbox account representative.*

1. **Add** a Custom URL Category for the SaaS application you are managing (**Objects** > **Custom Objects** > **URL Category**).
2. Specify a **Name** for the category.
3. **Add** the domains specific to the SaaS application you are managing or for which you want to insert the username and domain in the headers. See Domains used by the Predefined SaaS Application Types for a list of the domains that you use for each of the predefined SaaS applications. See Insert Username in HTTP Headers for more information on configuring the firewall to include the username and domain in the HTTP headers.

   Each domain name can be up to 254 characters and you can identify a maximum of 50 domains for each entry. The domain list supports wildcards (for example, **`*.example.com`**). As a best practice, do not nest wildcards (for example, **`*.*.*`**) and do not overlap domains within the same URL profile.
4. For SaaS application management, Create a Decryption Policy Rule and, as you follow this procedure, configure the following:
   - In the **Service/URL Category** tab, **Add** the **URL Category** that you created in the previous step.
   - In the **Options** tab, make sure the **Action** is set to **Decrypt** and that the **Type** is set to **SSL Forward Proxy**.

STEP 2 | Edit or add a URL filtering profile.

STEP 3 | Select **HTTP Header Insertion** in the **URL Filtering Profile** dialog.

STEP 4 | **Add** an entry.

1. Specify a **Name** (up to 100 characters) for this entry.
2. Select a predefined **Type**.

   This populates the **Domains** and **Headers** lists.
3. For each **Header**, enter a **Value**.
4. (Optional) Select **Log** to enable logging of insertion activity for the headers.

   Allowed traffic is not logged, so header insertions are not logged for allowed traffic.
5. Click **OK** to save your changes.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 5 | **Add** or edit a Security Policy rule (**Policies** > **Security**) to include the HTTP header insertion URL filtering profile.

- For SaaS application management, allow users to access the SaaS application for which you are configuring this header insertion rule.
- To include the username and domain in the HTTP headers, apply the URL filtering profile to the security policy rule for HTTP or HTTPS traffic.
  1. Choose the URL filtering profile (**Actions** > **URL Filtering**) that you edited or created in Step 2.
  2. Click **OK** to save and then **Commit** your changes.

STEP 6 | Verify that the firewall correctly inserts the header.

- For Saas application management, from an endpoint, confirm that access to the SaaS application is working in the way you expect.
  1. Try to access an account or content that you expect to be able to access. If you cannot access the SaaS account or content, then the configuration is not working.
  2. Try to access an account or content that you expect will be blocked. If you can access the SaaS account or content, then the configuration is not working.
  3. If both of the previous steps work as expected, then you can View Logs (if you configured logging in step 4.4) and you should see the recorded HTTP header insertion activity.

Create Custom HTTP Header Insertion Entries

STEP 1 | If there are no upstream devices already decrypting HTTPS traffic, configure SSL Forward Proxy.

1. **Add** a custom URL category for the SaaS application you are managing (**Objects** > **Custom Objects** > **URL Category**).
2. Specify a **Name** for the category.
3. **Add** the domains specific to the SaaS application you are managing.
4. Create a Decryption Policy Rule and, as you follow this procedure, configure the following:
   - In the **Service/URL Category** tab, **Add** the **URL Category** that you created in the previous step.
   - In the **Options** tab, make sure the **Action** is set to **Decrypt** and that the **Type** is set to **SSL Forward Proxy**.

STEP 2 | Edit or create a URL Filtering profile.

STEP 3 | Select **HTTP Header Insertion** in the URL Filtering Profile dialog.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 4 | **Add** an entry.

1. Specify a **Name** for this entry.
2. Select **Custom** as the **Type**.
3. **Add** domains to the **Domains** list.

   You can add up to 50 domains and each domain name can have up to 256 characters; wildcards are supported (for example, *.example.com).

   HTTP header insertion occurs when a domain in this list matches the domain in the Host header of the HTTP request.
4. **Add** headers to the **Headers** list.

   You can add up to 5 headers, and each header can have up to 100 characters but cannot contain any spaces.
5. For each header, enter a **Value**.

   Each header value can have up to 512 characters.
6. (Optional) **Log** insertion activity for the headers.
7. Click **OK** to save your changes.

STEP 5 | **Add** or edit a Security policy rule (**Policies** > **Security**) that allows users to access the SaaS application for which you are configuring this header insertion rule.

1. Choose the URL Filtering profile (**Actions** > **URL Filtering**) that you edited or created in Step 2.
2. Click **OK** to save and then **Commit** your changes.

STEP 6 | Verify that access to the SaaS application is working in the way you expect. From an endpoint that is connected to your network:

1. Try to access an account or content that you expect to be able to access. If you cannot access the SaaS account or content, then the configuration is not working.
2. Try to access an account or content that you expect will be blocked. If you can access the SaaS account or content, then the configuration is not working.
3. If both of the previous steps work as expected, then you can View Logs (if you configured logging in step 4.6) and you should see the recorded HTTP header insertion activity.

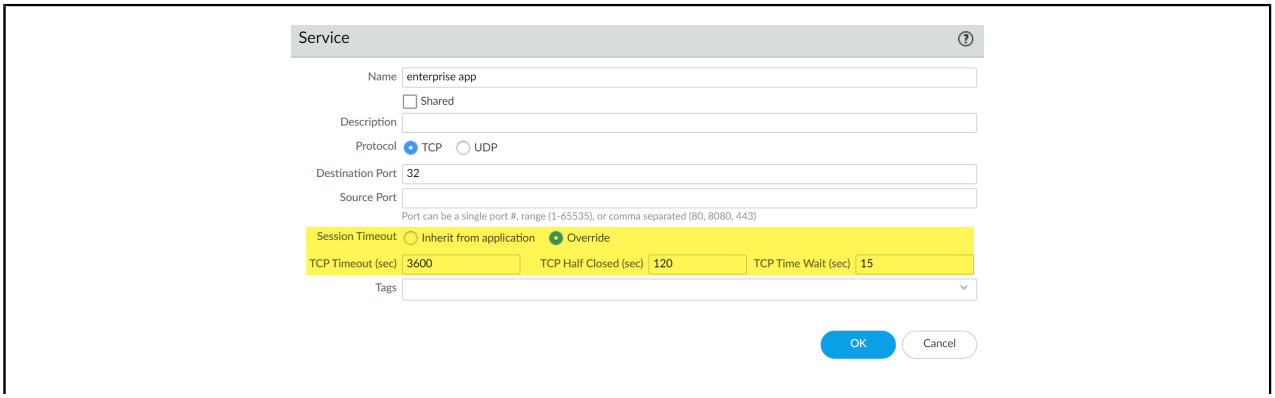
 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Maintain Custom Timeouts for Data Center Applications

Easily maintain custom timeouts for applications as you move from a port-based policy to an application-based policy. Use this method to maintain custom timeouts instead of overriding App-ID (losing application visibility) or creating a custom App-ID (expending time and research).

To get started, configure custom timeout settings as part of a service object:



Then add the service object in a policy rule to apply the custom timeouts to the application(s) the rule enforces.

The following steps describe how apply custom timeouts to applications; to apply custom timeouts to user groups, you can follow the same steps but just make sure to add the service object to the security policy rule that enforces the users to whom you want the timeout to apply.

STEP 1 | Select **Objects** > **Services** to add or modify a service object.

You can also create service objects as you are defining match criteria for a security policy rule: select **Policies** > **Security** > **Service/URL Category** and **Add** a new Service object to apply to the application traffic the rule governs.

STEP 2 | Select the protocol for the service to use (TCP or UDP).

STEP 3 | Enter the destination port number or a range of port numbers used by the service.

STEP 4 | Define the session timeout for the service.

- **Inherit from application** (default)—No service-based timeouts are applied; instead, apply the application timeout.
- **Override**—Define a custom session timeout for the service.

STEP 5 | If you chose to override the application timeout and define a custom session timeout, continue to:

- Enter a **TCP Timeout** value to set the Maximum length of time in seconds that a TCP session can remain open after data transmission has started. When this time expires, the session closes. The value range is 1 - 604800, and the default value is 3600 seconds.
- Enter a **TCP Half Closed** value to set the maximum length of time in seconds that a session remains in the session table between receiving the first FIN packet and receiving

https://t.me/learningnets

the second FIN packet or RST packet. If the timer expires, the session closes. The value range is 1 - 604800, and the default value is 120 seconds.

- Enter a **TCP Wait Time** value to set the maximum length of time in seconds that a session remains in the session table after receiving the second FIN packet or a RST packet. When the timer expires, the session closes. The value range is 1 - 600, and the default value is 15 seconds.

STEP 6 | Click **OK** to save the service object.

STEP 7 | Select **Policies** > **Security** and **Add** or modify a policy rule to govern the application traffic you want to control.

STEP 8 | Select **Service/URL Category** and **Add** the service object you just created to the security policy rule.

STEP 9 | Click **OK** and **Commit** your changes.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

paloalto TECHDOCS

Device-ID

- Device-ID Overview
- Prepare to Deploy Device-ID
- Configure Device-ID
- Manage Device-ID
- CLI Commands for Device-ID

https://t.me/learningnets

Device-ID Overview

Whether or not your environment supports a “Bring Your Own Device” (BYOD) policy, you likely already have a large number of devices in your network; maybe even more than you realize. Combined with the need for scalability as the number of users and their accompanying devices on your network increases, not to mention the growing infrastructure of the Internet of Things (IoT), this presents a constantly growing area of risk with many possibilities for exploitation by malicious users. Additionally, once you identify these devices, how do you secure them from vulnerabilities such as outdated operating software? Using Device-ID™ on your firewall or to push policy from Panorama, you can get device context for events on your network, obtain policy rule recommendations for those devices, write policies based on devices, and enforce Security policy based on the recommendations.

Similar to how User-ID provides user-based policy and App-ID provides app-based policy, Device-ID provides policy rules that are based on a device, regardless of changes to its IP address or location. By providing traceability for devices and associating network events with specific devices, Device-ID allows you to gain context for how events relate to devices and write policies that are associated with devices, instead of users, locations, or IP addresses, which can change over time. You can use Device-ID in Security, Decryption, Quality of Service (QoS) and Authentication policies.

Device-ID requires an IoT Security license, a Cortex Data Lake (CDL) license, and the device certificate.

If you use PAN-OS version 8.1.0 through PAN-OS 9.1.x on a firewall, the IoT Security license provides device classification, behavior analysis, and threat analysis for your devices. If you use PAN-OS 10.1 or later, you can use Device-ID to obtain IP address-to-device mappings to view device context for network events, use IoT Security to obtain policy rule recommendations for these devices, and gain visibility for devices in reports and the ACC.

You can create a device-based Security policy on any Panorama or firewall that uses PAN-OS version 10.1 or later. To enforce the Security policy, the device must have a valid IoT Security license.

To identify and classify devices, the IoT Security app uses metadata from logs, network protocols, and sessions on the firewall. This does not include private or sensitive information or data that is not relevant for device identification. Metadata also forms the basis of the expected behavior for the device, which then establishes the criteria for the policy rule recommendation that defines what traffic and protocols to allow for that device.

When a firewall imports security policy rule recommendations and IP address-to-device mappings from IoT Security, the firewall sends its device certificate to an edge server to authenticate itself. The edge server authenticates itself to the firewall by sending its own certificate. The firewall uses Online Certificate Status Protocol (OCSP) to validate the server’s certificate by checking it against the following sites using HTTP on TCP port 80:

- ocsp.int-x3.letsencrypt.org
- isrg.trustid.ocscp.identrust.com
- crl.identrust.com

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Panorama performs the same check to validate the edge server's certificate when Panorama imports policy rule recommendations from IoT Security.

After IoT Security identifies and classifies the devices in your network using the Palo Alto Networks firewalls already in your network, so you don't have to implement new devices or third-party solutions, Device-ID can leverage this data to match devices with policy rules and provide device context for network events. Through the visibility that the firewall or Panorama provides for traffic, apps, users, devices, and threats, you can instantly trace network events back to individual devices and obtain Security policy rule recommendations for securing those devices.

All firewall platforms that support PAN-OS 10.1 support Device-ID and the IoT Security app with the exception of the VM-50 series, the VM-200, the CN series, and Prisma Access.

There are six levels of classification (also known as attributes) for devices:

Attribute	Example
Category	`ATM Machine`; `3D Printer`
Profile	`Palo Alto Networks Device`
Model	`iPad`
OS Version	`iOS 9.9.3`
OS Family	`Android`; `iOS`
Vendor	`ASUS`; `Philips`

To obtain policy rule recommendations for devices in your network, the firewall observes traffic to generate Enhanced Application logs (EALs). The firewall then forwards the EALs to the Cortex Data Lake (CDL) for processing. The IoT Security app on the hub receives logs from CDL for analysis, provides IP address-to-device mappings, and generates the latest policy rule recommendations for your devices. Using the IoT Security app, you can review these policy rule recommendations and create a Security policy for these devices. After you activate the policy rules in the IoT Security app, import them to the firewall or Panorama and commit your Security policy.

The firewall must be able to observe DHCP broadcast and unicast traffic on your network to identify devices. The more traffic the firewall can observe, the more accurate the policy rule recommendations are for the device and the more rapid and accurate the IP address-to-device mappings are for the device. When a device sends DHCP traffic to obtain an IP address, the firewall observes this type of request, it generates EALs to send to the Cortex Data Lake for processing and then analysis by IoT Security.

To observe traffic on an L2 interface, you must configure a VLAN for that interface. By allowing the firewall to treat the interface as an L3 interface for a DHCP relay, it can observe the DHCP broadcast traffic without impacting traffic or performance.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Because the firewall needs to both detect the devices based on their traffic and then enforce Security policy for those devices, the firewall acts as both a *sensor* to collect metadata from devices and an *enforcer* by enforcing your Security policy for the devices. The IoT Security app automatically detects new devices as soon as they send DHCP traffic and can identify 95% of devices within the first week.

Each application has an individual recommendation that you import to the firewall or Panorama as a rule. When you import the recommendation, the firewall or Panorama creates at least two objects to define the device behavior from the recommendation:

- A source device object that identifies the device where the traffic originates
- One or more destination objects that identify the permitted destinations for the traffic, which can be a device, IP address, or Fully Qualified Domain Name (FQDN)

If any of the device objects already exist on the firewall or Panorama, the firewall or Panorama updates the device object instead of creating a new device objects. You can use these device objects in Security, authentication, decryption, and Quality of Service (QoS) policies.

Additionally, the firewall assigns two tags to each rule:

- One that identifies the source device, including the category (such as `NetworkDevice - TrendNet`).
- One that indicates that the rule is an IoT policy rule recommendation (`IoTSecurityRecommended`).



Because the tags that the firewall assigns to the rule are the only way to restore your mappings if they become out of sync, do not edit or remove the tags.

For optimal deployment and operation of Device-ID, we recommend the following best practices:

- Deploy Device-ID on firewalls that are centrally located in your network. For example, if you have a large environment, deploy Device-ID on a firewall that is upstream from the IP address management (IPAM) device. If you have a small environment, deploy Device-ID on a firewall that is acting as a DHCP server.
- During initial deployment, allow Device-ID to collect metadata from your network for at least fourteen days. If devices are not active daily, the identification process may take longer.
- Write device-based policy in order of your most critical devices to least critical. Prioritize by:
  1. Class (secure networked devices first)
  2. Critical devices (such as servers or MRI machines)
  3. Environment-specific devices (such as fire alarms and badge readers)
  4. Consumer-facing IoT devices (such as a smart watch or smart speaker)
- Enable Device-ID on a per-zone basis for internal zones only.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Prepare to Deploy Device-ID

To prepare your network for Device-ID deployment, complete the following predeployment tasks to enable your firewall to generate and send Enhanced Application logs (EALs) to the Cortex Data Lake for processing and analysis by IoT Security for policy rule recommendation generation.

STEP 1 | If you have not already done so, install the device certificate on your firewall or Panorama.

> *If you use Panorama to manage multiple firewalls, Palo Alto Networks strongly recommends upgrading all firewalls in your Device-ID deployment to PAN-OS 10.0 or a later version. If you create a rule that uses* ***Device*** *as a match criteria and Panorama pushes the rule to a firewall that uses PAN-OS 9.1 or an earlier version, the firewall omits the* ***Device*** *match criteria because it is not supported, which may cause issues with policy rule traffic matching.*

STEP 2 | Activate your Cortex Data Lake (CDL) instance and connect your firewall to the instance.

1. Activate a Cortex Data Lake instance.
2. Connect your firewall to Cortex Data Lake.

STEP 3 | (L2 interfaces only) Create a VLAN interface for each L2 interface so the firewall can observe the DHCP broadcast traffic.

STEP 4 | (Optional) Configure a service route to allow the necessary traffic for Device-ID and IoT Security.

By default, the firewall uses the management interface. To use a different interface, complete the following steps.

1. Select **Device** > **Setup** > **Services** then select **Service Route Configuration**.
2. **Customize** a service route.
3. Select the **IPv4** protocol.

   > *Device-ID and IoT Security do not support IPv6.*

4. Select **Data Services** in the Service column.
5. Select a **Source Interface** and **Source Address**.
6. Click **OK** twice.

STEP 5 | Use App-IDs to allow the necessary traffic for Device-ID and IoT Security.

Purpose	App-ID
Retrieve policy rule recommendations and allow traffic between the IoT Security app and your firewall or Panorama.	**paloalto-iot-security**
Allow traffic for all EALs and all session logs.	**paloalto-logging-service**

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Purpose	App-ID
Retrieve IoT Security dynamic updates and Device Dictionary updates.	**paloalto-updates**

If you have a third-party firewall between a Palo Alto Networks next-generation firewall using Device-ID and the internet, verify that the next-generation firewall can access `iot.services-edge.paloaltonetworks.com:443` *if it's in the Americas region and* `eu.iot.services-edge.paloaltonetworks.com:443` *if it's in the EU region.*

STEP 6 | If there's a third-party firewall between the internet and Panorama and Panorama-managed next-generation firewalls, make sure it allows the necessary traffic for Device-ID and IoT Security.

Purpose	Address	TCP Port
(PAN-OS versions 10.0.3 and later) Receive the regional FQDN allowing next-generation firewalls to retrieve IP address-to-device mappings and policy rule recommendations from IoT Security.	**enforcer.iot.services-edge.paloaltonetworks.com**	443
(PAN-OS versions 10.0.0—10.0.2 and later) Let next-generation firewalls receive policy rule recommendations and IP address-to-device mappings from IoT Security.	Americas region **iot.services-edge.paloaltonetworks.com** EU region **eu.iot.services-edge.paloaltonetworks.com**	443
(PAN-OS versions 10.0.0 and later) Let next-generation firewalls download device dictionary files from the update server.	**updates.paloaltonetworks.com**	443
(PAN-OS versions 10.0.0 and later) Let Panorama send queries for logs to the logging service.	Americas region **iot.services-edge.paloaltonetworks.com** EU region	443

©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Purpose	Address	TCP Port
	eu.iot.services-edge.paloaltonetworks.com	
(IoT Security subscription + Cortex Data Lake) Forward logs to Cortex Data Lake.	See TCP Ports and FQDNs Required for Cortex Data Lake.	

PAN-OS versions 10.0.0 - 10.0.2 connect to the edge services FQDN in the Americas region by default (`iot.services-edge.paloaltonetworks.com`). For firewalls running these PAN-OS versions to connect to the FQDN in the EU region (`eu.iot.services-edge.paloaltonetworks.com`), you must manually configure it. For PAN-OS versions 10.0.3 and later, firewalls automatically discover the correct FQDN to use based on the region set during the IoT Security onboarding process. There is no need to set it manually.

STEP 7 | If there's a third-party firewall between the internet and next-generation firewalls (without Panorama), make sure it allows the necessary traffic for Device-ID and IoT Security.

Purpose	Address	TCP Port
(PAN-OS versions 10.0.3 and later) Receive the regional FQDN allowing next-generation firewalls to retrieve IP address-to-device mappings and policy rule recommendations from IoT Security.	**enforcer.iot.services-edge.paloaltonetworks.com**	443
(PAN-OS versions 10.0.0—10.0.2) Let next-generation firewalls receive policy rule recommendations and IP address-to-device mappings from IoT Security.	Americas region **iot.services-edge.paloaltonetworks.com** EU region **eu.iot.services-edge.paloaltonetworks.com**	443
(PAN-OS versions 10.0.0 and later) Let next-generation firewalls download device dictionary files from the update server.	**updates.paloaltonetworks.com**	443
(IoT Security subscription + Cortex Data Lake) Forward logs to Cortex Data Lake.	See TCP Ports and FQDNs Required for Cortex Data Lake.	

©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 8 | Configure your firewall to observe and generate logs for DHCP traffic then forward the logs for processing and analysis by IoT Security.

- If the firewall is acting as a DHCP server:
  1. Enable Enhanced Application logging.
  2. Create a log forwarding profile to forward the logs to the CDL for processing.
  3. (Not supported on the PA-3200, PA-5200, PA-5450, or PA-7000) Enable the **DHCP Broadcast Session** option (**Device** > **Setup** > **Session** > **Session Settings**).
  4. Create a Security policy rule to allow **dhcp** as the **Application** type.
- If the firewall is not a DHCP server, configure an interface as a DHCP relay agent so that the firewall can generate EALs for the DHCP traffic it receives from clients.
- If your DHCP server is on the same network segment as the interface your firewall, deploy a virtual wire interface in front of the DHCP server to ensure the firewall generates EALs for all packets in the initial DHCP exchange with minimal performance impact.
  1. Configure a virtual wire interface with corresponding zones and enable the **Multicast Firewalling** option (**Network** > **Virtual Wires** > **Add**).
  2. Configure a rule to allow DHCP traffic to and from the DHCP server between the virtual wire zones. The policy must allow all existing traffic that the server currently observes and use the same log forwarding profile as the rest of your rules.
  3. To allow the DHCP servers to check if an IP address is active before assigning it as a lease to a new request, configure a rule to allow pings from the DHCP server to the rest of the subnet.
  4. Configure a rule to allow all other traffic to and from the DHCP server that does not forward logs for traffic matches.
  5. Configure the DHCP server host to use the first virtual wire interface and the network switch to use the second virtual wire interface. To minimize cabling, you can use an isolated VLAN in the switching infrastructure instead of connecting the DHCP server host directly to the firewall.
- If you want to use a tap interface to gain visibility into DHCP traffic that the firewall doesn't usually observe due to the current configuration or topology of the network, use the following configuration as a best practice.
  1. Configure a tap interface and corresponding zone.
  2. Configure a rule to match DHCP traffic that uses the same log forwarding profile as the rest of your rules.
  3. To minimize the session load on the firewall, configure a rule to drop all other traffic.
  4. Connect the tap interface to the port mirror on the network switch.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 9 | Add session log types to the log forwarding profile.

If there are no existing entries in the log forwarding profile, selecting the **Enable enhanced application logging to Cortex Data Lake (including traffic and url logs)** option adds all logs types.

1. **Add** a new profile and enter a name.
2. Select **traffic** as the **Log type**.
3. Select **All logs** as the **Filter**.
4. Select the **Cortex Data Lake** option.
5. Click **OK**.
6. Repeat substeps 1-5 for the **threat** and, if you have a subscription, **wildfire** log types.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Configure Device-ID

Complete the following tasks to import the IP address-to-device mappings and policy rule recommendations from IoT Security to your firewall or Panorama.



If you use Panorama to manage multiple firewalls, Palo Alto Networks strongly recommends upgrading all firewalls in your Device-ID deployment to PAN-OS 10.0 or a later version. If you create a rule that uses ***Device*** *as a match criteria and Panorama pushes the rule to a firewall that uses PAN-OS 9.1 or an earlier version, the firewall omits the* ***Device*** *match criteria because it is not supported, which may cause issues with policy rule traffic matching.*

STEP 1 | Activate your IoT Security license on the hub.

1. Follow the instructions that you received in your email to activate your IoT Security license.
2. Initialize your IoT Security app. For more information, refer to Get Started with IoT Security and the IoT Security Best Practices.
3. Apply the license to the firewalls you want to use to enforce the IoT Security policy.
4. Refresh your license on the firewall or Panorama.

STEP 2 | Define your IoT Security policy on the IoT Security app.

1. On the IoT Security app, select the source device object.
2. **Create** a new set of policy rules for the source device object.

   For more information about creating security policies with the IoT Security app, refer to Recommend Security Policies.
3. **Activate** the policy rules to confirm your changes.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 3 | Import the IP address-to-device mappings and policy rule recommendations to the firewall or Panorama.

1. Import the policy rule recommendation.
   - On the firewall, select **Device** > **Policy Recommendation** > **IoT**.
   - For Panorama, select **Panorama** > **Policy Recommendation** > **IoT** then push the policy rules to the firewalls that Panorama manages.

   After you push the policy to the firewalls, you must Sync Policy Rules on the firewalls to create the policy rule recommendation-to-policy rule mapping.

   When you select Policy Recommendation, the firewall or Panorama communicates with IoT Security to obtain the latest policy rule recommendations. The policy rule recommendations are not cached on the firewall or Panorama.

   Because IoT Security creates the policy rule recommendation using the trusted behavior for the device, the default action for the rule is allow.
2. Select the **Source Device Profile**.
3. Verify that the **Destination Device Profile** and permitted **Applications** are correct.
4. Select **Import Policy Rules** to import the policy rules.
5. (Panorama only) Select the **Location** of the device group where you want to import the policy rules.
6. Enter a **Name** for the policy rules.
7. (Panorama only) Select the **Destination Type** (**Pre-Rulebase** or **Post-Rulebase**).
8. Select **After Rule** to define the placement of the rule in the rulebase.
   - **No Rule Selection**—Places the rule at the top of the rulebase.
   - **Default One**—Places the rule after the listed rule.

   In your Security policy, Device-ID rules must precede any existing rules that apply to the devices.
9. Repeat this process for each policy rule recommendation to create rules to allow access for each device object to the necessary destination(s).
10. Click **OK** and **Commit** your changes.

STEP 4 | Enable Device-ID in each zone where you want to use Device-ID to detect devices and enforce your Security policy.

By default, Device-ID maps all subnetworks in the zones where you enable it. You can modify which subnetworks Device-ID maps in the **Include List** and **Exclude List**.

As a best practice, enable Device-ID in the source zone to detect devices and enforce security policy. Only enable Device-ID for internal zones.

1. Select **Network** > **Zones**.
2. Select the zone where you want to enable Device-ID.
3. **Enable Device Identification** then click **OK**.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 5 | **Commit** your changes.

STEP 6 | Verify your Security policy is correct.

1. Select **Policies** then select the rule you created from the policy rule recommendation.

   IoT Security assigns a **Description** that contains the source device object and **Tags** to identify the source device object and that this rule is a recommendation from IoT Security.

   Device object names must be unique.

2. Select the **Source** tab, then verify the **Source Device Profile**.
3. Select the **Destination** tab and verify the **Destination Device Profile**.
4. Select the **Application** tab and verify the **Applications**.
5. Select the **Actions** tab and verify the **Action** (default is **Allow**).
6. Use Explore to verify CDL receives your logs and review which logs CDL receives.

STEP 7 | Create custom device objects for any devices that do not have IoT Security policy rule recommendations.

For example, you cannot secure devices such as laptops and smartphones using policy rule recommendations, so you must manually create device objects for these types of devices to use in your Security policy. For more information on custom device objects, see Manage Device-ID.

STEP 8 | Use the device objects to enforce policy rules and to monitor and identify potential issues.

The following list includes some example use cases for device objects.

- Use source device objects and destination device objects in Security, Authentication, QoS, & decryption policies.
- Use the decryption log to identify failures and which assets are the most critical to decrypt.
- View device object activity in ACC to track new devices and device behavior.
- Use device objects to create a custom report (for example, for incident reports or audits).

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Manage Device-ID

Perform the following tasks as needed to ensure your policy rule recommendations and device objects are current or to restore policy rule recommendation mappings.

STEP 1 | Update your policy rule recommendation whenever the **New Updates Available** column displays **Yes** for that recommendation.

As devices gain new capabilities, IoT Security updates the policy rule recommendations to advise what additional traffic or protocols the firewall or Panorama should allow. Check IoT Security daily for updates and update your policy rule recommendations as soon as possible.

1. On the IoT Security app, **Edit** the policy rules then click **Next**.
2. Select the new recommendation then click **Next**.
3. **Save** your changes.
4. On the firewall or Panorama, click **Import Policy Rules** then click **Yes** to confirm that you want to overwrite the current rule.

   This action overwrites the recommendation for the rule, not the rule itself.
5. (Panorama only) Repeat the previous step for all device groups.
6. **Commit** your changes.

STEP 2 | Review, update, and maintain the device objects in the Device Dictionary.

You must create device objects for any devices that do not have an IoT Security policy rule recommendation. For example, you cannot secure devices such as laptops and smartphones using IoT Security policy rule recommendations, so you must create device objects for these types of devices and use them in your Security policy to secure these devices.

1. Select **Objects** > **Devices**.
2. **Add** a device object.
3. **Browse** the list or **Search** using keywords.

   The search results can include multiple types of device object attributes (for example, both **Category** and **Profile**).
4. To add a custom device object, enter a **Name** and optionally a **Description** for the device object.

   Always use a unique name for each device object. Do not change the tags in the description for device objects from policy rule recommendations.
5. (Panorama only) Select the **Shared** option to make this device object available to other device groups.
6. Select the attributes for the device object (**Category**, **OS**, **Profile**, **Osfamily**, **Model**, and **Vendor**).
7. Click **OK** to confirm your changes.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 3 | In some cases (for example, if you restore a previous configuration), the policy rule recommendation-to-policy rule mappings may become out of sync. You must also sync the mappings on each firewall after you push the policy rules from Panorama to the firewalls that Panorama manages. To sync the mappings:

- On the firewall, select **Device** > **Policy Recommendation** > **IoT** > **Sync Policy Rules**
- For Panorama, select **Panorama** > **Policy Recommendation** > **IoT** > **Sync Policy Rules**.

The firewall or Panorama scans all of the rules in the rulebase to check for tags that identify a rule as an IoT Security policy rule recommendation, obtains the source device object information, and repopulates the local policy rule recommendation database.

STEP 4 | Delete any policy rule recommendations that are no longer needed.

If a policy rule recommendation no longer applies, you can remove the policy rule recommendation. You must also remove the rule for the policy rule recommendation to enforce the updated Security policy.

1. On the IoT Security app, select **Delete**.
2. Click **Mark as Removed** to select this recommendation for removal.
3. Remove the mapping.
   - On the firewall, select **Device** > **Policy Recommendation** > **IoT** > **Remove Policy Mapping**.
   - For Panorama, select **Device** > **Policy Recommendation** > **IoT** > **Remove Policy Mapping** then select the **Location** from which you want to remove the mapping.
4. Click **Yes** to confirm the mapping removal.
5. Select **Policies** > **Security**. For Panorama, select **Policies** > **Security** > **Pre-Rules/Post-Rules**.
6. Select the rule for the policy rule recommendation you want to remove then select **Delete**.
7. **Commit** your changes.

STEP 5 | Use CLI commands to troubleshoot any issues between the firewall and IoT Security.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

CLI Commands for Device-ID

Use the following CLI commands to view information for troubleshooting any issues between the firewall and IoT Security. In general, CLI commands that include **eal** show counters for outgoing data and CLI commands that include **icd** show counters for incoming data.

Example	Command
View Enhanced Application Logging (EAL) counters, such as the number of connections between the firewall and the Cortex Data Lake and the volume of the logs.	**show iot eal all**
View more details about the connection between the firewall and Cortex Data Lake.	**show iot eal conn**
View a summary of the EAL counters by plane (dataplane or management plane), such as the PAN-OS version and serial number.	**show iot eal dpi-eal**
View EAL counters by plane (dataplane or management plane) and by protocol.	**show iot eal dpi-stats all**
View EAL counters by protocol.	**show iot eal dpi-stats subtype dhcp|http**
View a summary of Host Information Profile (HIP) Match report counters.	**show iot eal hipreport-eal**
View EAL log response time counters.	**show iot eal response-time**
View details for the health of the connection to the edge service between the firewall and the IoT Security app and counters for the IP address-to-device mappings and policy rule recommendations.	**show iot icd statistics all**
View counters for the connection to the edge service.	**show iot icd statistics conn**
View counters for the IP address-to-device mappings.	**show iot icd statistics verdict**
View all IP address-to-device mappings on the firewall.	**show iot ip-device-mapping-mp all**
View the IP address-to-device mapping for a specific IP address.	**show iot ip-device-mapping-mp ip *IP-address***

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Example	Command
View a list of IP address-to-device mappings on the dataplane.	**show iot ip-device-mapping all**
Clear the IP address-to-device mappings on the management plane.	**debug iot clear-all type device**
Clear the IP address-to-device mappings on the dataplane.	**clear user-cache all**

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Threat Prevention

The Palo Alto Networks® next-generation firewall protects and defends your network from commodity threats and advanced persistent threats (APTs). The multi-pronged detection mechanisms of the firewall include a signature-based (IPS/Command and Control/Antivirus) approach, heuristics-based (bot detection) approach, sandbox-based (WildFire) approach, and Layer 7 protocol analysis-based (App-ID) approach.

Commodity threats are exploits that are less sophisticated and more easily detected and prevented using a combination of antivirus, anti-spyware, and vulnerability protection features along with URL filtering and Application identification capabilities on the firewall.

Advanced threats are perpetuated by organized cyber adversaries who use sophisticated attack vectors to target your network, most commonly for intellectual property theft and financial data theft. These threats are more evasive and require intelligent monitoring mechanisms for detailed host and network forensics on malware. The Palo Alto Networks next-generation firewall together with WildFire™ and Panorama™ provide a comprehensive solution that intercepts and breaks the attack chain and provides visibility to prevent security infringement on your network infrastructure—both mobile and virtualized.



After you implement your threat prevention configurations, Export Configuration Table Data to create a PDF or CSV report of your configurations to use for internal review or for auditing.

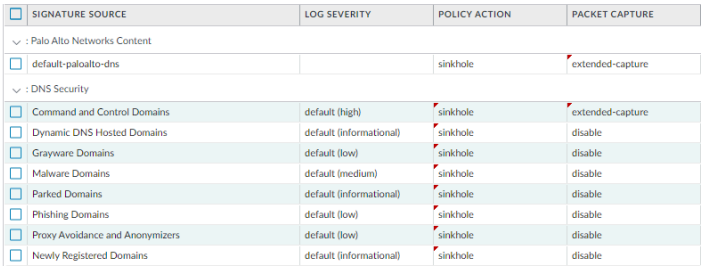
- Best Practices for Securing Your Network from Layer 4 and Layer 7 Evasions
- Set Up Antivirus, Anti-Spyware, and Vulnerability Protection
- DNS Security
- Use DNS Queries to Identify Infected Hosts on the Network
- Set Up Data Filtering
- Predefined Data Filtering Patterns
- Create a Data Filtering Profile
- WildFire Inline ML
- Set Up File Blocking
- Prevent Brute Force Attacks
- Customize the Action and Trigger Conditions for a Brute Force Signature
- Enable Evasion Signatures
- Monitor Blocked IP Addresses
- Threat Signature Categories
- Create Threat Exceptions
- Custom Signatures
- Learn More About and Assess Threats
- Share Threat Intelligence with Palo Alto Networks
- Threat Prevention Resources

https://t.me/learningnets

Best Practices for Securing Your Network from Layer 4 and Layer 7 Evasions

To monitor and protect your network from most Layer 4 and Layer 7 attacks, here are a few recommendations.

- Upgrade to the most current PAN-OS software version and content release version to ensure that you have the latest security updates. See Install Content and Software Updates.
- Enable DNS Security (requires a Threat Prevention and DNS Security subscription license) to sinkhole malicious DNS requests. Palo Alto Networks recommends using the following DNS Security category configuration settings in your Anti-Spyware profile:



SIGNATURE SOURCE	LOG SEVERITY	POLICY ACTION	PACKET CAPTURE
: Palo Alto Networks Content			
default-paloalto-dns		sinkhole	extended-capture
: DNS Security			
Command and Control Domains	default (high)	sinkhole	extended-capture
Dynamic DNS Hosted Domains	default (informational)	sinkhole	disable
Grayware Domains	default (low)	sinkhole	disable
Malware Domains	default (medium)	sinkhole	disable
Parked Domains	default (informational)	sinkhole	disable
Phishing Domains	default (low)	sinkhole	disable
Proxy Avoidance and Anonymizers	default (low)	sinkhole	disable
Newly Registered Domains	default (informational)	sinkhole	disable

  - For the log severity settings, use the default settings:
  - For the policy action, set all signature sources to **sinkhole**.
  - For packet capture, set Command and Control Domains to **extended-capture**. Leave all other categories to default settings.

  For more information on related anti-spyware settings, see Best Practice Internet Gateway Anti-Spyware Profile.
- Set up the firewall to act as a DNS proxy and enable evasion signatures:

  DNS proxy is not part of the firewall security policy engine; instead, it directs the firewall to resolve DNS hostnames, while maintaining domain to IP mapping, which is crucial for preventing TLS/HTTP evasion.

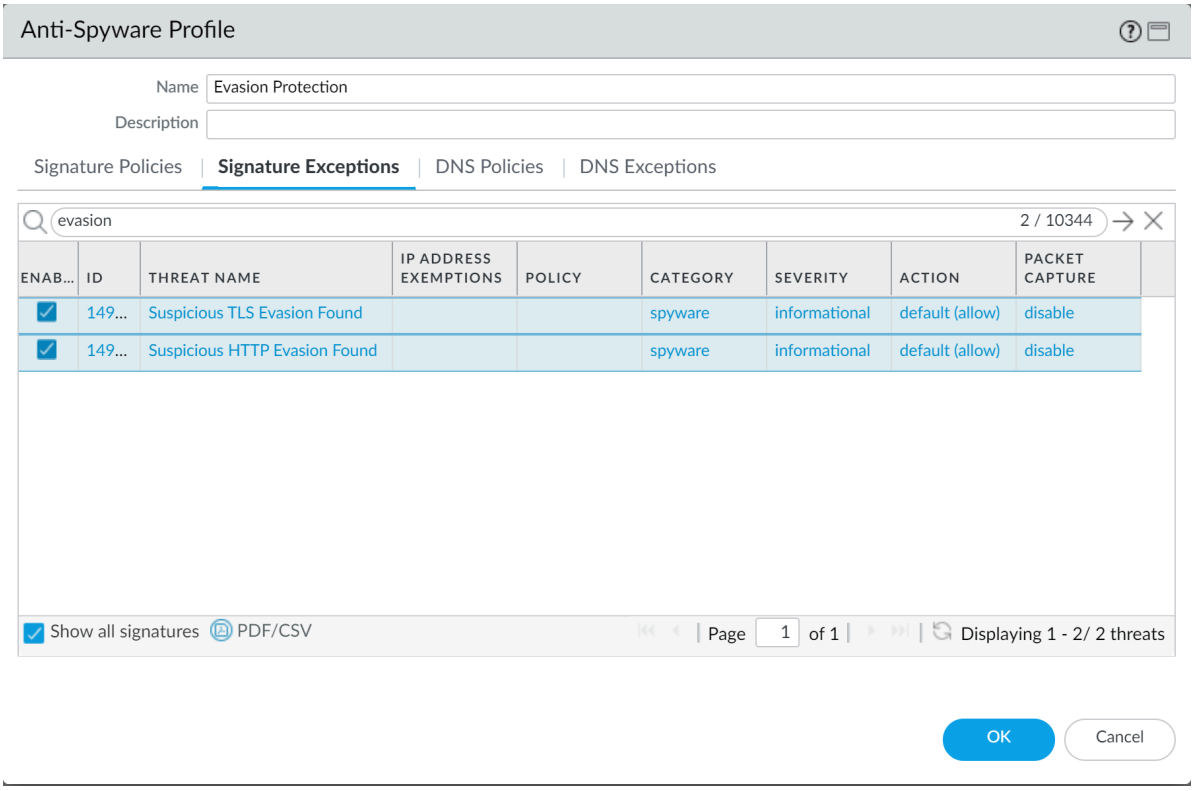
  - Configure a DNS Proxy Object.

    When acting as a DNS proxy, the firewall resolves DNS requests and caches hostname-to-IP address mappings to quickly and efficiently resolve future DNS queries.
  - Enable Evasion Signatures

    Evasion signatures that detect crafted HTTP or TLS requests can send alerts when clients connect to a domain other than the domain specified in the original DNS request. Make sure to configure DNS proxy before you enable evasion signatures. Without DNS proxy, evasion signatures can trigger alerts when a DNS server in the DNS load balancing configuration

https://t.me/learningnets

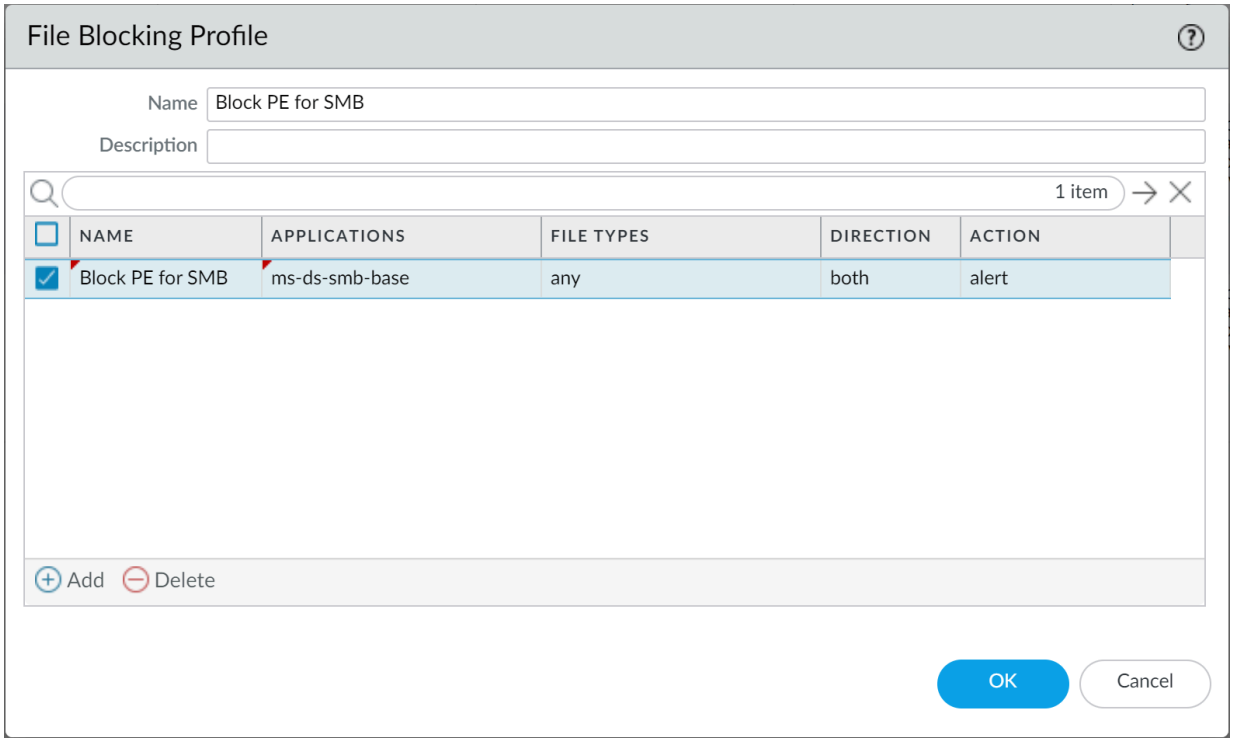
returns different IP addresses—for servers hosting identical resources—to the firewall and client in response to the same DNS request.



- For servers, create Security policy rules to allow only the application(s) that you sanction on each server. Verify that the standard port for the application matches the listening port on the server. For example, to ensure that only SMTP traffic is allowed to your email server, set the Application to **smtp** and set the Service to **application-default**. If your server uses only a subset of the standard ports (for example, if your SMTP server uses only port 587 while the SMTP application has standard ports defined as 25 and 587), create a new custom service that includes only port 587 and use that new service in your security policy rule instead of application-default. Additionally, make sure you restrict access to specific source and destinations zones and sets of IP addresses.
- Block all unknown applications and traffic using the Security policy. Typically, the only applications classified as unknown traffic are internal or custom applications on your network and potential threats. Unknown traffic can be either non-compliant applications or protocols that are anomalous or abnormal or it can be known applications that are using non-standard ports, both of which should be blocked. See Manage Custom or Unknown Applications.

https://t.me/learningnets

- Set Up File Blocking to block Portable Executable (PE) file types for internet-based SMB (Server Message Block) traffic from traversing trust to untrust zones (ms-ds-smb applications).

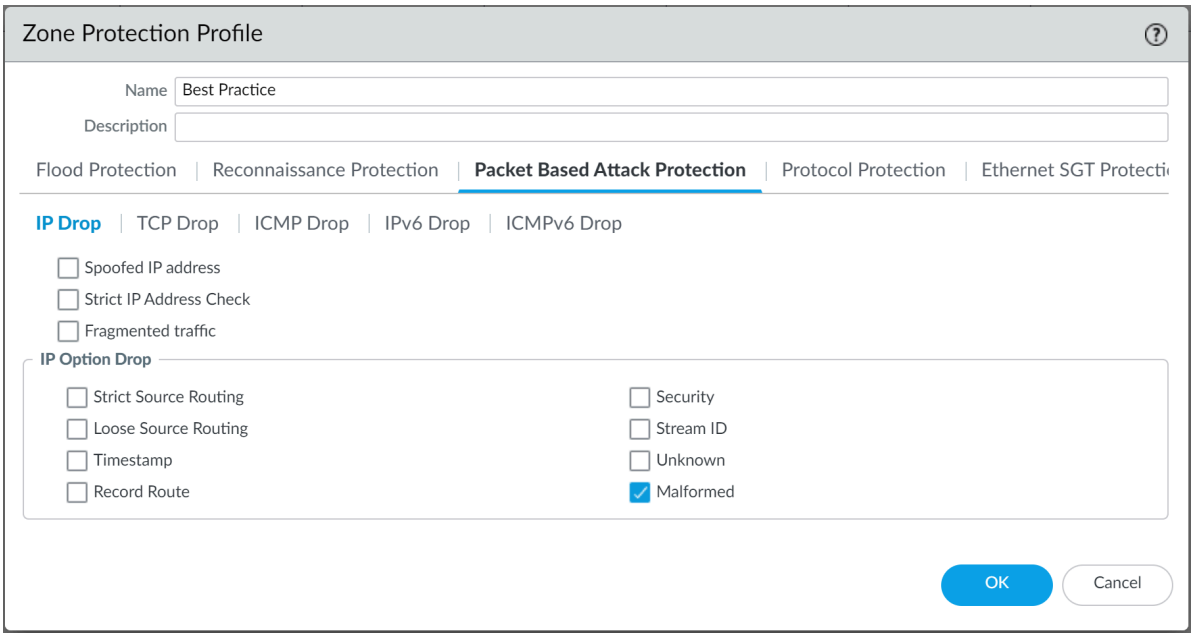


- Block malicious variants of PE (portable executables), ELF and MS Office files, and PowerShell and shell scripts in real-time. Enabling WildFire Inline ML allows you to dynamically analyze files using machine learning on the firewall. This additional layer of antivirus protection complements the WildFire-based signatures to provide extended coverage for files of which signatures do not already exist.

©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

- Create a Zone Protection profile that is configured to protect against packet-based attacks (**Network** > **Network Profiles** > **Zone Protection**):
  - Select the option to drop **Malformed** IP packets (**Packet Based Attack Protection** > **IP Drop**).



  - Enable the drop **Mismatched overlapping TCP segment** option (**Packet Based Attack Protection** > **TCP Drop**).

    By deliberately constructing connections with overlapping but different data in them, attackers attempt to cause misinterpretation of the intent of the connection and deliberately induce false positives or false negatives. Attackers also use IP spoofing and sequence number prediction to intercept a user's connection and inject their own data into that connection. Selecting the **Mismatched overlapping TCP segment** option specifies that PAN-OS discards frames with mismatched and overlapping data. Received segments are

https://t.me/learningnets

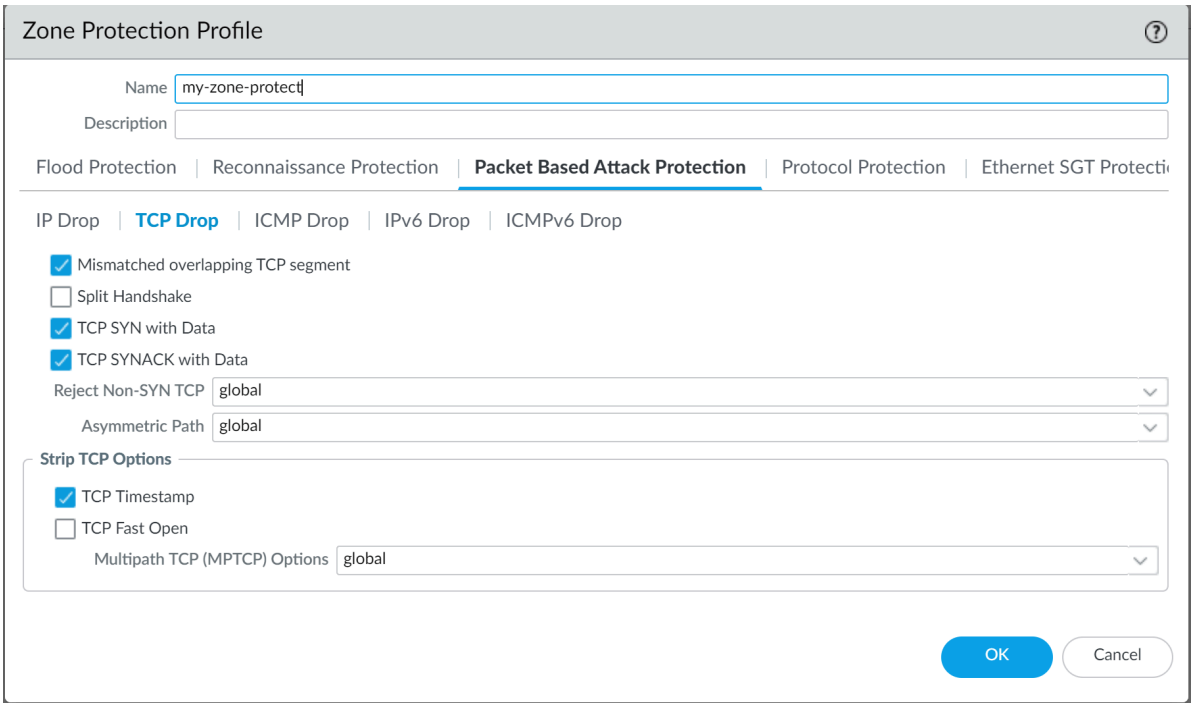
discarded when they are contained within another segment, when they overlap with part of another segment, or when they contain another complete segment.

- Enable the drop **TCP SYN with Data** and drop **TCP SYNACK with Data** options (**Packet Based Attack Protection** > **TCP Drop**).

  Dropping SYN and SYN-ACK packets that contain data in the payload during a three-way handshake increases security by blocking malware contained in the payload and preventing it from extracting unauthorized data before the TCP handshake is completed.

- Strip TCP timestamps from SYN packets before the firewall forwards the packet (**Packet Based Attack Protection** > **TCP Drop**).

  When you enable the **Strip TCP Options - TCP Timestamp** option, the TCP stack on both ends of the TCP connection will not support TCP timestamps. This prevents attacks that use different timestamps on multiple packets for the same sequence number.

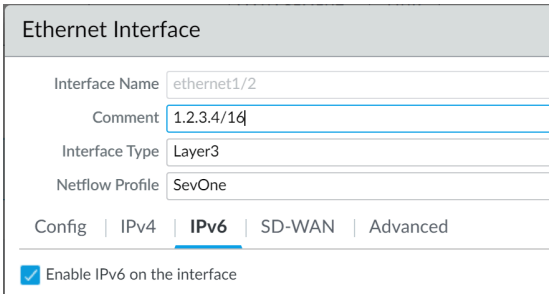


 ©2022 Palo Alto Networks, Inc.

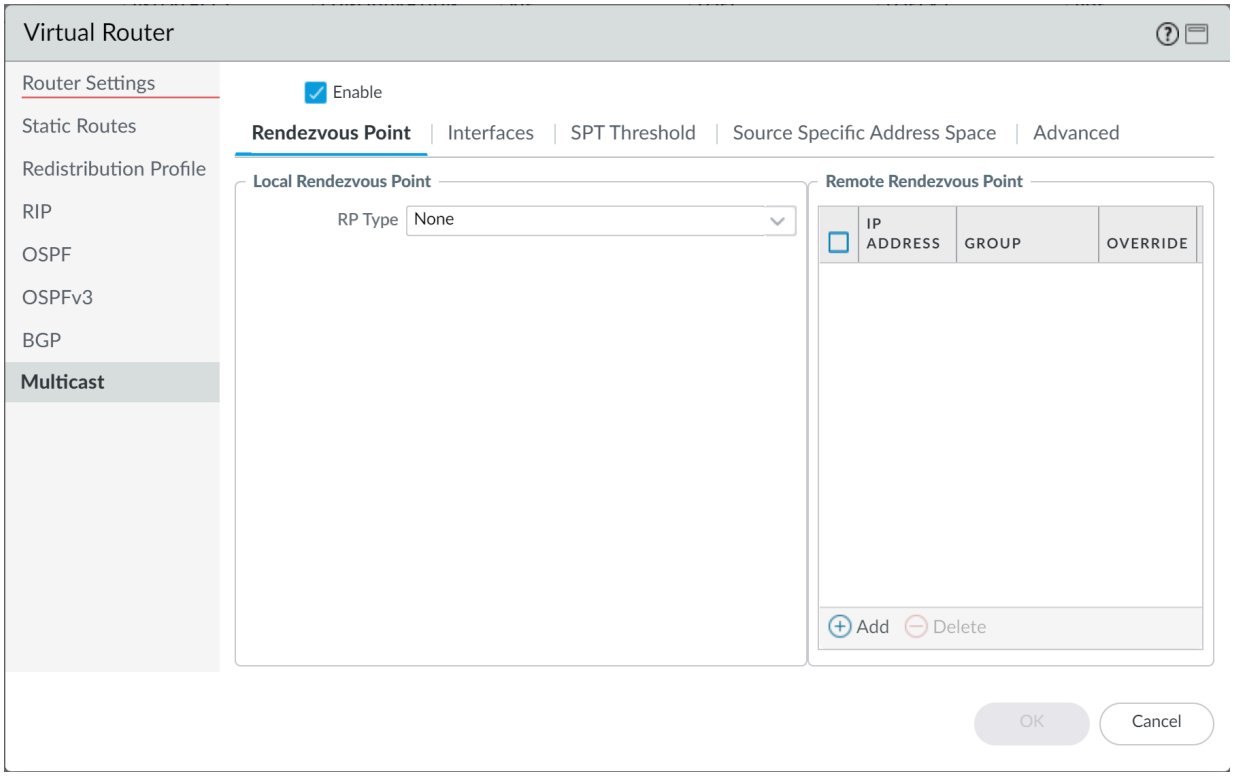
https://t.me/learningnets

- If you configure IPv6 addresses on your network hosts, be sure to enable support for IPv6 if not already enabled (**Network** > **Interfaces** > **Ethernet** > **IPv6**).

  Enabling support for IPv6 allows access to IPv6 hosts and also filters IPv6 packets encapsulated in IPv4 packets, which prevents IPv6 over IPv4 multicast addresses from being leveraged for network reconnaissance.



- Enable support for multicast traffic so that the firewall can enforce policy on multicast traffic (**Network** > **Virtual Router** > **Multicast**).



https://t.me/learningnets

- Disable the options to **Forward datagrams exceeding UDP content inspection queue** and **Forward segments exceeding TCP content inspection queue** (**Device** > **Setup** > **Content-ID** > **Content-ID Settings**).

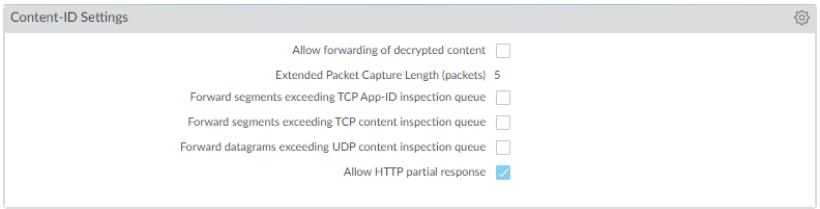
  By default, when the TCP or UDP content inspection queues are full, the firewall skips content inspection for TCP segments or UDP datagrams that exceed the queue limit of 64. Disabling this option ensures content inspection for all TCP and UDP datagrams that the firewall allows. Only under specific circumstances—for example, if the firewall platform is not sized appropriately to align with a use case—could disabling this setting impact performance.

- Disable the **Allow HTTP partial response** (**Device** > **Setup** > **Content-ID** > **Content-ID Settings**).

  The HTTP partial response option allows a client to fetch only part of a file. When a next-generation firewall in the path of a transfer identifies and drops a malicious file, it terminates the TCP session with an RST packet. If the web browser implements the HTTP header range option, it can start a new session to fetch only the remaining part of the file, which prevents the firewall from triggering the same signature again due to the lack of context into the initial session and, at the same time, allows the web browser to reassemble the file and deliver the malicious content. Disabling this option prevents this from happening.

  Allow HTTP partial response is enabled on the firewall by default. This provides maximum availability but increases the risk of a successful cyberattack. For maximum security, disable this option to prevent the web browser from starting a new session to fetch the rest of a file after the firewall terminates the original session due to malicious activity. Disabling HTTP partial response affects HTTP-based data transfers which use the RANGE header, which may cause service anomalies for certain applications. After you disable HTTP partial response, validate the operation of your business-critical applications.

  If you experience HTTP data transfer disruption on a business-critical application, you can create an Application Override policy for that specific application. Because Application Override bypasses App-ID (including threat and content inspection), create an Application Override policy for only the specific business-critical application, and specify sources and destinations to limit the rule (principle of least privilege access). Do not create Application Override policy unless you must. For information about Application Override policies, refer to https://knowledgebase.paloaltonetworks.com/KCSArticleDetail?id=kA10g000000ClVLCA0.



- Create a Vulnerability Protection Profile that blocks protocol anomalies and all vulnerabilities with low and high severities.

  A protocol anomaly occurs when a protocol behavior deviates from standard and compliant usage. For example, a malformed packet, poorly written application, or an application running on a non-standard port would all be considered protocol anomalies, and could be used as evasion tools.

  If yours is a mission-critical network, where the business's highest priority is application availability, you should begin by alerting on protocol anomalies for a period of time to ensure

https://t.me/learningnets

that no critical internal applications are using established protocols in a non-standard way. If you find that certain critical applications trigger protocol anomaly signatures, you can then exclude those applications from protocol anomaly enforcement. To do this, add another rule to the Vulnerability Protection Profile that allows protocol anomalies and attach the profile to the security policy rule that enforces traffic to and from the critical applications.

Make sure that Vulnerability Protection Profile rules and security policy rules that allow protocol anomalies for critical internal applications are listed above rules that block protocol anomalies. Traffic is evaluated against security policy rules and associated Vulnerability Protection Profiles rules from top to bottom, and is enforced based on the first matching rule.

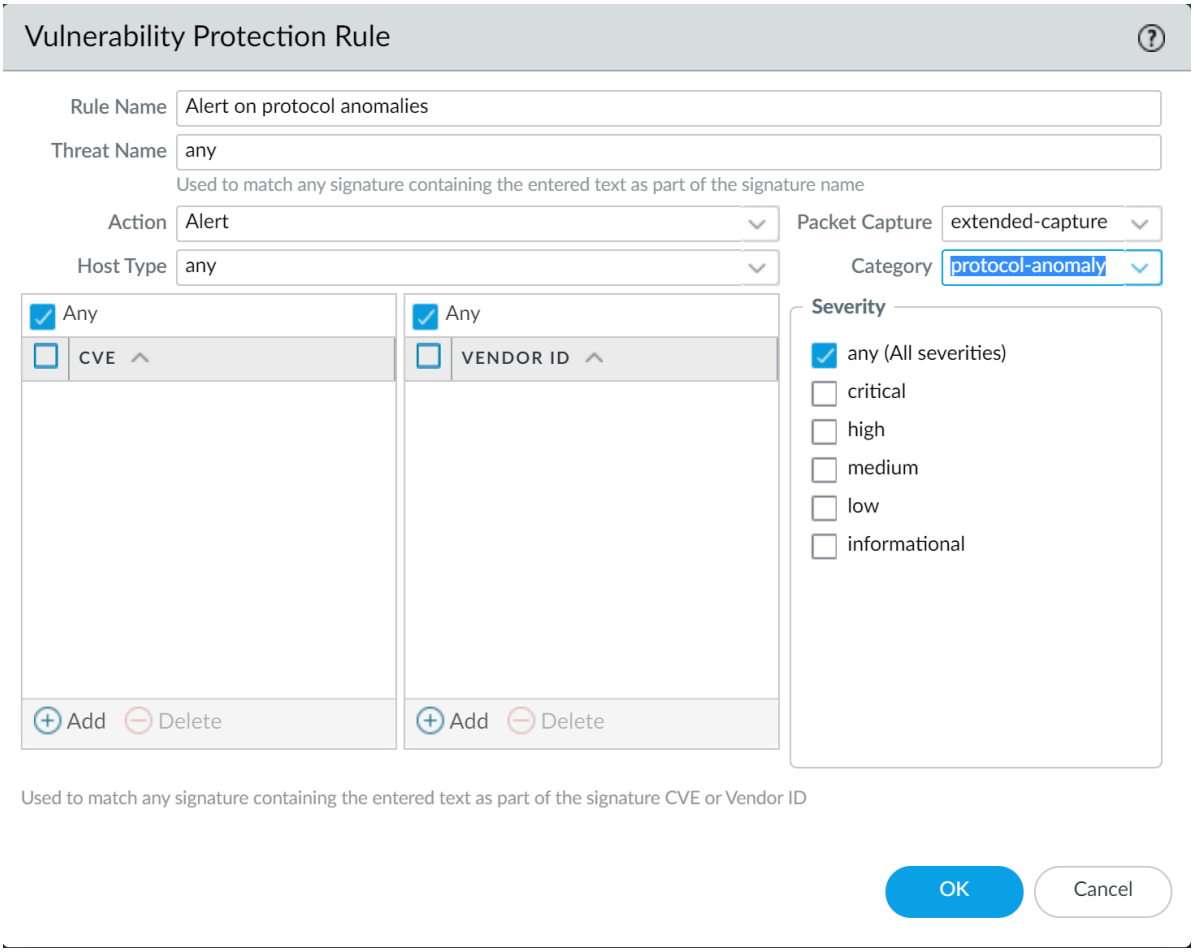
- Begin by alerting on protocol anomalies:

  Create a Vulnerability Protection Profile rule with the **Action** set to Alert, the **Category** set to protocol-anomaly, and the **Severity** set to Any. Monitor your traffic to determine if any critical internal applications are using established protocols in non-standard ways. If you find

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

this to be true, continue to allow protocol anomalies for those applications, and then block protocol anomalies for all other applications.

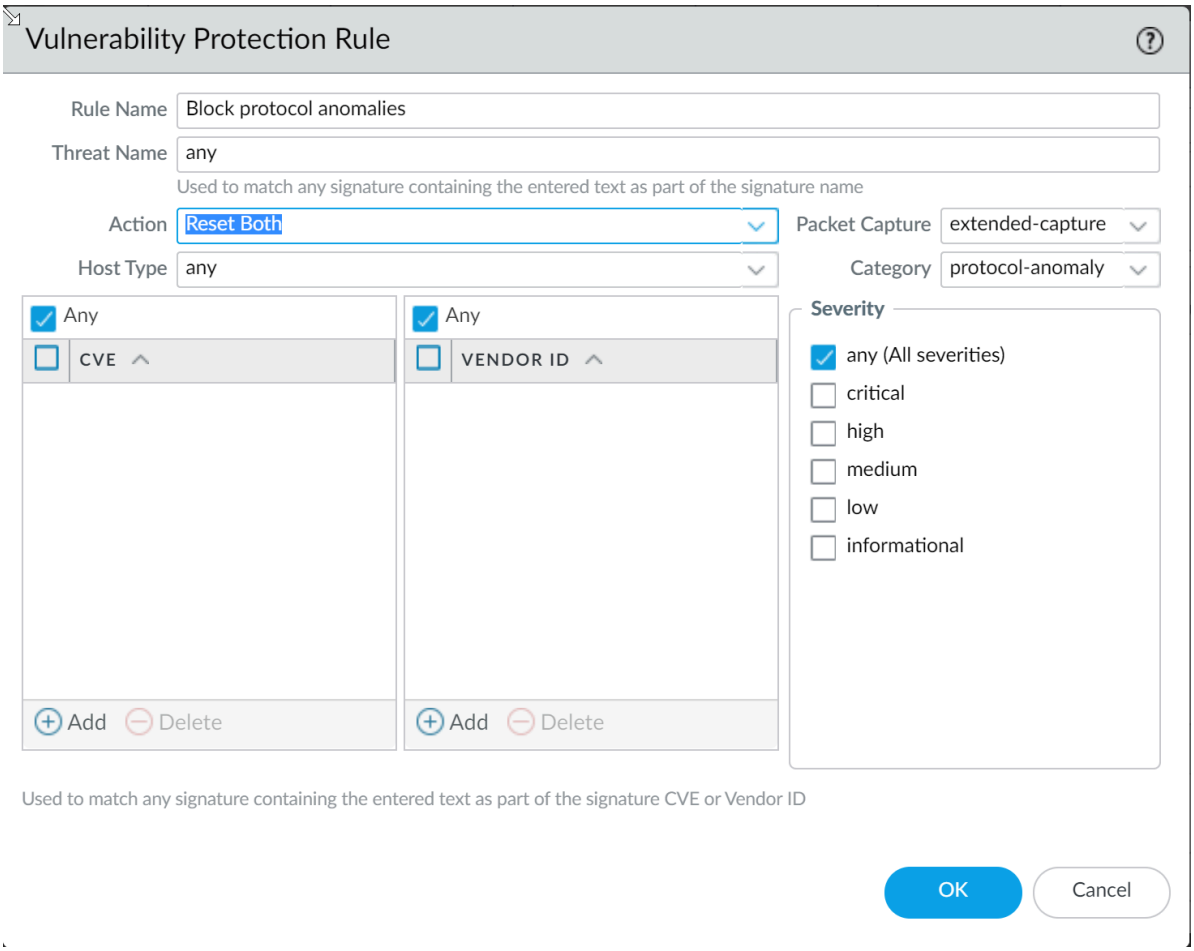


- Block protocol anomalies:

  Create a Vulnerability Protection Profile rule with the **Category** set to protocol-anomaly, the rule **Action** set to Reset Both, and the **Severity** set to Any.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets



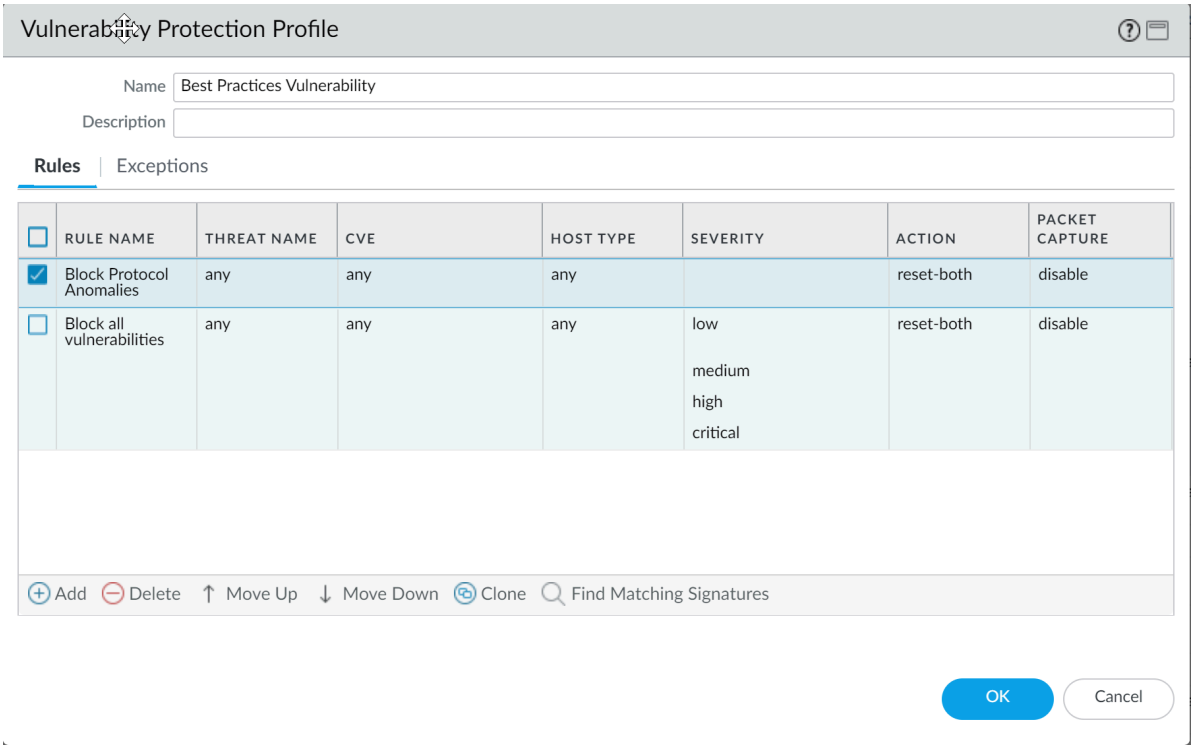
- Optionally allow protocol anomalies for critical applications that use established protocols in a non-standard way. To do this, create a Vulnerability Protection Profile rule that allows protocol anomalies: set the rule **Action** to Allow, the **Category** to protocol-anomaly, and the

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Severity to any. Attach the Vulnerability Protection Profile rule to the security policy rule that enforces traffic to and from critical applications.

- Add another rule to the Vulnerability Protection profile to block all vulnerabilities with low and higher severity. This rule must be listed after the rule that blocks protocol anomalies.



- ☐ Continue to attach the following security profiles to your Security policy rules to provide signature-based protection:
  - An Anti-Spyware profile to block all spyware with severity low and higher.
  - An Antivirus profile to block all content that matches an antivirus signature.

©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Set Up Antivirus, Anti-Spyware, and Vulnerability Protection

Every Palo Alto Networks next-generation firewall comes with predefined Antivirus, Anti-Spyware, and Vulnerability Protection profiles that you can attach to Security policy rules. There is one predefined Antivirus profile, **default**, which uses the default action for each protocol (block HTTP, FTP, and SMB traffic and alert on SMTP, IMAP, and POP3 traffic). There are two predefined Anti-Spyware and Vulnerability Protection profiles:

- **default**—Applies the default action to all client and server critical, high, and medium severity spyware/vulnerability protection events. It does not detect low and informational events.
- **strict**—Applies the block response to all client and server critical, high and medium severity spyware/vulnerability protection events and uses the default action for low and informational events.

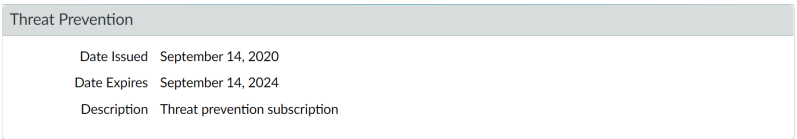
To ensure that the traffic entering your network is free from threats, attach the predefined profiles to your basic web access policies. As you monitor the traffic on your network and expand your policy rulebase, you can then design more granular profiles to address your specific security needs.

Use the following workflow to set up the default Antivirus, Anti-Spyware, and Vulnerability Protection Security Profiles.

*Palo Alto Networks defines a default action for all anti-spyware and vulnerability protection signatures. To see the default action, select **Objects** > **Security Profiles** > **Anti-Spyware** or **Objects** > **Security Profiles** > **Vulnerability Protection** and then select a profile. Click the Exceptions tab and then click Show all signatures to view the list of the signatures and the corresponding default **Action**. To change the default action, create a new profile and specify an **Action**, and/or add individual signature exceptions to **Exceptions** in the profile.*

STEP 1 | Verify that you have a Threat Prevention subscription.

The Threat Prevention subscription bundles the antivirus, anti-spyware, and vulnerability protection features in one license. To verify that you have an active Threat Prevention subscription, select **Device** > **Licenses** and verify that the **Threat Prevention** expiration date is in the future.



STEP 2 | Download the latest content.

1. Select **Device** > **Dynamic Updates** and click **Check Now** at the bottom of the page to retrieve the latest signatures.
2. In the **Actions** column, click **Download** and install the latest Antivirus updates and then download and then **Install** the latest Applications and Threats updates.

https://t.me/learningnets

STEP 3 | Schedule content updates.

Review the Best Practices for Applications and Threats Content Updates for important information on deploying updates.

1. Select **Device** > **Dynamic Updates** and then click **Schedule** to automatically retrieve signature updates for **Antivirus** and **Applications and Threats**.
2. Specify the frequency and timing for the updates:
   - **download-only**—The firewall automatically downloads the latest updates per the schedule you define but you must manually **Install** them.
   - **download-and-install**—The firewall automatically downloads and installs the updates per the schedule you define.
3. Click **OK** to save the update schedule; a commit is not required.
4. (Optional) Define a **Threshold** to indicate the minimum number of hours after an update becomes available before the firewall will download it. For example, setting the **Threshold** to **10** means the firewall will not download an update until it is at least 10 hours old regardless of the schedule.
5. (HA only) Decide whether to **Sync To Peer**, which enables peers to synchronize content updates after download and install (the update schedule does not sync across peers; you must manually configure the schedule on both peers).

   There are additional considerations for deciding if and how to **Sync To Peer** depending on your HA deployment:

   - **Active/Passive HA**—If the firewalls are using the MGT port for content updates, then schedule both firewalls to download and install updates independently. However, if the firewalls are using a data port for content updates, then the passive firewall will not download or install updates unless and until it becomes active. To keep the schedules in sync on both firewalls when using a data port for updates, schedule updates on both firewalls and then enable **Sync To Peer** so that whichever firewall is active downloads and installs the updates and also pushes the updates to the passive firewall.
   - **Active/Active HA**—If the firewalls are using the MGT interface for content updates, then select **download-and-install** on both firewalls but do not enable **Sync To Peer**. However, if the firewalls are using a data port, then select **download-and-install** on both firewalls and enable **Sync To Peer** so that if one firewall goes into the active-secondary state, the active-primary firewall will download and install the updates and push them to the active-secondary firewall.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 4 | (Optional) Create custom security profiles for antivirus, anti-spyware, and vulnerability protection.

Alternatively, you can use the predefined default or strict profiles.

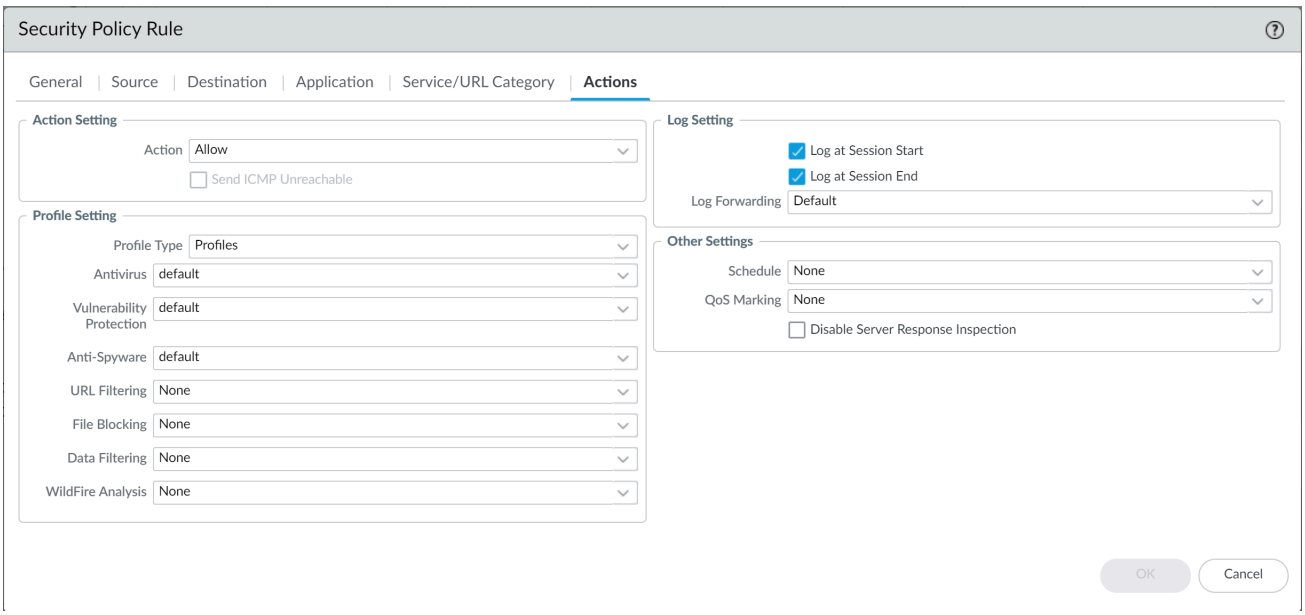
Transition safely *to* best practice Security profiles *for the best security posture.*

- To create custom Antivirus Profiles, select **Objects** > **Security Profiles** > **Antivirus** and **Add** a new profile. Use the Antivirus profile transition steps to safely reach your goal.
- To create custom Anti-Spyware Profiles, select **Objects** > **Security Profiles** > **Anti-Spyware** and **Add** a new profile. Use the Anti-Spyware profile transition steps to safely reach your goal.
- To create custom Vulnerability Protection Profiles, select **Objects** > **Security Profiles** > **Vulnerability Protection** and **Add** a new profile. Use the Vulnerability Protection profile transition steps to safely reach your goal.

STEP 5 | Attach security profiles to your Security policy rules.

When you configure the firewall with a Security policy rule that uses a Vulnerability Protection profile to block connections, the firewall automatically blocks that traffic in hardware (see Monitor Blocked IP Addresses).

1. Select **Policies** > **Security** and select the rule you want to modify.
2. In the **Actions** tab, select **Profiles** as the **Profile Type**.
3. Select the security profiles you created for **Antivirus**, **Anti-Spyware**, and **Vulnerability Protection**.



STEP 6 | Commit your changes.

Click **Commit**.

https://t.me/learningnets

DNS Security

DNS Security is a continuously evolving threat prevention service designed to protect and defend your network from advanced threats using DNS. By leveraging advanced machine learning and predictive analytics, the service provides real-time DNS request analysis and rapidly produces and distributes DNS signatures that are specifically designed to defend against malware using DNS for C2 and data theft. Combined with an extensible cloud architecture, it provides access to a scalable threat intelligence system to keep your network protections up to date.

- About DNS Security
- Cloud-Delivered DNS Signatures and Protections
- DNS Security Analytics
- Enable DNS Security
- DNS Security Data Collection and Logging

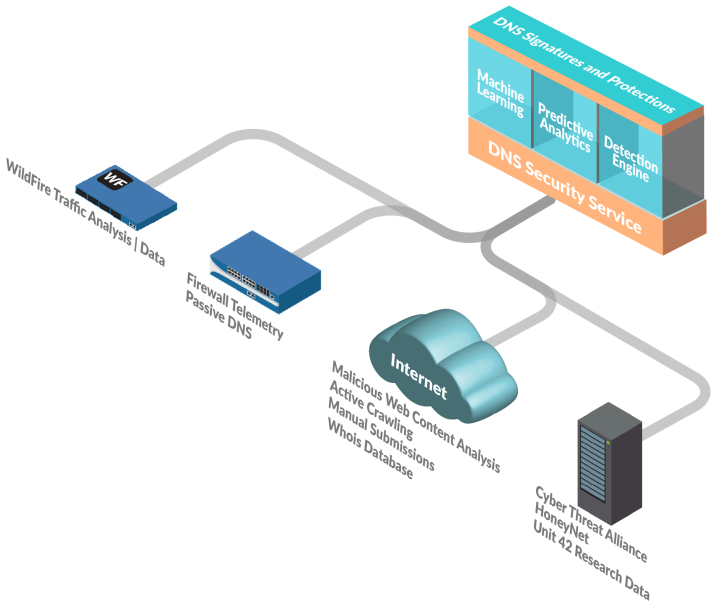
About DNS Security

With an active Threat Prevention license, customers can configure their firewalls to sinkhole DNS requests using a list of domains generated by Palo Alto Networks. These locally-accessed, customizable DNS signature lists are packaged with antivirus and WildFire updates and include the most relevant threats for policy enforcement and protection at the time of publication. For improved coverage against threats using DNS, the DNS Security subscription enables users to access real-time protections using advanced predictive analytics. Using techniques such as DGA/DNS tunneling detection and machine learning, threats hidden within DNS traffic can be proactively identified and shared through an infinitely scalable cloud service. Because the DNS signatures and protections are stored in a cloud-based architecture, you can access the full database of ever-expanding signatures that have been generated using a multitude of data sources. This allows you to defend against an array of threats using DNS in real-time against newly generated malicious domains. To combat future threats, updates to the analysis, detection, and prevention capabilities of the DNS Security service will be available through content releases.

To access the DNS Security service, you must have a valid Threat Prevention and DNS Security license.

The following workflow describes how the DNS Security service uses various data sources to generate DNS signatures:

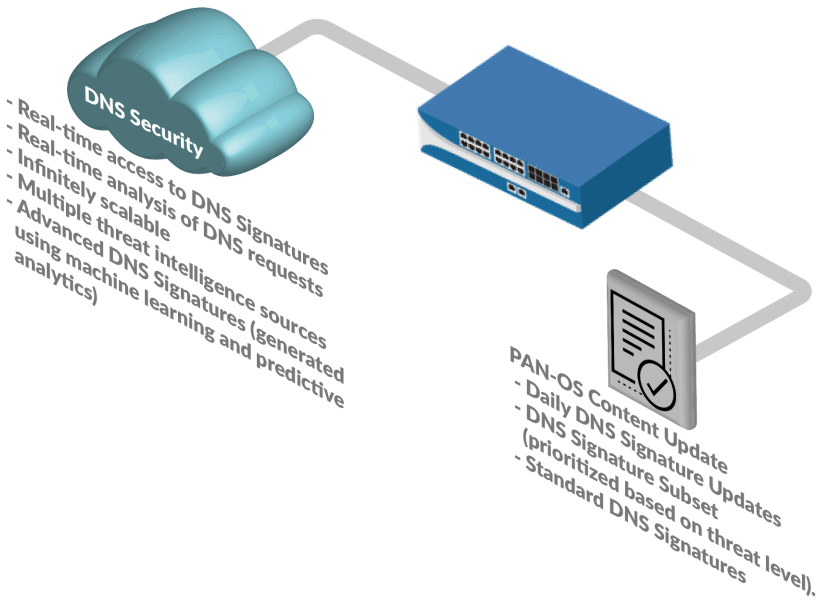
https://t.me/learningnets



Cloud-Delivered DNS Signatures and Protections

As a cloud-based service, DNS Security allows you to access an infinitely scalable DNS signature and protections source to defend your organization from malicious domains. Domain signatures and protections generated by Palo Alto Networks are derived from a multitude of sources, including WildFire traffic analysis, passive DNS, active web crawling & malicious web content analysis, URL sandbox analysis, Honeynet, DGA reverse engineering, telemetry data, whois, the Unit 42 research organization, and third party data sources such as the Cyber Threat Alliance. This on-demand cloud database provides users with access to the complete Palo Alto Network's DNS signature set, including signatures generated using advanced analysis techniques, as well as real-time DNS request analysis. Locally available, downloadable DNS signature sets (packaged with the antivirus and WildFire updates) come with a hard-coded capacity limitation of 100k signatures and do not include signatures generated through advanced analysis. To better accommodate the influx of new DNS signatures being produced on a daily basis, the cloud-based signature database provides users with instant access to newly added DNS signatures without the need to download updates. If network connectivity goes down or is otherwise unavailable, the firewall uses the onbox DNS signature set.

https://t.me/learningnets



DNS Security Analytics

The DNS Security service operates real-time DNS request analysis using predictive analytics and machine learning on multiple DNS data sources. This is used to generate protections for DNS-based threats, which are accessible in real-time through configuration of the Anti-Spyware Security profile attached to a Security policy rule. Each DNS threat category (the DNS Signature Source) allows allow you to define separate policy actions as well as a log severity level for a specific signature type. This enables you to create specific security policies based on the nature of the threat, according to your network security protocols. Palo Alto Networks also generates and maintains a list of explicitly allowable domains based on metrics from PAN-DB and Alexa. These allow list domains are frequently accessed and known to be free of malicious content. The DNS Security categories and the allow list are updated and extensible through PAN-OS content releases.

You can view your organization's DNS statistics data generated by the DNS Security Cloud service using AutoFocus. This provides a fast, visual assessment describing the breakdown of DNS requests passing through your network based on the available DNS categories. Alternatively, you can retrieve domain information, as well as the transaction details, such as latency and TTL using the **`test dns-proxy dns-signature fqdn <domain>`** command.

Upon upgrade to PAN-OS 10.0 and later, the DNS Security source gets redefined into new categories to provide extended granular controls; as a result, the new categories will overwrite the previously defined action and acquire default settings. Make sure to reapply any sinkhole, log severity, and packet captures settings appropriate for the newly defined DNS Security Categories.

The DNS Security service currently supports detection of the following DNS threat categories:

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

The universal threat ID number (indicated as ID in the Threat logs) maps to the specific DNS detection mechanism used by DNS Security to classify domains. This shows the precise categorization of the domain, alongside the broadly defined threat category, that it belongs to.

- **Command and Control Domains**—C2 includes URLs and domains used by malware and/or compromised systems to surreptitiously communicate with an attacker's remote server to receive malicious commands or exfiltrate data (this includes DNS tunneling detection and DGA detection), or deplete resources on a target authoritative DNS servers (such as NXNSattack).
  - **DNS Tunnel Detection** (UTID: 109001001/109001002)—DNS tunneling can be used by attackers to encode data of non-DNS programs and protocols within DNS queries and responses. This provides attackers with an open back channel with which they can transfer files or remotely access the system. DNS tunnel detection uses machine learning to analyze the behavioral qualities of DNS queries, including n-gram frequency analysis of domains, entropy, query rate, and patterns to determine if the query is consistent with a DNS tunneling-based attack. This includes certain next-generation DNS tunneling malware that exfiltrates data slowly across multiple domains to avoid detection, such as TriFive and Snugy. Combined with the firewall's automated policy actions, this allows you to quickly detect C2 or data theft hidden in DNS tunnels and to automatically block it, based on your defined policy rules.
  - **DGA Domain Detection** (UTID: 109000001)—Domain generation algorithms (DGAs) are used to auto-generate domains, typically in large numbers within the context of establishing a malicious command-and-control (C2) communications channel. DGA-based malware (such as Pushdo, BankPatch, and CryptoLocker) limit the number of domains from being blocked by hiding the location of their active C2 servers within a large number of possible suspects, and can be algorithmically generated based on factors such as time of day, cryptographic keys, dictionary-derived naming schemes, and other unique values. While most domains generated by a DGA do not resolve as a valid domain, they must all be identified to fully defend against a given threat. DGA analysis determines whether a domain is likely to have been generated by a machine, rather than a person, by reverse-engineering and analyzing other frequently used techniques found in DGAs. Palo Alto Networks then uses these characteristics to identify and block previously unknown DGA-based threats in real-time.
  - **NXNSAttack** (UTID: 109010007)—The NXNSAttack vulnerability present in the DNS protocol affects all recursive DNS resolvers and can be used by malicious actors to launch DDos-like amplification attacks to disrupt the normal operation of vulnerable authoritative DNS servers. NXNSAttack can introduce massive traffic spikes on an authoritative DNS server by forcing the recursive DNS resolver to issue a large number of invalid requests to potentially shut down the server.
  - **DNS Rebinding** (UTID: 109010009)—DNS rebinding attacks lure users to an attacker-controlled domain configured with a short TTL parameter to manipulate how domain names are resolved to exploit and bypass the same-origin policy in browsers. This enables malicious actors to use the client machine as an intermediary to attack or access a resource contained within a private network.
  - **DNS Infiltration** (UTID: 109010003)—DNS infiltration includes DNS queries that enable malicious actors to hide and resolve minute payloads via a response to fraudulent A (IPv4) and AAAA (IPv6) record requests. When the client resolves multiple subdomains, each containing a A/AAAA record with an encoded component, the data contained within them can be consolidated to form a malicious payload, which can then be executed on the client

https://t.me/learningnets

machine. After executing the payload, it can introduce secondary payloads to establish a DNS tunnel or additional exploits.

- **Dynamic DNS Hosted Domains** (UTID: 109020002) —Dynamic DNS (DDNS) services provide mapping between hostnames and IP addresses in near real-time to keep changing IP addresses linked to a specific domain, when static IPs are unavailable. This provides attackers a method of infiltrating networks by using DDNS services to change the IP addresses that host command-and-control servers. Malware campaigns and exploit kits can utilize DDNS services as part of their payload distribution strategy. By utilizing DDNS domains as part of their hostname infrastructure, adversaries can change the IP address associated with given DNS records and more easily avoid detection. DNS Security detects exploitative DDNS services by filtering and cross-referencing DNS data from various sources to generate candidate lists which are then further validated to maximize accuracy.
- **Malware Domains** (UTID: 109003001)—Malicious domains host and distribute malware and can include websites that attempt to install various threats (such as executables, scripts, viruses, drive-by downloads). Malicious domains are distinguishable from C2 domains in that they deliver malicious payloads into your network via an external source, whereas with C2, infected endpoints typically attempt to connect to a remote server to retrieve additional instructions or other malicious content.
- **Newly Registered Domains** (UTID: 109020001)—Newly registered domains are domains that have been recently added by a TLD operator or had change in ownership within the last 32 days. While new domains can be created for legitimate purposes, the vast majority are often used to facilitate malicious activities, such as operating as C2 servers or used to distribute malware, spam, PUP/adware. Palo Alto Networks detects newly registered domains by monitoring specific feeds (domain registries and registrars) and using zone files, passive DNS, WHOIS data to detect registration campaigns.
- **Phishing Domains** (UTID: 109010001) —Phishing domains attempt to lure users into submitting sensitive data, such as personal information or user credentials, by masquerading as legitimate websites through phishing or pharming. These malicious activities can be conducted through social engineering campaigns (whereby a seemingly trusted source manipulates users into submitting personal information via email or other forms of electronic communications) or through web traffic redirection, which directs users to fraudulent sites that appear legitimate.
- **Grayware Domains** (UTID: 109010002) —(Available with installation of PAN-OS content release 8290 and later). Grayware domains generally do not pose a direct security threat,

https://t.me/learningnets

however, they can facilitate vectors of attack, produce various undesirable behaviors, or might simply contain questionable/offensive content.These can include websites and domains that:

- Attempt to trick users into granting remote access.
- Contain adware and other unsolicited applications (such as cryptominers, hijackers, and PUPs [potentially unwanted programs]).
- Deploy domain identification concealment actions using fast flux techniques (fastflux detection - 109010005).
- Demonstrate malicious behavior and usage as evidenced through DNS security predicative analytics (malicious NRD - UTID: 109010006).
- Take advantage of user errors when entering web page addresses (typosquatting domains).
- Redirect traffic from a legitimate source to a malicious website due to an improperly configured or stale DNS record on an authoritative DNS server that has not been removed or otherwise corrected (dangling DNS - UTID: 109010008).
- Promote illegal activities or scams.
- Include wildcard DNS entries, which can be used to evade block lists or enable wildcard DNS attacks by routing traffic to malicious websites (Wildcard abuse - UTID: 109002001).
- Indicate the presence of DNS traffic with anomalous characteristics when compared to established baseline profiles built from collected DNS data (Anomaly detection).
- Have been registered months or years in advance and left in a state of dormancy to bypass reputation checks when they become active (Strategically-aged domains - UTID: 109002002).

- **Parked Domains** (UTID: 109010003)—(Available with installation of PAN-OS content release 8318 and later) Parked domains are typically inactive websites that host limited content, often in the form of click-through ads which may generate revenue for the host entity, but generally do not contain content that is useful to the end user. While they often function as a legitimate placeholder or as nothing more than a benign nuisance, they could also be used as a possible vector for distribution of malware.
- **Proxy Avoidance and Anonymizers** (UTID: 109010004)—(Available with installation of PAN-OS content release 8340 and later) Proxy Avoidance and Anonymizers is traffic to services that are used to bypass content filtering policies. Users who attempt to circumvent an organization's content filtering policies via anonymizer proxy services are blocked at the DNS level.

Enable DNS Security

To enable DNS sinkholing for domain queries using DNS security, you must activate your DNS Security subscription, create (or modify) an Anti-Spyware policy to reference the DNS Security service, configure the log severity and policy settings for each DNS signature category, and then attach the profile to a security policy rule.

STEP 1 | Activate Subscription Licenses.

https://t.me/learningnets

STEP 2 | Verify that the *paloalto-dns-security* App-ID in your security policy is configured to enable traffic from the DNS security cloud security service.

If your firewall deployment routes your management traffic though an Internet-facing perimeter firewall configured to enforce App-ID security policies, you must allow the App-IDs on the perimeter firewall; failure to do so will prevent DNS security connectivity.

STEP 3 | Configure DNS Security signature policy settings to send malware DNS queries to the defined sinkhole.

If you use an external dynamic list as a domain allow list, it does not have precedence over the DNS Security domain policy actions. As a result, when there is a domain match to an entry in the EDL and a DNS Security domain category, the action specified under DNS Security is still applied, even when the EDL is explicitly configured with an action of Allow. If you want to add DNS domain exceptions, either configure an EDL with an Alert action or add them to the DNS Domain/FQDN Allow List located in the DNS Exceptions tab (step 8).

1. Select **Objects** > **Security Profiles** > **Anti-Spyware**.
2. Create or modify an existing profile, or select one of the existing default profiles and clone it.
3. **Name** the profile and, optionally, provide a description.
4. Select the **DNS Policies** tab.
5. In the **Signature Source** column, beneath the DNS Security heading, there are individually configurable DNS signature sources, which allow you to define separate policy actions as well as a log severity level.

   Palo Alto Networks recommends changing your default DNS Policies settings for signature sources to ensure optimum coverage as well as to assist with incidence response and remediation. Follow the best practices for configuring your DNS Security settings as outlined in Best Practices for Securing Your Network from Layer 4 and Layer 7 Evasions.

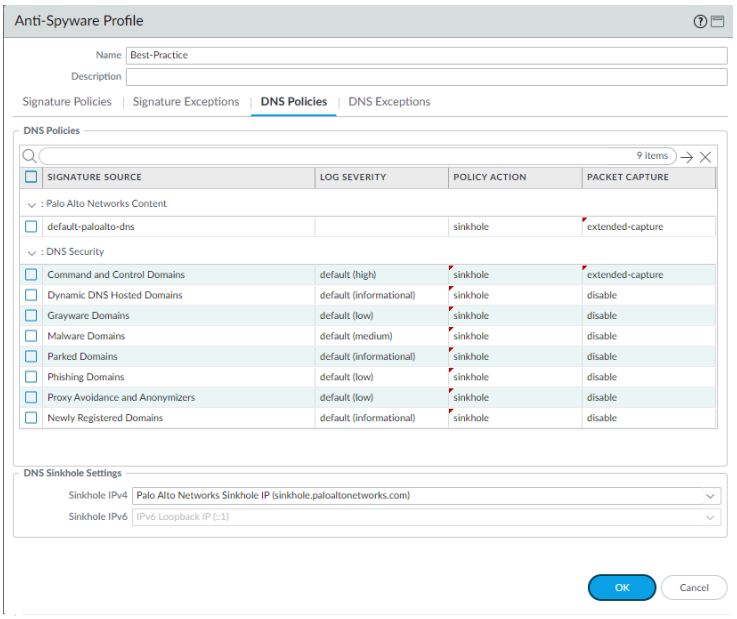
   - Specify the log severity level that is recorded when the firewall detects a domain matching a DNS signature. For more information about the various log severity levels, refer to Threat Severity Levels.
   - Select an action to be taken when DNS lookups are made to known malware sites for the DNS Security signature source. The options are alert, allow, block, or sinkhole. Verify that the action is set to sinkhole.
   - In the **Packet Capture** drop-down, select **single-packet** to capture the first packet of the session or **extended-capture** to set between 1-50 packets. You can then use the packet captures for further analysis.
6. In the **DNS Sinkhole Settings** section, verify that **Sinkhole** is enabled. For your convenience, the default Sinkhole address (sinkhole.paloaltonetworks.com) is set to

https://t.me/learningnets

access a Palo Alto Networks server. Palo Alto Networks can automatically refresh this address through content updates.

If you want to modify the **Sinkhole IPv4** or **Sinkhole IPv6** address to a local server on your network or to a loopback address, see Configure the Sinkhole IP Address to a Local Server on Your Network.

7. Click **OK** to save the Anti-Spyware profile.



STEP 4 | Attach the Anti-Spyware profile to a Security policy rule.

1. Select **Policies** > **Security**.
2. Select or create a **Security Policy Rule**.
3. On the **Actions** tab, select the **Log at Session End** check box to enable logging.
4. In the Profile Setting section, click the **Profile Type** drop-down to view all **Profiles**. From the **Anti-Spyware** drop-down and select the new or modified profile.
5. Click **OK** to save the policy rule.

https://t.me/learningnets

STEP 5 | Test that the policy action is enforced.

1. Access the following test domains to verify that the policy action for a given threat type is being enforced:
   - C2—test-c2.testpanw.com
   - DNS Tunneling—test-dnstun.testpanw.com
   - DGA—test-dga.testpanw.com
   - Dynamic DNS—test-ddns.testpanw.com
   - Malware—test-malware.testpanw.com
   - Newly Registered Domains—test-nrd.testpanw.com
   - Phishing—test-phishing.testpanw.com
   - Grayware—test-grayware.testpanw.com
   - Parked—test-parked.testpanw.com
   - Proxy Avoidance and Anonymizers—test-proxy.testpanw.com
2. To monitor the activity on the firewall:
   1. Select **ACC** and add a URL Domain as a global filter to view the Threat Activity and Blocked Activity for the domain you accessed.
   2. Select **Monitor** > **Logs** > **Threat** and filter by `(action eq sinkhole)` to view logs on sinkholed domains.

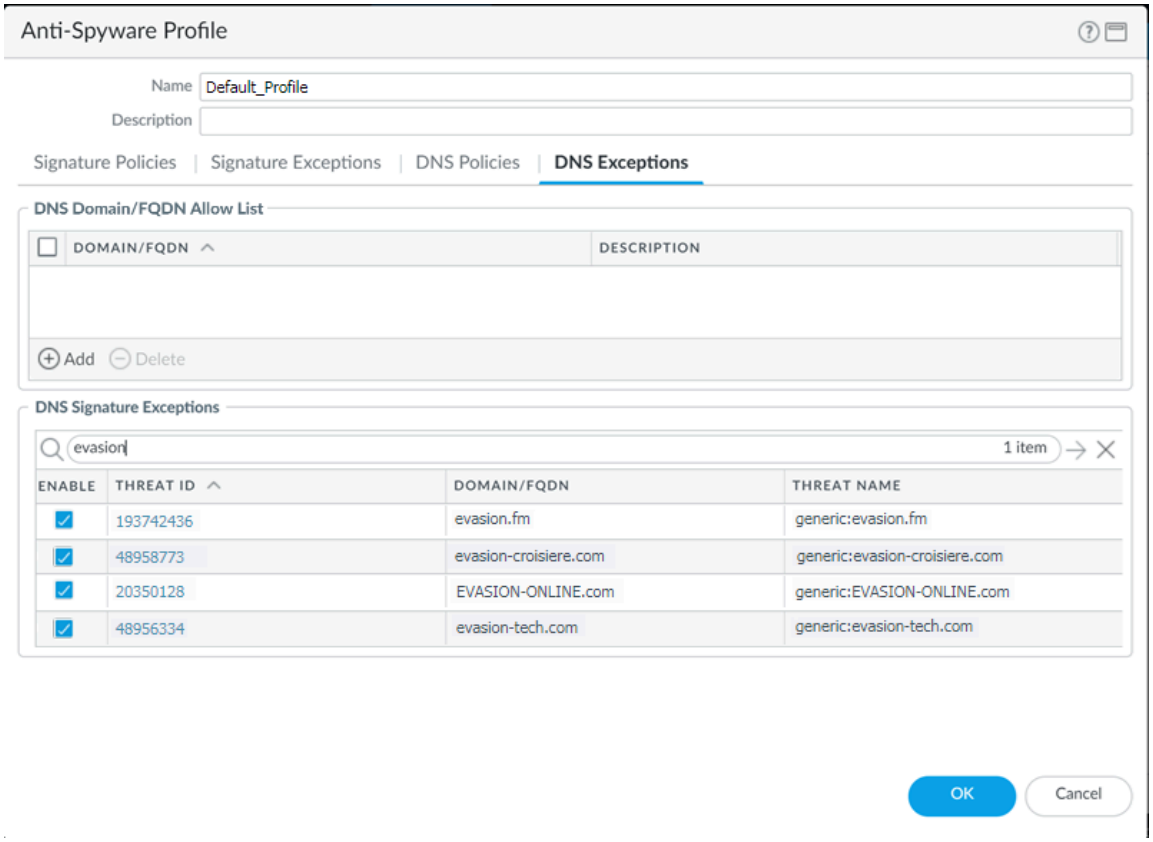
STEP 6 | Identify Infected Traffic Hosts in the Traffic Logs

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 7 | (Optional) Add domain signature exceptions in cases where false-positives occur.

1. Select **Objects** > **Security Profiles** > **Anti-Spyware**.
2. Select a profile to modify.
3. **Add** or modify the Anti-Spyware profile from which you want to exclude the threat signature, and select **DNS Exceptions**.
4. Search for a DNS signature to exclude by entering the name or FQDN.
5. Select the checkbox for each **Threat ID** of the DNS signature that you want to exclude from enforcement.
6. Click **OK** to save your new or modified Anti-Spyware profile.



STEP 8 | (Optional) Add an allow list to specify a list of DNS domains / FQDNs to be explicitly allowed.

1. Select **Objects** > **Security Profiles** > **Anti-Spyware**.
2. Select a profile to modify.
3. **Add** or modify the Anti-Spyware profile from which you want to exclude the threat signature, and select **DNS Exceptions**.
4. To **Add** a new **FQDN Allow List**, provide the DNS domain or FQDN location and a description.
5. Click **OK** to save your new or modified Anti-Spyware profile.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 9 | (Optional) Verify your firewall's connectivity to the DNS Security service. If you cannot reach the service, verify that the following domain is not being blocked: dns.service.paloaltonetworks.com.

Use the following CLI command on the firewall to verify your firewall's connection availability to the DNS Security service.

```
show dns-proxy dns-signature info
```

For example:

```
show dns-proxy dns-signture info

Cloud URL: dns.service.paloaltonetworks.com:443

Telemetry URL: io.dns.service.paloaltonetworks.com:443

Last Result: None

Last Server Address:

Parameter Exchange: Interval 300 sec

Allow List Refresh: Interval 43200 sec

Request Waiting Transmission: 0

Request Pending Response: 0

Cache Size: 0
```

STEP 10 | (Optional) Retrieve a specified domain's transaction details, such as latency, TTL, and the signature category.

Use the following CLI command on the firewall to review the details about the list.

```
test dns-proxy dns-signature fqdn
```

For example:

```
test dns-proxy dns-signature fqdn www.yahoo.com

DNS Signature Query [ www.yahoo.com ]

Completed in 178 ms

DNS Signature Response

Entries: 2

Domain                              Category        GTID
     TTL
```

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

```
--------------------------------------------------------------------------------
*.yahoo.com                          Benign          0
    86400
  www.yahoo.com                      Benign          0
    3600
```

STEP 11 | (Optional) Configure the DNS signature lookup timeout setting. If the firewall is unable to retrieve a signature verdict in the allotted time due to connectivity issues, the request, including all subsequent DNS responses, are passed through. You can check the average latency to verify that the requests fall within the configured period. If the average latency exceeds the configured period, consider updating the setting to a value that is higher than the average latency to prevent requests from timing out.

1. In the CLI, issue the following command to view the average latency.

   ```
   show dns-proxy dns-signature
   counters
   ```

   The default timeout is 100 milliseconds.

2. Scroll down through the output to the latency section under the Signature query API heading and verify that the average latency falls within the defined timeout period. This latency indicates the amount of time it takes, on average, to retrieve a signature verdict from the DNS security service. Additional latency statistics for various latency periods can be found below the averages.

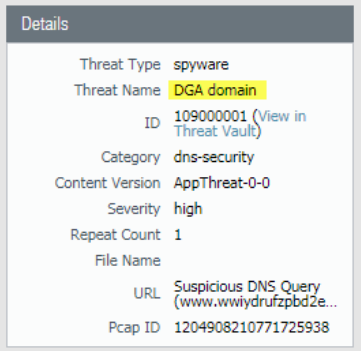
   ```
   Signature query API:
     .
     .
     .
        [latency   ] :
           max    1870 (ms)  min    16(ms)  avg    27(ms)
             50 or less : 47246
             100 or less : 113
             200 or less : 25
             400 or less : 15
                    else : 21
   ```

3. If the average latency is consistency above the default timeout value, you can raise the setting so that the requests fall within a given period. Select **Device > Content-ID** and update the **Realtime Signature Lookup** setting.
4. Commit the changes.

To view sinkholed DNS queries, refer to the firewall threat logs (**Monitor > Logs**, then select the log type from the list):

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets



DNS Security Data Collection and Logging

The DNS Security service collects server response and request information based on your firewall security policy rules, associated action, and the DNS query details when performing domain lookups. The firewall forwards supplemental DNS data to the DNS Security cloud servers and is used by Palo Alto Networks services to provide more accurate domain information (such as provider ASN, hosting information, and geolocation identification). While this supplemental data is not necessary to operate the DNS Security service, it provides the resources to generate improved analytics, DNS detection, and prevention capabilities. This action occurs in less than 30 seconds after collection and batching does not impact firewall performance. In cases where the firewall is experiencing a high load, DNS data collection scales down as needed to maintain expected performance levels.

The firewall can submit the following data fields:

Field	Description
Action	Displays the policy action taken on the DNS query.
Type	Displays the DNS record type.
Response	The IP address that the domain in the DNS query got resolved to.
Response Code	The DNS response code that was received as an answer to your DNS query.
Source IP	The IP address of the system that made the DNS request.
Source User	When the firewall User-ID feature is enabled, the identity of the DNS requester is shown.
Source Zone	The configured source zone referenced in your security policy rule.

DNS expanded data collection is bypassed for domains added to the Allow list in DNS Exceptions.

Data fields that can be used to potentially identify users (Source IP, Source User, and Source Zone) can be withheld from automatic submission using the following CLI command: **set**

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

deviceconfig setting ctd cloud-dns-privacy-mask yes. You must **commit** the changes for the update to take effect.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Use DNS Queries to Identify Infected Hosts on the Network

The DNS sinkhole action in Anti-Spyware profiles enables the firewall to forge a response to a DNS query for a known malicious domain or to a custom domain, so that you can identify hosts on your network that have been infected with malware. A compromised host might initiate communication with a command-and-control (C2) server—once the connection is made, an attacker can remotely control the infected host, in order to further infiltrate the network or exfiltrate data.

DNS queries to any domain included in the Palo Alto Networks DNS signatures list is sinkholed to a Palo Alto Networks server IP address.

The firewall has two sources of DNS signatures that it can use to identify malicious and C2 domains:

- (Requires Threat Prevention) Local DNS signatures—This is a limited, on-box set of DNS signatures that the firewall can use to identify malicious domains. The firewall gets new DNS signatures as part of daily antivirus updates.
- (Requires DNS Security) DNS Security signatures—The firewall accesses the Palo Alto Networks DNS Security cloud service to check for malicious domains against the complete database of DNS signatures. Certain signatures—that only DNS Security provides—can uniquely detect C2 attacks that use machine learning techniques, like domain generation algorithms (DGAs) and DNS tunneling.

DNS queries to domains in the local DNS signature set or the DNS Security signature set are redirected to a Palo Alto Networks server, and the host is unable to access the malicious domain. The following topics provide details on how to enable DNS sinkholing so that you can identify infected hosts.

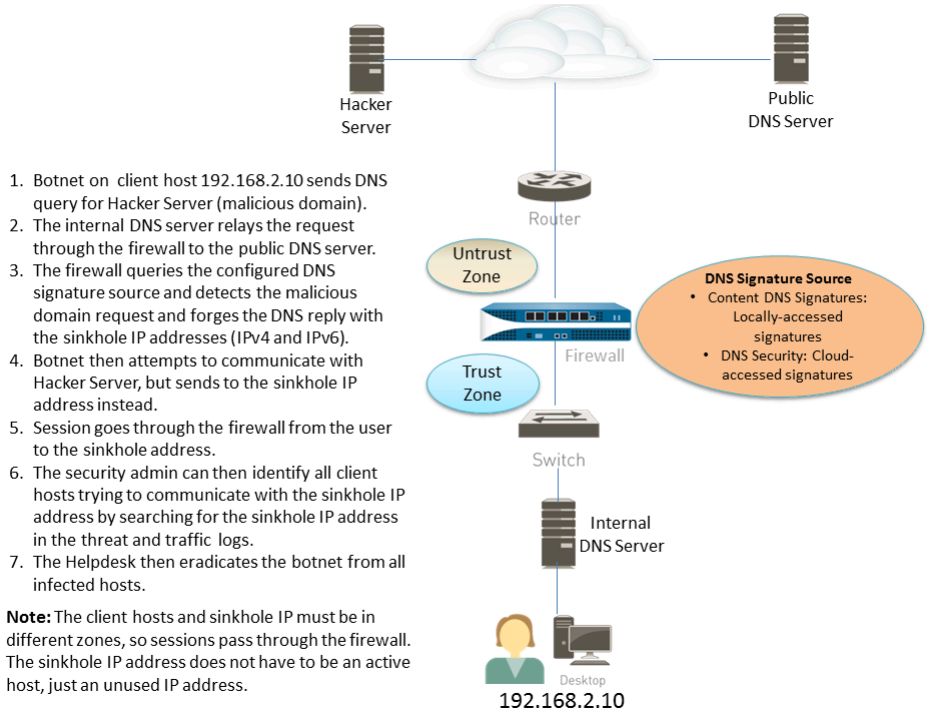
- Learn How DNS Sinkholing Works.
- Configure DNS Sinkholing.
- Configure DNS Sinkholing for a List of Custom Domains.
- Enable DNS Security to sinkhole C2 domains.
- Configurethe Sinkhole IPAddress to a Local Server on Your Network.
- See Infected Hosts that Attempted to Connect to a Malicious Domain.

How DNS Sinkholing Works

DNS sinkholing helps you to identify infected hosts on the protected network using DNS traffic in situations where the firewall cannot see the infected client's DNS query (that is, the firewall cannot see the originator of the DNS query). In a typical deployment where the firewall is north of the local DNS server, the threat log will identify the local DNS resolver as the source of the traffic rather than the actual infected host. Sinkholing malware DNS queries solves this visibility problem by forging responses to the client host queries directed at malicious domains, so that clients attempting to connect to malicious domains (for command-and-control, for example) will instead attempt to connect to a default Palo Alto Networks sinkhole IP address (or to IP address

https://t.me/learningnets

that you define if you choose to Configure DNS Sinkholing for a List of Custom Domains). Infected hosts can then be easily identified in the traffic logs.



Configure DNS Sinkholing

To enable DNS sinkholing, attach the default Anti-Spyware profile to a security policy rule (see Set Up Antivirus, Anti-Spyware, and Vulnerability Protection). DNS queries to any domain included in the Palo Alto Networks DNS signature source that you specify are resolved to the default Palo Alto Networks sinkhole IP address. The IP addresses currently are IPv4—sinkhole.paloaltonetworks.com and a loopback address IPv6 address—::1. These address are subject to change and can be updated with content updates.

STEP 1 | Enable DNS sinkholing for the custom list of domains in an external dynamic list.

1. Select **Objects** > **Security Profiles** > **Anti-Spyware**.
2. Modify an existing profile, or select one of the existing default profiles and clone it.
3. **Name** the profile and select the **DNS Policies** tab.
4. Verify that **default-paloalto-dns** is present in the **Signature Source**.
5. (Optional) In the **Packet Capture** drop-down, select **single-packet** to capture the first packet of the session or **extended-capture** to set between 1-50 packets. You can then use the packet captures for further analysis.

https://t.me/learningnets

STEP 2 | Verify the sinkholing settings on the Anti-Spyware profile.

1. On the **DNS Policies** tab, verify that the **Policy Action** on DNS queries is **sinkhole**.
2. In the DNS Sinkhole Settings section, verify that **Sinkhole** is enabled. For your convenience, the default Sinkhole IP address is set to access a Palo Alto Networks server. Palo Alto Networks can automatically refresh this IP address through content updates.

   If you want to modify the **Sinkhole IPv4** or **Sinkhole IPv6** address to a local server on your network or to a loopback address, see Configure the Sinkhole IP Address to a Local Server on Your Network.
3. Click **OK** to save the Anti-Spyware profile.

STEP 3 | Attach the Anti-Spyware profile to a Security policy rule.

1. Select **Policies** > **Security** and select a security policy rule.
2. On the **Actions** tab, select the **Log at Session Start** check box to enable logging.
3. In the Profile Setting section, click the **Profile Type** drop-down to view all **Profiles**. From the **Anti-Spyware** drop-down and select the new profile.
4. Click **OK** to save the policy rule.

STEP 4 | Test that the policy action is enforced by monitoring the activity on the firewall.

1. Select **ACC** and add a URL Domain as a global filter to view the Threat Activity and Blocked Activity for the domain you accessed.
2. Select **Monitor** > **Logs** > **Threat** and filter by `(action eq sinkhole)` to view logs on sinkholed domains.

Configure DNS Sinkholing for a List of Custom Domains

To enable DNS Sinkholing for a custom list of domains, you must create an External Dynamic List that includes the domains, enable the sinkhole action in an Anti-Spyware profile and attach the profile to a security policy rule. When a client attempts to access a malicious domain in the list, the firewall forges the destination IP address in the packet to the default Palo Alto Networks server or to a user-defined IP address for sinkholing.

For each custom domain included in the external dynamic list, the firewall generates DNS-based spyware signatures. The signature is named Custom Malicious DNS Query <domain name>, and is of type spyware with medium severity; each signature is a 24-byte hash of the domain name.

Each firewall model supports a maximum of 50,000 domain names total in one or more external dynamic lists but no maximum limit is enforced for any one list.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 1 | Enable DNS sinkholing for the custom list of domains in an external dynamic list.

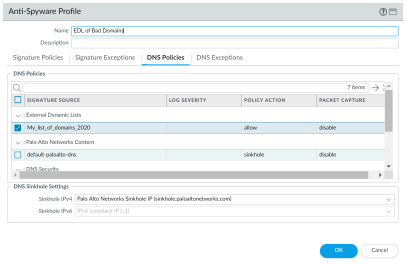
1. Select **Objects** > **Security Profiles** > **Anti-Spyware**.
2. Modify an existing profile, or select one of the existing default profiles and clone it.
3. **Name** the profile and select the **DNS Policies** tab.
4. Select an EDL from the **External Dynamic Lists** signature source.

   *If you have already created an external dynamic list of type: **Domain List**, you can select it from here. The list does not display external dynamic lists of type URL or IP Address that you may have created.*
5. Configure the external dynamic list from the Anti-Spyware profile (see Configure the Firewall to Access an External Dynamic List). The **Type** is preset to **Domain List**.
6. (Optional) In the **Packet Capture** drop-down, select **single-packet** to capture the first packet of the session or **extended-capture** to set between 1-50 packets. You can then use the packet captures for further analysis.

STEP 2 | Verify the sinkholing settings on the Anti-Spyware profile.

1. On the **DNS Policies** tab, verify that the **Policy Action** on DNS queries is **sinkhole**.
2. In the DNS Sinkhole Settings section, verify that **Sinkhole** is enabled. For your convenience, the default Sinkhole IP address is set to access a Palo Alto Networks server. Palo Alto Networks can automatically refresh this IP address through content updates.

   If you want to modify the **Sinkhole IPv4** or **Sinkhole IPv6** address to a local server on your network or to a loopback address, see Configure the Sinkhole IP Address to a Local Server on Your Network.



3. Click **OK** to save the Anti-Spyware profile.

STEP 3 | Attach the Anti-Spyware profile to a Security policy rule.

1. Select **Policies** > **Security** and select a security policy rule.
2. On the **Actions** tab, select the **Log at Session Start** check box to enable logging.
3. In the Profile Setting section, click the **Profile Type** drop-down to view all **Profiles**. From the **Anti-Spyware** drop-down and select the new profile.
4. Click **OK** to save the policy rule.

https://t.me/learningnets

STEP 4 | Test that the policy action is enforced.

1. View External Dynamic List Entries that belong to the domain list, and access a domain from the list.
2. To monitor the activity on the firewall:
   1. Select **ACC** and add a URL Domain as a global filter to view the Threat Activity and Blocked Activity for the domain you accessed.
   2. Select **Monitor** > **Logs** > **Threat** and filter by `(action eq sinkhole)` to view logs on sinkholed domains.

STEP 5 | Verify whether entries in the external dynamic list are ignored or skipped.

Use the following CLI command on the firewall to review the details about the list.

```
request system external-list show type domain name <list_name>
```

For example:

```
request system external-list show type domain name
My_List_of_Domains_2015
vsys1/EBLDomain:
Next update at : Thu May 21 10:15:39 2015
Source : https://1.2.3.4/My_List_of_Domains_2015
Referenced : Yes
Valid : Yes
Number of entries : 3
domains:www.example.com
baddomain.com
qqq.abcedfg.com
```

STEP 6 | (Optional) Retrieve the external dynamic list on-demand.

To force the firewall to retrieve the updated list on-demand instead of at the next refresh interval (the **Repeat** frequency you defined for the external dynamic list), use the following CLI command:

```
request system external-list refresh type domain name  <list_name>
```

As an alternative, you can use the firewall interface to Retrieve an External Dynamic List from the Web Server.

Configure the Sinkhole IP Address to a Local Server on Your Network

By default, sinkholing is enabled for all Palo Alto Networks DNS signatures, and the sinkhole IP address is set to access a Palo Alto Networks server. Use the instructions in this section if you want to set the sinkhole IP address to a local server on your network.

You must obtain both an IPv4 and IPv6 address to use as the sinkhole IP addresses because malicious software may perform DNS queries using one or both of these protocols. The DNS

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

sinkhole address must be in a different zone than the client hosts to ensure that when an infected host attempts to start a session with the sinkhole IP address, it will be routed through the firewall.

The sinkhole addresses must be reserved for this purpose and do not need to be assigned to a physical host. You can optionally use a honey-pot server as a physical host to further analyze the malicious traffic.

The configuration steps that follow use the following example DNS sinkhole addresses:

IPv4 DNS sinkhole address—10.15.0.20

IPv6 DNS sinkhole address—fd97:3dec:4d27:e37c:5:5:5:5

STEP 1 | Configure the sinkhole interface and zone.

Traffic from the zone where the client hosts reside must route to the zone where the sinkhole IP address is defined, so traffic will be logged.

Use a dedicated zone for sinkhole traffic, because the infected host will be sending traffic to this zone.

1. Select **Network** > **Interfaces** and select an interface to configure as your sinkhole interface.
2. In the **Interface Type** drop-down, select **Layer3**.
3. To add an IPv4 address, select the **IPv4** tab and select **Static** and then click **Add**. In this example, add 10.15.0.20 as the IPv4 DNS sinkhole address.
4. Select the **IPv6** tab and click **Static** and then click **Add** and enter an IPv6 address and subnet mask. In this example, enter fd97:3dec:4d27:e37c::/64 as the IPv6 sinkhole address.
5. Click **OK** to save.
6. To add a zone for the sinkhole, select **Network** > **Zones** and click **Add**.
7. Enter zone **Name**.
8. In the **Type** drop-down select **Layer3**.
9. In the **Interfaces** section, click **Add** and add the interface you just configured.
10. Click **OK**.

STEP 2 | Enable DNS sinkholing.

By default, sinkholing is enabled for all Palo Alto Networks DNS signatures. To change the sinkhole address to your local server, see Step Verify the sinkholing settings on the Anti-Spyware profile. in Configure DNS Sinkholing for a List of Custom Domains.

STEP 3 | Edit the security policy rule that allows traffic from client hosts in the trust zone to the untrust zone to include the sinkhole zone as a destination and attach the Anti-Spyware profile.

Editing the Security policy rule(s) that allows traffic from client hosts in the trust zone to the untrust zone ensures that you are identifying traffic from infected hosts. By adding the

https://t.me/learningnets

sinkhole zone as a destination on the rule, you enable infected clients to send bogus DNS queries to the DNS sinkhole.

1. Select **Policies** > **Security**.
2. Select an existing rule that allows traffic from the client host zone to the untrust zone.
3. On the **Destination** tab, **Add** the Sinkhole zone. This allows client host traffic to flow to the sinkhole zone.
4. On the **Actions** tab, select the **Log at Session Start** check box to enable logging. This will ensure that traffic from client hosts in the Trust zone will be logged when accessing the Untrust or Sinkhole zones.
5. In the **Profile Setting** section, select the **Anti-Spyware** profile in which you enabled DNS sinkholing.
6. Click **OK** to save the Security policy rule and then **Commit**.

STEP 4 | To confirm that you will be able to identify infected hosts, verify that traffic going from the client host in the Trust zone to the new Sinkhole zone is being logged.

In this example, the infected client host is 192.168.2.10 and the Sinkhole IPv4 address is 10.15.0.20.

1. From a client host in the trust zone, open a command prompt and run the following command:

```
C:\>ping <sinkhole address>
```

The following example output shows the ping request to the DNS sinkhole address at 10.15.0.2 and the result, which is `Request timed out` because in this example the sinkhole IP address is not assigned to a physical host:

```
C:\>ping 10.15.0.20
Pinging 10.15.0.20 with 32 bytes of data:
Request timed out.
Request timed out.
Ping statistics for 10.15.0.20:
Packets: Sent = 4, Received = 0, Lost = 4 (100% loss)
```

2. On the firewall, select **Monitor** > **Logs** > **Traffic** and find the log entry with the Source 192.168.2.10 and Destination 10.15.0.20. This will confirm that the traffic to the sinkhole IP address is traversing the firewall zones.



*You can search and/or filter the logs and only show logs with the destination 10.15.0.20. To do this, click the IP address (10.15.0.20) in the **Destination** column, which will add the filter (addr.dst in 10.15.0.20) to the search field. Click the Apply Filter icon to the right of the search field to apply the filter.*

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 5 | Test that DNS sinkholing is configured properly.

You are simulating the action that an infected client host would perform when a malicious application attempts to call home.

1. Find a malicious domain that is included in the firewall's current Antivirus signature database to test sinkholing.
   1. Select **Device** > **Dynamic Updates** and in the **Antivirus** section click the **Release Notes** link for the currently installed antivirus database. You can also find the antivirus release notes that list the incremental signature updates under Dynamic Updates on the Palo Alto Networks support site.
   2. In the second column of the release note, locate a line item with a domain extension (for example, .com, .edu, or .net). The left column will display the domain name. For example, Antivirus release 1117-1560, includes an item in the left column named "tbsbana" and the right column lists "net".

      The following shows the content in the release note for this line item:

      ```
      conficker:tbsbana 1
      variants: net
      ```
2. From the client host, open a command prompt.
3. Perform an NSLOOKUP to a URL that you identified as a known malicious domain.

   For example, using the URL `track.bidtrk.com`:

   ```
   C:\>nslookup
    track.bidtrk.com
    Server: my-local-dns.local
    Address: 10.0.0.222
    Non-authoritative answer:
    Name: track.bidtrk.com.org
    Addresses: fd97:3dec:4d27:e37c:5:5:5:510.15.0.20
   ```

   In the output, note that the NSLOOKUP to the malicious domain has been forged using the sinkhole IP addresses that we configured (10.15.0.20). Because the domain matched a malicious DNS signature, the sinkhole action was performed.
4. Select **Monitor** > **Logs** > **Threat** and locate the corresponding threat log entry to verify that the correct action was taken on the NSLOOKUP request.
5. Perform a ping to **`track.bidtrk.com`**, which will generate network traffic to the sinkhole address.

See Infected Hosts that Attempted to Connect to a Malicious Domain

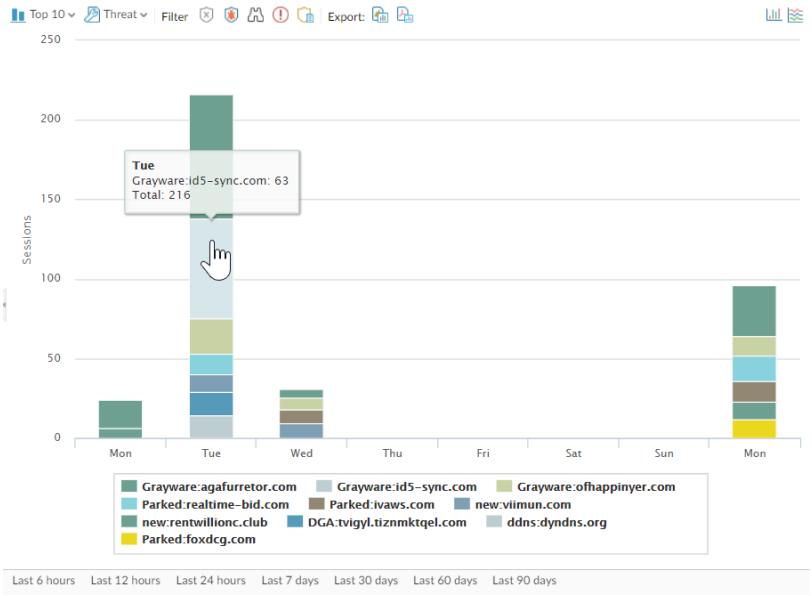
After you have configured DNS sinkholing and verified that traffic to a malicious domain goes to the sinkhole address, you should regularly monitor traffic to the sinkhole address, so that you can track down the infected hosts and eliminate the threat.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

- Use App Scope to identify infected client hosts.
  1. Select **Monitor** > **App Scope** and select **Threat Monitor**.
  2. Click the **Show spyware** button along the top of the display page.
  3. Select a time range.

     The following screenshot shows three instances of Suspicious DNS queries, which were generated when the test client host performed an NSLOOKUP on a known malicious domain. Click the graph to see more details about the event.



- Configure a custom report to identify all client hosts that have sent traffic to the sinkhole IP address, which is 10.15.0.20 in this example.

  Forward to an SNMP manager, Syslog server and/or Panorama to enable alerts on these events.

  In this example, the infected client host performed an NSLOOKUP to a known malicious domain that is listed in the Palo Alto Networks DNS Signature database. When this occurred, the query was sent to the local DNS server, which then forwarded the request through the firewall to an external DNS server. The firewall security policy with the Anti-Spyware profile configured matched the query to the DNS Signature database, which then forged the reply using the sinkhole address of 10.15.0.20 and fd97:3dec:4d27:e37c:5:5:5:5. The client attempts to start a session and the traffic log records the activity with the source host and the destination address, which is now directed to the forged sinkhole address.

  Viewing the traffic log on the firewall allows you to identify any client host that is sending traffic to the sinkhole address. In this example, the logs show that the source address 192.168.2.10 sent the malicious DNS query. The host can then be found and cleaned. Without the DNS sinkhole option, the administrator would only see the local DNS server as the system that performed the query and would not see the client host that is infected. If you attempted

https://t.me/learningnets

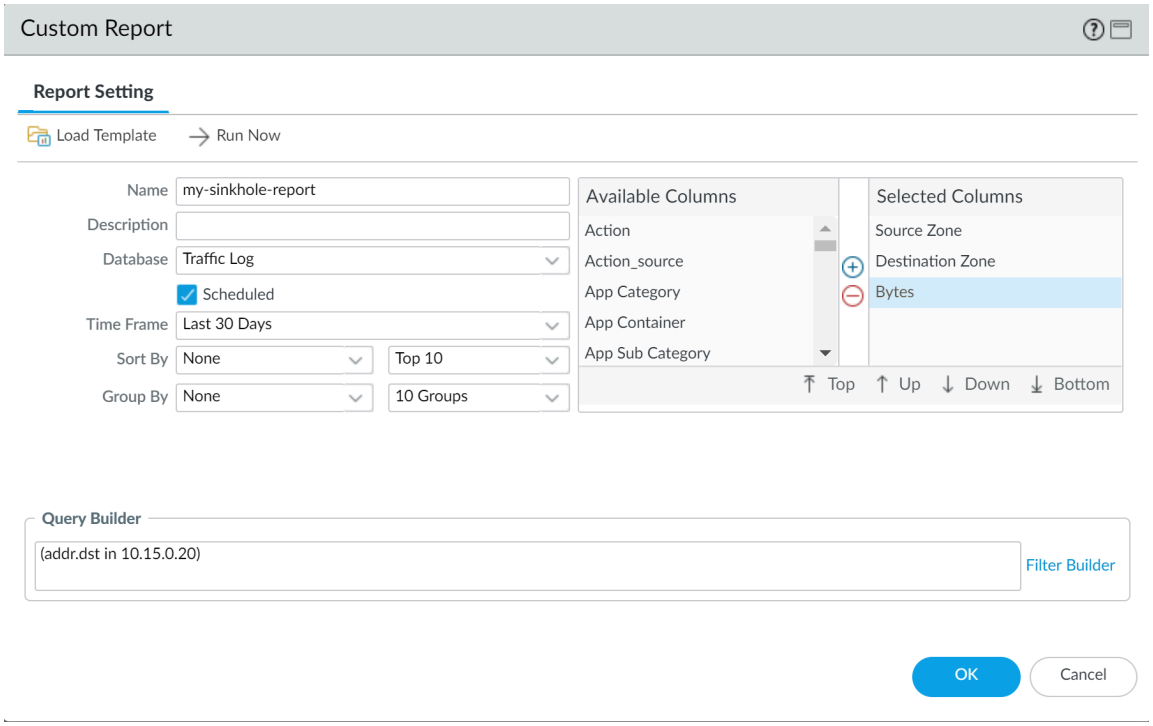
to run a report on the threat log using the action “Sinkhole”, the log would show the local DNS server, not the infected host.

1. Select **Monitor** > **Manage Custom Reports**.
2. Click **Add** and **Name** the report.
3. Define a custom report that captures traffic to the sinkhole address as follows:
   - **Database**—Select **Traffic Log**.
   - **Scheduled**—Enable **Scheduled** and the report will run every night.
   - **Time Frame**—30 days
   - **Selected Columns**—Select **Source address** or **Source User** (if you have User-ID configured), which will identify the infected client host in the report, and **Destination address**, which will be the sinkhole address.
   - In the section at the bottom of the screen, create a custom query for traffic to the sinkhole address (10.15.0.20 in this example). You can either enter the destination address in the **Query Builder** window (**addr.dst in 10.15.0.20**) or select the following

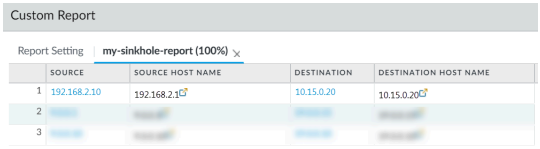
 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

in each column and click **Add**: Connector = and, Attribute = Destination Address, Operator = in, and Value = 10.15.0.20. Click **Add** to add the query.



4. Click **Run Now** to run the report. The report will show all client hosts that have sent traffic to the sinkhole address, which indicates that they are most likely infected. You can now track down the hosts and check them for spyware.



5. To view scheduled reports that have run, select **Monitor** > **Reports**.

https://t.me/learningnets

Data Filtering

Use Data Filtering Profiles to prevent sensitive, confidential, and proprietary information from leaving your network. Predefined patterns, built-in settings, and customizable options make it easy for you to protect files that contain certain file properties (such as a document title or author), credit card numbers, regulated information from different countries (like social security numbers), and third-party data loss prevention (DLP) labels.

- **Predefined Data Patterns**—Easily filter common patterns, including credit card numbers. Predefined data filtering patterns also identify specific (regulated) information from different countries of the world, such as social security numbers (United States), INSEE Identification numbers (France), and New Zealand Internal Revenue Department Identification Numbers. Many of the predefined data filtering patterns enable compliance for standards such as HIPAA, GDPR, Gramm-Leach-Bliley Act.
- **Built-In Support for Azure Information Protection and Titus Data Classification**—Predefined file properties allow you to filter content based on Azure Information Protection and Titus labels. Azure Information Protection labels are stored in metadata, so make sure that you know the GUID of the Azure Information Protect label that you want the firewall to filter.
- **Custom Data Patterns for Data Loss Prevention (DLP) Solutions**—If you're using a third-party, endpoint DLP solution that populates file properties to indicate sensitive content, you can create a custom data pattern to identify the file properties and values tagged by your DLP solution and then log or block the files that your Data Filtering profile detects based on that pattern.

Create a Data Filtering Profile

Data Filtering profiles can keep sensitive information from leaving your network.

To get started, you'll first create a data pattern that specifies the information types and fields that you want the firewall to filter. Then, you attach that pattern to a data filtering profile, which specifies how you want to enforce the content that the firewall filters. Add the data filtering profile to a security policy rule to start filtering traffic matching the rule.

STEP 1 | Define a new data pattern object to detect the information you want to filter.

1. Select **Objects** > **Custom Objects** > **Data Patterns** and **Add** a new object.
2. Provide a descriptive **Name** for the new object.
3. (Optional) Select **Shared** if you want the data pattern to be available to:
   - **Every virtual system (vsys) on a multi-vsys firewall**—If cleared (disabled), the data pattern is available only to the Virtual System selected in the **Objects** tab.
   - **Every device group on Panorama**—If cleared (disabled), the data pattern is available only to the Device Group selected in the **Objects** tab.
4. (Optional—Panorama only) Select **Disable override** to prevent administrators from overriding the settings of this data pattern object in device groups that inherit the object.

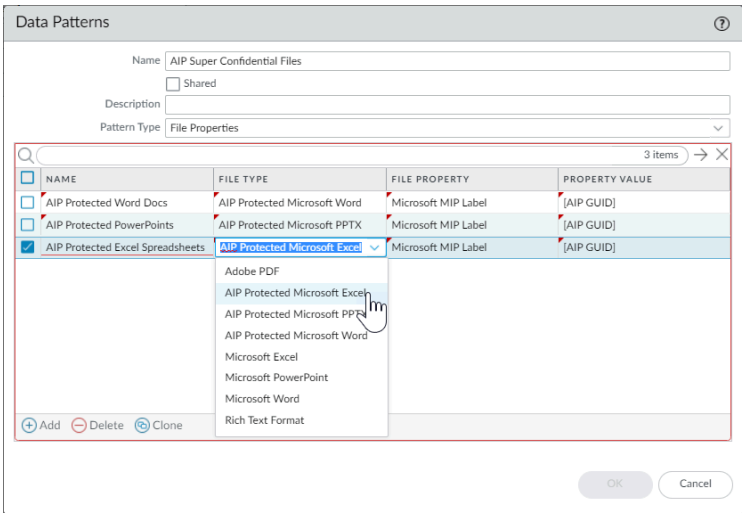
https://t.me/learningnets

This selection is cleared by default, which means administrators can override the settings for any device group that inherits the object.

5. (Optional—Panorama only) Select **Data Capture** to automatically collect the data that is blocked by the filter.

   *Specify a password for Manage Data Protection on the Settings page to view your captured data (**Device** > **Setup** > **Content-ID** > **Manage Data Protection**).*

6. Set the **Pattern Type** to one of the following:
   - **Predefined Pattern**—Filter for credit card, social security numbers, and personally identifiable information for several compliance standards including HIPAA, GDPR, Gramm-Leach-Bliley Act.
   - **Regular Expression**—Filter for custom data patterns.
   - **File Properties**—Filter based on file properties and the associated values.
7. **Add** a new rule to the data pattern object.
8. Specify the data pattern according to the **Pattern Type** you selected for this object:
   - **Predefined**—Select the **Name** and choose the predefined data pattern on which to filter.
   - **Regular Expression**—Specify a descriptive **Name**, select the **File Type** (or types) you want to scan, and then enter the specific **Data Pattern** you want the firewall to detect.
   - **File Properties**—Specify a descriptive **Name**, select the **File Type** and **File Property** you want to scan, and enter the specific **Property Value** that you want the firewall to detect.
     - **To filter Titus classified documents**: Select one of the non-AIP protected file types, and set the **File Property** to TITUS GUID. Enter the Titus label GUID as the **Property Value**.
     - **For Azure Information Protection labeled documents**: Select any **File Type** except Rich Text Format. For the file type you choose, set the **File Property** to Microsoft MIP Label, and enter the Azure Informatin Protect label GUID as the **Property Value**.



9. Click **OK** to save the data pattern.

https://t.me/learningnets

STEP 2 | Add the data pattern object to a data filtering profile.

1. Select **Objects** > **Security Profiles** > **Data Filtering** and **Add** or modify a data filtering profile.
2. Provide a descriptive **Name** for the new profile.
3. **Add** a new profile rule and select the Data Pattern you created in Step .
4. Specify **Applications**, **File Types**, and what **Direction** of traffic (upload or download) you want to filter based on the data pattern.

   The file type you select must be the same file type you defined for the data pattern earlier, or it must be a file type that includes the data pattern file type. For example, you could define both the data pattern object and the data filtering profile to scan all Microsoft Office documents. Or, you could define the data pattern object to match to only Microsoft PowerPoint Presentations while the data filtering profile scans all Microsoft Office documents.

   If a data pattern object is attached to a data filtering profile and the configured file types do not align between the two, the profile will not correctly filter documents matched to the data pattern object.
5. Set the **Alert Threshold** to specify the number of times the data pattern must be detected in a file to trigger an alert.
6. Set the **Block Threshold** to block files that contain at least this many instances of the data pattern.
7. Set the **Log Severity** recorded for files that match this rule.
8. Click **OK** to save the data filtering profile.

STEP 3 | Apply the data filtering settings to traffic.

1. Select **Policies** > **Security** and **Add** or modify a security policy rule.
2. Select **Actions** and set the Profile Type to **Profiles**.
3. Attach the Data Filtering profile you created in Step 2 to the security policy rule.
4. Click **OK**.

STEP 4 | (Recommended) Prevent web browsers from resuming sessions that the firewall has terminated.

This option ensures that when the firewall detects and then drops a sensitive file, a web browser cannot resume the session in an attempt to retrieve the file.

1. Select **Device** > **Setup** > **Content-ID** and edit Content-ID Settings.
2. Clear the **Allow HTTP partial response**.
3. Click **OK**.

STEP 5 | Monitor files that the firewall is filtering.

Select **Monitor** > **Data Filtering** to view the files that the firewall has detected and blocked based on your data filtering settings.

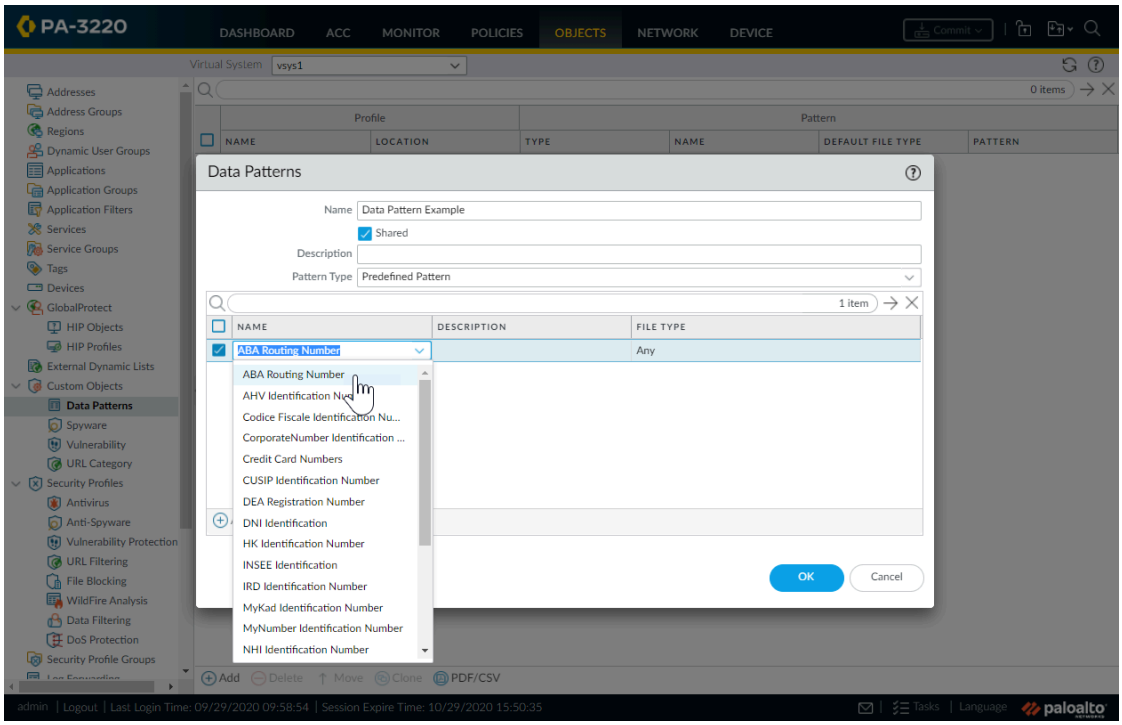
 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Predefined Data Filtering Patterns

To comply with standards such as HIPAA, GDPR, and the Gramm-Leach-Bliley Act, the firewall provides predefined data patterns. You can use these patterns to prevent common types of sensitive information, like credit cards and social security numbers, from leaving your network.

You can find predefined data patterns by selecting **Objects > Custom Objects > Data Patterns** and clicking **Add** a new object. Then, set the **Pattern Type** to **Predefined Pattern** and **Add** a new rule to the data pattern object. Select a data pattern from the list that appears under **Name**.



If the type of information you want to protect is not covered in the list of predefined patterns, you can use regular expressions to create custom patterns.

The following is a list of available data patterns:

Pattern	Description
Credit Card Numbers	16-digit credit card numbers
Social Security Numbers	9-digit social security numbers with dashes
Social Security Numbers (without dash separator)	9-digit social security numbers without dashes
ABA Routing Number	The American Banking Association Routing Number
AHV Identification Number	Swiss Alters und Hinterlassenenversicherungsnummer

©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Pattern	Description
Codice Fiscale Identification Number	Italian Fiscal Tax Code Card Identification Number
CorporateNumber Identification Number	Japanese National Tax Agency Corporate Number
CUSIP Identification Number	Committee on Uniform Security Identification Procedures Identification Number
DEA Registration Number	U.S. Drug Enforcement Administration Registration Number
DNI Identification Number	Spanish Documento nacional de identidad Identification Number number
HK Identification Number	Hong Kong Residents Identification Number
INSEE Identification Number	French National Institute of Statistics and Economic Studies identification number
IRD Identification Number	New Zealand Internal Revenue Department Identification Number
MyKad Identification Number	Malaysia MyKad Identity Card Identification Number
MyNumber Identification Number	Japanese Social Security and Tax Number System Identification Number
NHI Identification Number	New Zealand National Health Index Number
NIF Identification Number	Spanish Tax Identification Number
NIN Identification Number	Taiwan Identification Card Number
NRIC Identification Number	Singapore National Registration Identity Card Identification Number
Permanent Account Identification Number	India Permanent Account Number of Indian nationals
PRC Identification Number	People's Republic of China Resident Identification Number
PRN Identification Number	Republic of South Korea Resident Registration Number

©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Pattern	Description
Republic of South Korea Resident Registration	Republic of South Korea Resident Registration Number

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

WildFire Inline ML

The WildFire inline ML option present in the Antivirus profile enables the firewall dataplane to apply machine learning on PE (portable executable), ELF (executable and linked format) and MS Office files, and PowerShell and shell scripts in real-time. This layer of antivirus protection complements the WildFire-based signatures to provide extended coverage for files of which signatures do not already exist. Each inline ML model dynamically detects malicious files of a specific type by evaluating file details, including decoder fields and patterns, to formulate a high probability classification of a file. This protection extends to currently unknown as well as future variants of threats that match characteristics that Palo Alto Networks has identified as malicious. To keep up with the latest changes in the threat landscape, inline ML models are added or updated via content releases. Before you can enable WildFire inline ML, you must possess an active WildFire subscription.

Inline ML-based protection can also be enabled to detect malicious URLs in real-time as part of your URL Filtering configuration. For more information, refer to: URL Filtering Inline ML

WildFire inline ML is not supported on the VM-50 or VM50L virtual appliance.

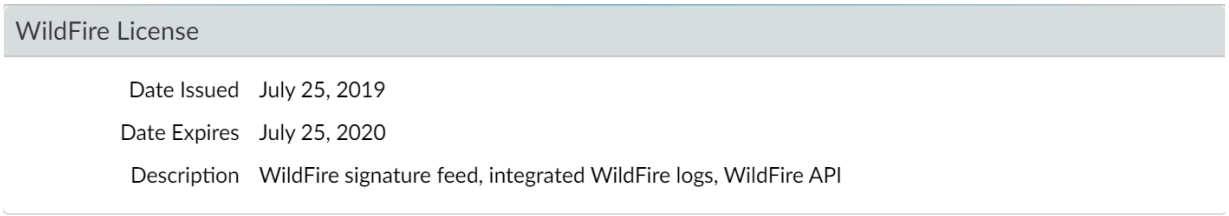
Configure WildFire Inline ML

To enable your WildFire inline ML configuration, attach the Antivirus profile configured with the inline ML settings to a security policy rule (see Set Up Antivirus, Anti-Spyware, and Vulnerability Protection).

WildFire inline ML is not currently supported on the VM-50 or VM50L virtual appliance.

STEP 1 | To take advantage of WildFire inline ML, you must have an active WildFire subscription to analyze Windows executables.

Verify that you have a WildFire subscription. To verify which subscriptions that you currently have licenses for, select **Device** > **Licenses** and verify that the appropriate licenses display and have not expired.



 ©2022 Palo Alto Networks, Inc.

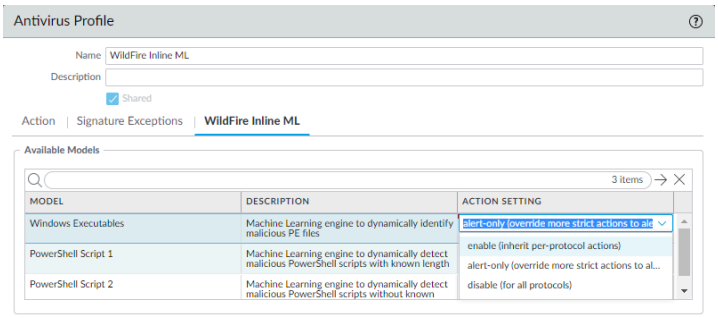
https://t.me/learningnets

STEP 2 | Create a new or update your existing Antivirus security profile(s) to use the real-time WildFire inline ML models.

1. Select an existing **Antivirus Profile** or create a new one (select **Objects > Security Profiles > Antivirus** and **Add** a new profile.
2. Configure your Antivirus profile.
3. Select the **WildFire Inline ML** tab and apply an **Action Setting** for each WildFire Inline ML model. This enforces the WildFire Inline ML Actions settings configured for each protocol on a per model basis. The following classification engines available: Windows Executables, PowerShell Scripts 1, and PowerShell Scripts 2.

   The following additional analysis models are available upon installation of the specified content update:

   - *Executable Linked Format (available with installation of PAN-OS content release 8367 and later)*
   - *MSOffice (available with installation of PAN-OS content release 8434 and later)*
   - *Shell Scripts (available with installation of PAN-OS content release 8543 and later)*



   - **enable (inherit per-protocol actions)**—WildFire inspects traffic according to your selections in the WildFire Inline ML Action column in the decoders section of the **Action** tab.
   - **alert-only (override more strict actions to alert)**—WildFire inspects traffic according to your selections in the WildFire Inline ML Action column in the decoders section of the **Action** tab and overrides any action with a severity level higher than `alert` (`drop`, `reset-client`, `reset-server`, `reset-both`) `alert`, which allows traffic to pass while still generating and saving an alert in the threat logs.
   - **disable (for all protocols)**—WildFire allows traffic to pass without any policy action.
4. Click **OK** to exit the Antivirus Profile configuration window and **Commit** your new settings.

STEP 3 | (Optional) Add file exceptions to your Antivirus security profile if you encounter false-positives. This is typically done for users who are not forwarding files to WildFire for analysis.

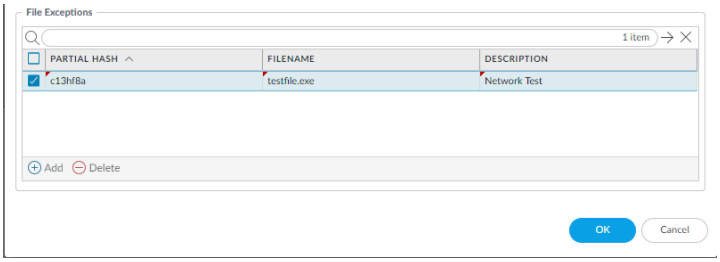
 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

You can add the file exception details directly to the exception list or by specifying a file from the threat logs.

If your WildFire Analysis security profile is configured to forward the filetypes analyzed using WildFire inline ML, false-positives are automatically corrected as they are received. If you continue to see ml-virus alerts for files that have been classified as benign by WildFire Analysis, please contact Palo Alto Networks Support.

- Add file exceptions directly to the exception list.
  1. Select **Objects > Security Profiles > Antivirus**.
  2. Select an Antivirus profile for which you want to exclude specific files and then select **WildFire Inline ML**.
  3. Add the hash, filename, and description of the file that you want to exclude from enforcement.



  4. Click **OK** to save the Antivirus profile and then **Commit** your updates.
- Add file exceptions from threat logs entries.
  1. Select **Monitor > Logs > Threat** and filter the logs for the **ml-virus** threat type. Select a threat log for a file that you wish to create a file exception for.
  2. Go to the **Detailed Log View** and scroll down to the **Details** pane then select **Create Exception**.



  3. Add a **Description** and click **OK** to add the file exception.
  4. The new file exception can be found **File Exceptions** list under **Objects > Security Profiles > Antivirus > WildFire Inline ML**.

STEP 4 | (Optional) Verify the status of your firewall's connectivity to the Inline ML cloud service.

Use the following CLI command on the firewall to view the connection status.

```
show mlav cloud-status
```

For example:

```
show mlav cloud-status

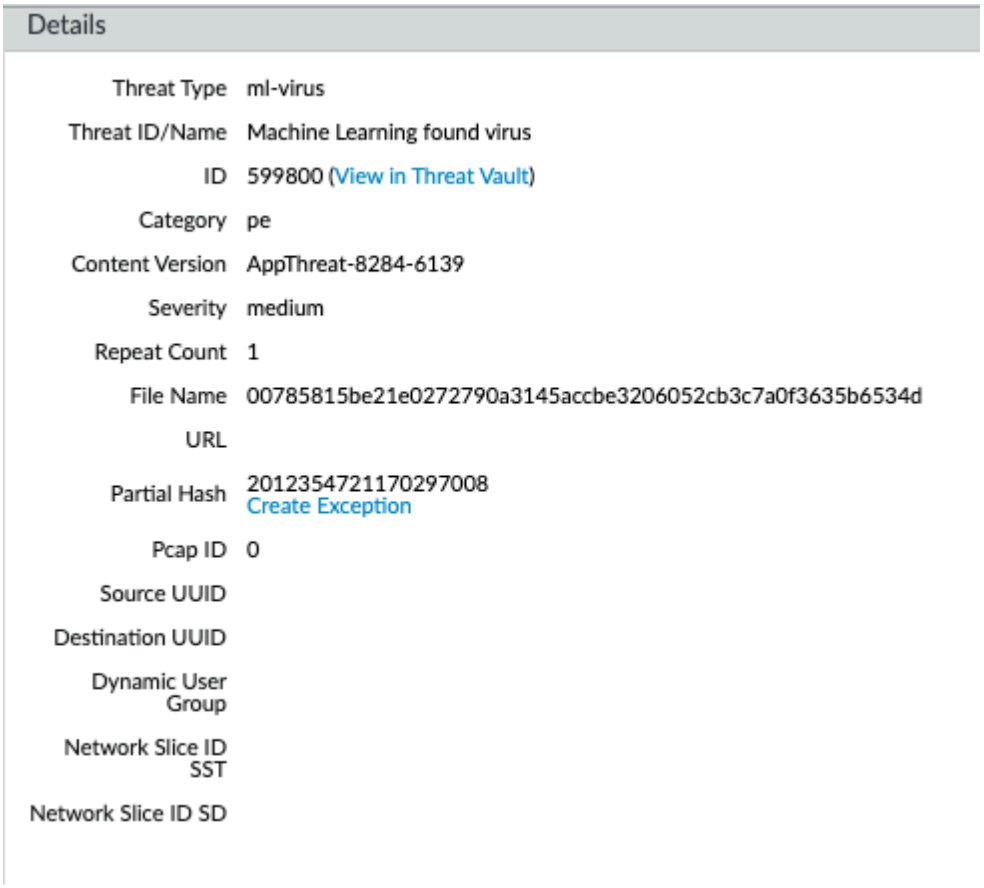
MLAV cloud
Current cloud server:          ml.service.paloaltonetworks.com
```


https://t.me/learningnets

```
Cloud connection:                connected
```

If you are unable to connect to the Inline ML cloud service, verify that the following domain is not being blocked: ml.service.paloaltonetworks.com.

To view information about files that have been detected using WildFire Inline ML, examine the threat logs (**Monitor > Logs > Threat**, then select the log type from the list). Files that have been analyzed using WildFire inline ML are labeled with the threat type **ml-virus**:



 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Set Up File Blocking

File Blocking Profiles allow you to identify specific file types that you want to want to block or monitor. For most traffic (including traffic on your internal network), block files that are known to carry threats or that have no real use case for upload/download. Currently, these include batch files, DLLs, Java class files, help files, Windows shortcuts (.lnk), and BitTorrent files. Additionally, to provide drive-by download protection, allow download/upload of executables and archive files (.zip and .rar), but force users to acknowledge that they are transferring a file so that they notice that the browser is attempting to download something they were not aware of. For policy rules that allow general web browsing, be stricter with your file blocking because the risk of users unknowingly downloading malicious files is much higher. For this type of traffic, attach a more strict file blocking profile that also blocks portable executable (PE) files.

You can define your own custom File Blocking profiles or choose one of the following predefined profiles when applying file blocking to a Security policy rule. You can clone and edit the predefined profiles, which are available with content release version 653 and later, and then follow File Blocking profile safe transition steps to preserve application availability as you transition to best practice file blocking settings:

- **basic file blocking**—Attach this profile to the Security policy rules that allow traffic to and from less sensitive applications to block files that are commonly included in malware attack campaigns or that have no real use case for upload/download. This profile blocks upload and download of PE files (.scr, .cpl, .dll, .ocx, .pif, .exe) , Java files (.class, .jar), Help files (.chm, .hlp) and other potentially malicious file types, including .vbe, .hta, .wsf, .torrent, .7z, .rar, .bat. Additionally, it prompts users to acknowledge when they attempt to download encrypted-rar or encrypted-zip files. This rule alerts on all other file types to give you complete visibility into all file types coming in and out of your network.
- **strict file blocking**—Use this stricter profile on the Security policy rules that allow access to your most sensitive applications. This profile blocks the same file types as the other profile, and additionally blocks flash, .tar, multi-level encoding, .cab, .msi, encrypted-rar, and encrypted-zip files.

These predefined profiles are designed to provide the most secure posture for your network. However, if you have business-critical applications that rely on some of the applications that are blocked in these default profiles, you can clone the profiles and modify them as necessary. Make sure you only use the modified profiles for those users who need to upload and/or download a risky file type. Additionally, to reduce your attack surface, make sure you are using other security measures to ensure that the files your users are uploading and downloading do not pose a threat to your organization. For example, if you must allow download of PE files, make sure you are sending all unknown PE files to WildFire foranalysis. Additionally, maintain a strict URL filtering policy to ensure that users cannot download content from web sites that have been known to host malicious content.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 1 | Create the file blocking profile.

1. Select **Objects** > **Security Profiles** > **File Blocking** and **Add** a profile.
2. Enter a **Name** for the file blocking profile such as **`Block_EXE`**.
3. (Optional) Enter a **Description**, such as **`Block users from downloading exe files from websites`**.
4. (Optional) Specify that the profile is **Shared** with:
   - **Every virtual system (vsys) on a multi-vsys firewall**—If cleared (disabled), the profile is available only to the Virtual System selected in the **Objects** tab.
   - **Every device group on Panorama**—If cleared (disabled), the profile is available only to the Device Group selected in the **Objects** tab.
5. (Optional—Panorama only) Select **Disable override** to prevent administrators from overriding the settings of this file blocking profile in device groups that inherit the profile. This selection is cleared by default, which means administrators can override the settings for any device group that inherits the profile.

STEP 2 | Configure the file blocking options.

1. **Add** and define a rule for the profile.
2. Enter a **Name** for the rule, such as **`BlockEXE`**.
3. Select **Any** or specify one or more specific **Applications** for filtering, such as **web-browsing**.

   Only web browsers can display the response page (continue prompt) that allows users to confirm their Choosing any other application results in blocked traffic for those applications because there is no prompt displayed to allow users to continue.
4. Select **Any** or specify one or more specific **File Types**, such as **exe**.
5. Specify the **Direction**, such as **download**.
6. Specify the **Action** (**alert**, **block**, or **continue**). For example, select **continue** to prompt users for confirmation before they are allowed to download an executable (.exe) file. Alternatively, you could **block** the specified files or you could configure the firewall to simply trigger an **alert** when a user downloads an executable file.
7. Click **OK** to save the profile.

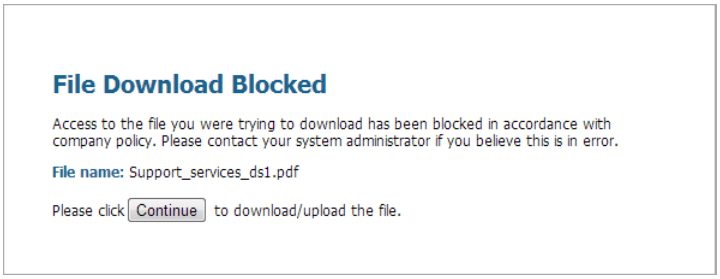
STEP 3 | Apply the file blocking profile to a security policy rule.

1. Select **Policies** > **Security** and either select an existing policy rule or **Add** a new rule as described in Set Up a Basic Security Policy.
2. On the **Actions** tab, select the file blocking profile you configured in the previous step. In this example, the profile name is **`Block_EXE`**.
3. **Commit** your configuration.

STEP 4 | To test your file blocking configuration, access an endpoint PC in the trust zone of the firewall and attempt to download an executable file from a website in the untrust zone; a response page should display. Click **Continue** to confirm that you can download the file. You

https://t.me/learningnets

can also set other actions, such as **alert** or **block**, which do not provide an option for the user to continue the download. The following shows the default response page for File Blocking:



STEP 5 | (Optional) Define custom file blocking response pages (**Device** > **Response Pages**). This allows you to provide more information to users when they see a response page. You can include information such as company policy information and contact information for a Helpdesk.

*When you create a file blocking profile with the **continue** action, you can choose only the **web-browsing** application. If you choose any other application, traffic that matches the security policy will not flow through the firewall because users are not prompted with an option to continue. Additionally, you need to configure and enable a decryption policy for HTTPS websites.*

*Check your logs to determine the application used when you test this feature. For example, if you are using Microsoft SharePoint to download files, even though you are using a web-browser to access the site, the application is actually `sharepoint-base`, or `sharepoint-document`. (It can help to set the application type to **Any** for testing.)*

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Prevent Brute Force Attacks

A brute force attack uses a large volume of requests/responses from the same source or destination IP address to break into a system. The attacker employs a trial-and-error method to guess the response to a challenge or a request.

The Vulnerability Protection profile on the firewall includes signatures to protect you from brute force attacks. Each signature has an ID, Threat Name, and Severity and is triggered when a pattern is recorded. The pattern specifies the conditions and interval at which the traffic is identified as a brute-force attack; some signatures are associated with another child signature that is of a lower severity and specifies the pattern to match against. When a pattern matches against the signature or child signature, it triggers the default action for the signature.

To enforce protection:

- Attach the Vulnerability Protection profile to a Security policy rule. See Set Up Antivirus, Anti-Spyware, and Vulnerability Protection.
- Install content updates that include new signatures to protect against emerging threats. See Install Content and Software Updates.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Customize the Action and Trigger Conditions for a Brute Force Signature

The firewall includes two types of predefined brute force signatures—parent signatures and child signatures. A child signature is a single occurrence of a traffic pattern that matches the signature. A parent signature is associated with a child signature and is triggered when multiple events occur within a specified time interval and that matches the traffic pattern defined in the child signature.

Typically, the default action for a child signature is *allow* because a single event is not indicative of an attack. This ensures that legitimate traffic is not blocked and avoids generating threat logs for non-noteworthy events. Palo Alto Networks recommends that you do not change the default action without careful consideration.

In most cases, the brute force signature is a noteworthy event due to its recurrent pattern. If needed, you can do one of the following to customize the action for a brute-force signature:

- Create a rule to modify the default action for all signatures in the brute force category. You can choose to allow, alert, block, reset, or drop the traffic.
- Define an exception for a specific signature. For example, you can search for and define an exception for a CVE.

For a parent signature, you can modify both the trigger conditions and the action; for a child signature, you can modify only the action.



To effectively mitigate an attack, specify the block-ip address action instead of the drop or reset action for most brute force signatures.

STEP 1 | Create a new Vulnerability Protection profile.

1. Select **Objects** > **Security Profiles** > **Vulnerability Protection** and **Add** a profile.
2. Enter a **Name** for the Vulnerability Protection profile.
3. (Optional) Enter a **Description**.
4. (Optional) Specify that the profile is **Shared** with:
   - **Every virtual system (vsys) on a multi-vsys firewall**—If cleared (disabled), the profile is available only to the Virtual System selected in the **Objects** tab.
   - **Every device group on Panorama**—If cleared (disabled), the profile is available only to the Device Group selected in the **Objects** tab.
5. (Optional—Panorama only) Select **Disable override** to prevent administrators from overriding the settings of this Vulnerability Protection profile in device groups that inherit the profile. This selection is cleared by default, which means administrators can override the settings for any device group that inherits the profile.

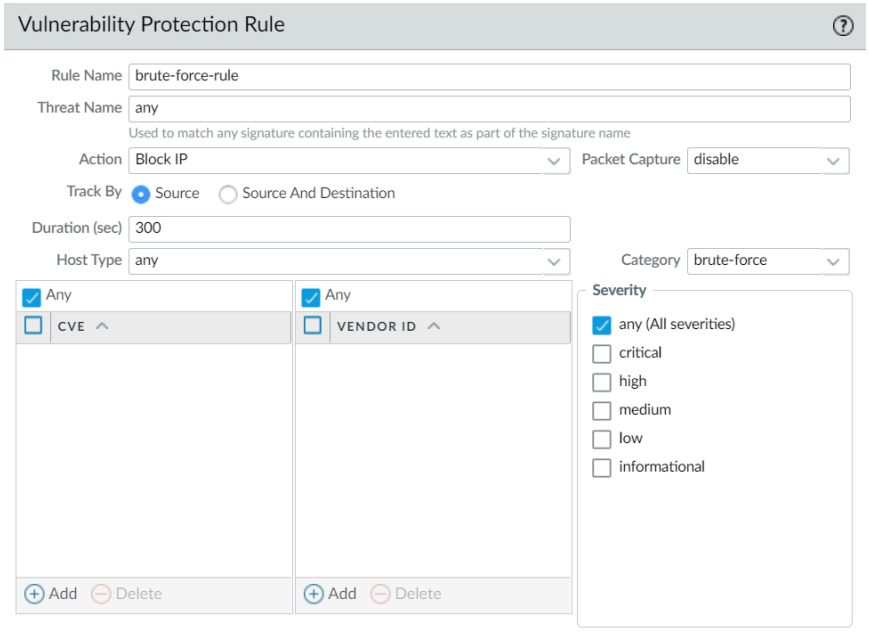
 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 2 | Create a rule that defines the action for all signatures in a category.

1. On the **Rules** tab, **Add** and enter a **Rule Name** for a new rule.
2. (Optional) Specify a specific threat name (default is **any**).
3. Set the **Action**. In this example, it is set to **Block IP**.

   If you set a Vulnerability Protection profile to Block IP, the firewall first uses hardware to block IP addresses. If attack traffic exceeds the blocking capacity of the hardware, the firewall then uses software blocking mechanisms to block the remaining IP addresses.
4. Set **Category** to **brute-force**.
5. (Optional) If blocking, specify the **Host Type** on which to block: **server** or **client** (default is **any**).
6. See Step 3 to customize the action for a specific signature.
7. See Step 4 to customize the trigger threshold for a parent signature.



8. Click **OK** to save the rule and the profile.

 ©2022 Palo Alto Networks, Inc.

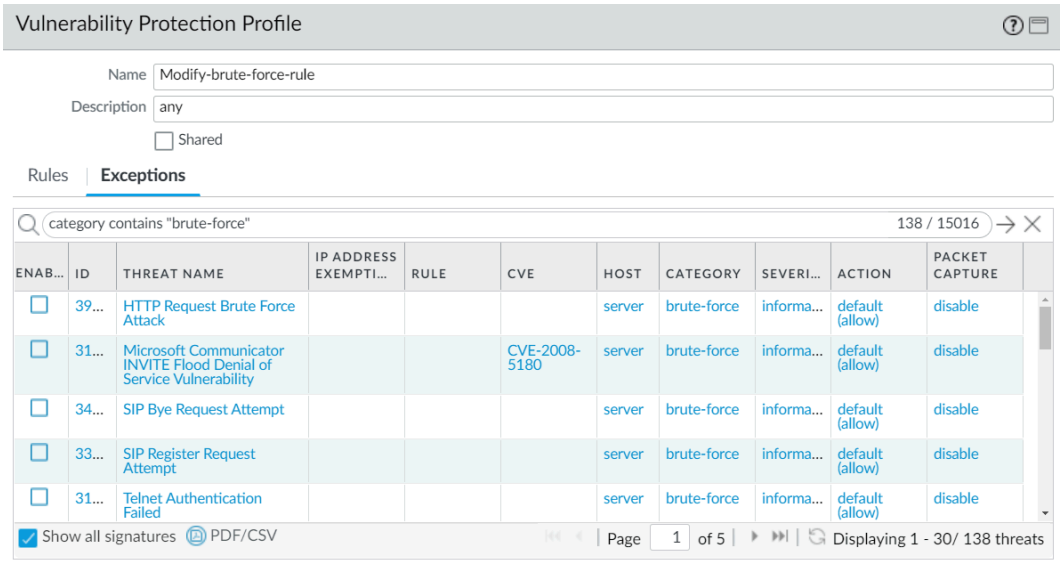
https://t.me/learningnets

STEP 3 | (Optional) Customize the action for a specific signature.

1. On the **Exceptions** tab, **Show all signatures** to find the signature you want to modify.

   To view all the signatures in the brute-force category, search for `category contains 'brute-force'`.

2. To edit a specific signature, click the predefined default action in the Action column.



3. Set the action: **Allow**, **Alert**, **Block Ip**, or **Drop**. If you select **Block Ip**, complete these additional tasks:

   1. Specify the **Time** period (in seconds) after which to trigger the action.
   2. Specify whether to **Track By** and block the IP address using the **IP source** or the **IP source and destination**.

4. Click **OK**.
5. For each modified signature, select the check box in the **Enable** column.
6. Click **OK**.

STEP 4 | Customize the trigger conditions for a parent signature.

A parent signature that can be edited is marked with this icon: .

In this example, the search criteria was brute force category and CVE-2008-1447.

1. Edit ( ) the time attribute and the aggregation criteria for the signature.
2. To modify the trigger threshold, specify the **Number of Hits** per number of **seconds**.
3. Specify whether to aggregate the number of hits (**Aggregation Criteria**) by **source**, **destination**, or **source-and-destination**.
4. Click **OK**.

https://t.me/learningnets

STEP 5 | Attach this new profile to a Security policy rule.

1. Select **Policies** > **Security** and **Add** or modify a Security policy rule.
2. On the **Actions** tab, select **Profiles** as the **Profile Type** for the Profile Setting.
3. Select your **Vulnerability Protection** profile.
4. Click **OK**.

STEP 6 | Commit your changes.

1. Click **Commit**.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Enable Evasion Signatures

Palo Alto Networks evasion signatures detect crafted HTTP or TLS requests, and can alert to instances where a client connects to a domain other than the domain specified in a DNS query. Evasion signatures are effective only when the firewall is also enabled to act as a DNS proxy and resolve domain name queries. As a best practice, take the following steps to enable evasion signatures.

STEP 1 | Enable a firewall intermediate to clients and servers to act as a DNS proxy.

Configure a DNS Proxy Object, including:

- Specify the interfaces on which you want the firewall to listen for DNS queries.
- Define the DNS servers with which the firewall communicates to resolve DNS requests.
- Set up static FQDN-to-IP address entries that the firewall can resolve locally, without reaching out to DNS servers.
- Enable caching for resolved hostname-to-IP-address mappings.

STEP 2 | Get the latest Applications and Threats content version (at least content version 579 or later).

1. Select **Device** > **Dynamic Updates**.
2. **Check Now** to get the latest Applications and Threats content update.
3. Download and Install Applications and Threats content version 579 (or later).

STEP 3 | Define how the firewall should enforce traffic matched to evasion signatures.

1. Select **Objects** > **Security Profiles** > **Anti-Spyware** and **Add** or modify an Anti-spyware profile.
2. Select **Exceptions** and select **Show all signatures**.
3. Filter signatures based on the keyword `evasion`.
4. For all evasion signatures, set the **Action** to any setting other than allow or the default action (the default action is for evasion signatures is allow). For example, set the **Action** for signature IDs 14978 and 14984 to **alert** or **drop**.
5. Click **OK** to save the updated Anti-spyware profile.
6. Attach the Anti-spyware profile to a security policy rule: Select **Policies** > **Security**, select the desired policy to modify and then click the **Actions** tab. In Profile Settings, click the drop-down next to **Anti-Spyware** and select the anti-spyware profile you just modified to enforce evasion signatures.

STEP 4 | Commit your changes.

Click **Commit**.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Monitor Blocked IP Addresses

The firewall maintains a block list of source IP addresses that it's blocking. When the firewall blocks a source IP address, such as when you configure either of the following policy rules, the firewall blocks that traffic in hardware before those packets use CPU or packet buffer resources:

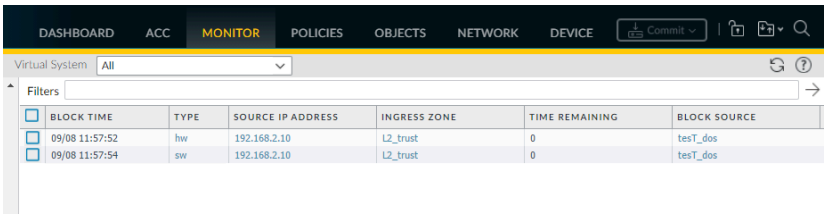
- A classified DoS Protection policy rule with the action to **Protect** (a classified DoS Protection policy specifies that incoming connections match a source IP address, destination IP address, or source and destination IP address pair, and is associated with a Classified DoS Protection profile, as described in DoS Protection Against Flooding of New Sessions).
- A Security Policy rule that uses a Vulnerability Protection profile

Hardware IP address blocking is supported on PA-3200 Series, PA-5200 Series, and PA-7000 Series firewalls.

You can view the block list, get detailed information about an IP address on the block list, or view counts of addresses that hardware and software are blocking. You can delete an IP address from the list if you think it shouldn't be blocked. You can change the source of detailed information about addresses on the list. You can also change how long hardware blocks IP addresses.

- View block list entries.
  1. Select **Monitor** > **Block IP List**.

     Entries on the block list indicate in the Type column whether they were blocked by hardware (hw) or software (sw).
  2. View at the bottom of the screen:
     - Count of **Total Blocked IPs** out of the number of blocked IP addresses the firewall supports.
     - Percentage of the block list the firewall has used.
  3. To filter the entries displayed, select a value in a column (which creates a filter in the **Filters** field) and Apply Filter (→). Otherwise, the firewall displays the first 1,000 entries.
  4. Enter a **Page** number or click the arrows at the bottom of the screen to advance through pages of entries.
  5. To view details about an address on the block list, hover over a Source IP address and click the down arrow link. Click the **Who Is** link, which displays Network Solutions Who Is information about the address.



 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

- Delete block list entries.

  Delete an entry if you determine the IP address shouldn't be blocked. Then revise the policy rule that caused the firewall to block the address.

  1. Select **Monitor** > **Block IP List**.
  2. Select one or more entries and click **Delete**.
  3. (Optional) Select **Clear All** to remove all entries from the list.

- Disable or re-enable hardware IP address blocking for troubleshooting purposes.

  While hardware IP address blocking is disabled, the firewall still performs any software IP address blocking you have configured.

  ```
  > set system setting hardware-acl-blocking [enable | disable]
  ```

  To conserve CPU and packet buffer resources, leave hardware IP address blocking enabled unless Palo Alto Networks technical support asks you to disable it, for example, if they are debugging a traffic flow.

- Tune the number of seconds that IP addresses blocked by hardware remain on the block list (range is 1-3,600; default is 1).

  ```
  > set system setting hardware-acl-blocking duration <seconds>
  ```

  Maintain a shorter duration for hardware block list entries than software block list entries to reduce the likelihood of exceeding the blocking capacity of the hardware.

- Change the default website for finding more information about an IP address from Network Solutions Who Is to a different website.

  ```
  # set deviceconfig system ip-address-lookup-url <url>
  ```

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

- View counts of source IP addresses blocked by hardware and software, for example to see the rate of an attack.

  View the total sum of IP address entries on the hardware block table and block list (blocked by hardware and software):

  ```
  > show counter global name flow_dos_blk_num_entries
  ```

  View the count of IP address entries on the hardware block table that were blocked by hardware:

  ```
  > show counter global name flow_dos_blk_hw_entries
  ```

  View the count of IP address entries on the block list that were blocked by software:

  ```
  > show counter global name flow_dos_blk_sw_entries
  ```

- View block list information per slot on a PA-7000 Series firewall.

  ```
  > show dos-block-table software filter slot <slot-number>
  ```

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Threat Signature Categories

There are three types of Palo Alto Networks threat signatures, each designed to detect different types of threats as the firewall scans network traffic:

- Antivirus signatures—Detect viruses and malware found in executables and file types.
- Anti-spyware signatures—Detects command-and-control (C2) activity, where spyware on an infected client is collecting data without the user's consent and/or communicating with a remote attacker.
- Vulnerability signatures—Detects system flaws that an attacker might otherwise attempt to exploit.

A signature's severity indicates the risk of the detected event, and a signature's default action (for example, block or alert) is how Palo Alto Networks recommends that you enforce matching traffic.

You must Set Up Antivirus, Anti-Spyware, and Vulnerability Protection to tell the firewall what action to take when it detects a threat, and you can easily use the default security profiles to start blocking threats based on Palo Alto Networks recommendations. For each signature type, category, and even specific signatures you can continue to modify or create new profiles to more granularly enforce potential threats.

The following table lists all possible signature categories by type—Antivirus, Spyware, and Vulnerability—and includes the content update (Applications and Threats, Antivirus, or WildFire) that provides the signatures in each category. You can also go to the Palo Alto Networks Threat Vault to Learn More About Threat Signatures.

Threat Category	Content Update that Provides These Signatures	Description
Antivirus Signatures		
apk	Antivirus WildFire or WildFire Private	Malicious Android Application (APK) files.
dmg	Antivirus WildFire or WildFire Private	Malicious Apple disk image (DMG) files, that are used with Mac OS X.
flash	Antivirus Wildfire or WildFire Private	Adobe Flash applets and Flash content embedded in web pages.
java-class	Antivirus	Java applets (JAR/class file types).

©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Threat Category	Content Update that Provides These Signatures	Description
macho	Antivirus Wildfire or WildFire Private	Mach object files (Mach-O) are executables, libraries, and object code that are native to Mac OS X.
office	Antivirus Wildfire or WildFire Private	Microsoft Office files, including documents (DOC, DOCX, RTF), workbooks (XLS, XLSX), and PowerPoint presentations (PPT, PPTX).
openoffice	Antivirus Wildfire or WildFire Private	Office Open XML (OOXML) 2007+ documents.
pdf	Antivirus Wildfire or WildFire Private	Portable Document Format (PDF) files.
pe	Antivirus Wildfire or WildFire Private	Portable executable (PE) files can automatically execute on a Microsoft Windows system and should be only allowed when authorized. These files types include: • Object code. • Fonts (FONs). • System files (SYS). • Driver files (DRV). • Windows control panel items (CPLs). • DLLs (dynamic-link libraries). • OCXs (libraries for OLE custom controls, or ActiveX controls). • SCRs (scripts that can be used to execute other files). • Extensible Firmware Interface (EFI) files, which run between an OS and firmware in order to facilitate device updates and boot operations. • Program information files (PIFs).
pkg	Antivirus Wildfire or WildFire Private	Apple software installer packages (PKGs), used with Mac OS X.
Spyware Signatures		

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Threat Category	Content Update that Provides These Signatures	Description
adware	Applications and Threats	Detects programs that display potentially unwanted advertisements. Some adware modifies browsers to highlight and hyperlink the most frequently searched keywords on web pages-these links redirect users to advertising websites. Adware can also retrieve updates from a command-and-control (C2) server and install those updates in a browser or onto a client system. Newly-released protections in this category are rare.
autogen	Antivirus	These payload-based signatures detect command-and-control (C2) traffic and are automatically-generated. Importantly, autogen signatures can detect C2 traffic even when the C2 host is unknown or changes rapidly.
backdoor	Applications and Threats	Detects a program that allows an attacker to gain unauthorized remote access to a system.
botnet	Applications and Threats	Indicates botnet activity. A botnet is a network of malware-infected computers (“bots”) that an attacker controls. The attacker can centrally command every computer in a botnet to simultaneously carry out a coordinated action (like launching a DoS attack, for example).
browser-hijack	Applications and Threats	Detects a plugin or software that is modifying browser settings. A browser hijacker might take over auto search or track users’ web activity and send this information to a C2 server. Newly-released protections in this category are rare.
cryptominer	Applications and Threats	(Sometimes known as cryptojacking or miners) Detects the download attempt or network traffic generated from malicious programs designed to use computing resources to mine cryptocurrencies without the user's knowledge. Cryptominer binaries are frequently delivered by a shell script downloader that attempts to determine system architecture and kill other miner processes on the system. Some miners execute within other processes, such as a web browser rendering a malicious web page.
data-theft	Applications and Threats	Detects a system sending information to a known C2 server.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Threat Category	Content Update that Provides These Signatures	Description
		Newly-released protections in this category are rare.
dns	Antivirus	Detects DNS requests to connect to malicious domains. dns and dns-wildfire signatures detect the same malicious domains; however, dns signatures are included in the daily Antivirus content update and dns-wildfire signatures are included in the WildFire updates that release protections every 5 minutes.
dns-security	Antivirus	Detects DNS requests to connect to malicious domains. dns-security includes signatures from dns and dns-wildfire in addition to the unique signatures generated by the DNS Security service.
dns-wildfire	Wildfire or WildFire Private	Detects DNS requests to connect to malicious domains. dns and dns-wildfire signatures detect the same malicious domains; however, dns signatures are included in the daily Antivirus content update and dns-wildfire signatures are included in the WildFire updates that release protections every 5 minutes.
downloader	Applications and Threats	(Also known as droppers, stagers, or loaders) Detects programs that use an internet connection to connect to a remote server to download and execute malware on the compromised system. The most common use case is for a downloader to be deployed as the culmination of *stage one* of a cyber attack, where the downloader's fetched payload execution is considered *second stage*. Shell scripts (Bash, PowerShell, etc.), trojans, and malicious lure documents (also known as maldocs) such as PDFs and Word files are common downloader types.
fraud	Applications and Threats	(Including form-jacking, phishing, and scams) Detects access to compromised websites that have been determined to be injected with malicious JavaScript code to collect sensitive user information. (for example, Name, address, email, credit card number, CVV, expiration date) from payment forms that are captured on the checkout pages of e-commerce websites.
hacktool	Applications and Threats	Detects traffic generated by software tools that are used by malicious actors to conduct reconnaissance, attack or gain access to vulnerable systems, exfiltrate data, or create a command and control channel to

https://t.me/learningnets

Threat Category	Content Update that Provides These Signatures	Description
		surreptitiously control a computer system without authorization. These programs are strongly associated with malware and cyber attacks. Hacking tools might be deployed in a benign manner when used in Red and Blue Team operations, penetration tests, and R&D. The use or possession of these tools may be illegal in some countries, regardless of intent.
keylogger	Applications and Threats	Detects programs that allow attackers to secretly track user activity, by logging keystrokes and capturing screenshots. Keyloggers use various C2 methods to periodically sends logs and reports to a predefined e-mail address or a C2 server. Through keylogger surveillance, an attacker could retrieve credentials that would enable network access.
networm	Applications and Threats	Detects a program that self-replicates and spreads from system to system. Net-worms might use shared resources or leverage security failures to access target systems.
phishing-kit	Applications and Threats	Detects when a user attempts to connect to a phishing kit landing page (likely after receiving an email with a link to the malicious site). A phishing website tricks users into submitting credentials that an attacker can steal to gain access to the network. *In addition to blocking access to phishing kit landing pages, enable Multi-Factor Authentication and Credential Phishing Prevention to prevent phishing attacks at all stages.*
post-exploitation	Applications and Threats	Detects activity that indicates the post-exploitation phase of an attack, where an attacker attempts to assess the value of a compromised system. This might include evaluating the sensitivity of the data stored on the system, and the system's usefulness in further compromising the network.
webshell	Applications and Threats	Detects web shells and web shell traffic, including implant detection and command and control interaction. Web shells must first be implanted by a malicious actor onto the compromised host, most often targeting a

https://t.me/learningnets

Threat Category	Content Update that Provides These Signatures	Description
		web server or framework. Subsequent communication with the web shell file frequently enables a malicious actor to establish a foothold in the system, conduct service and network enumeration, data exfiltration, and remote code execution in the context of the web server user. The most common web shell types are PHP, .NET, and Perl markup scripts. Attackers can also use web shell-infected web servers (the web servers can be both internet-facing or internal systems) to target other internal systems.
spyware	Applications and Threats	Detect outbound C2 communication. These signatures are either auto-generated or are manually created by Palo Alto Networks researchers. *Spyware and autogen signatures both detect outbound C2 communication; however, autogen signatures are payload-based and can uniquely detect C2 communications with C2 hosts that are unknown or change rapidly.*
Vulnerability Signatures		
brute force	Applications and Threats	A brute-force signature detects multiple occurrences of a condition in a particular time frame. While the activity in isolation might be benign, the brute-force signature indicates that the frequency and rate at which the activity occurred is suspect. For example, a single FTP login failure does not indicate malicious activity. However, many failed FTP logins in a short period likely indicate an attacker attempting password combinations to access an FTP server. You can tune the action and trigger conditions for brute force signatures.
code execution	Applications and Threats	Detects a code execution vulnerability that an attacker can leverage to run code on a system with the privileges of the logged-in user.
code-obfuscation	Applications and Threats	Detects code that has been transformed to conceal certain data while retaining its function. Obfuscated code is difficult or impossible to read, so it's not apparent what commands the code is executing or with which programs its designed to interact. Most

https://t.me/learningnets

Threat Category	Content Update that Provides These Signatures	Description
		commonly, malicious actors obfuscate code to conceal malware. More rarely, legitimate developers might obfuscate code to protect privacy, intellectual property, or to improve user experience. For example, certain types of obfuscation (like minification) reduce file size, which decreases website load times and bandwidth usage.
dos	Applications and Threats	Detects a denial-of-service (DoS) attack, where an attacker attempts to render a targeted system unavailable, temporarily disrupting the system and dependent applications and services. To perform a DoS attack, an attacker might flood a targeted system with traffic or send information that causes it to fail. DoS attacks deprive legitimate users (like employees, members, and account holders) of the service or resource to which they expect access.
exploit-kit	Applications and Threats	Detects an exploit kit landing page. Exploit kit landing pages often contain several exploits that target one or many common vulnerabilities and exposures (CVEs), for multiple browsers and plugins. Because the targeted CVEs change quickly, exploit-kit signatures trigger based on the exploit kit landing page, and not the CVEs. When a user visits a website with an exploit kit, the exploit kit scans for the targeted CVEs and attempts to silently deliver a malicious payload to the victim's computer.
info-leak	Applications and Threats	Detects a software vulnerability that an attacker could exploit to steal sensitive or proprietary information. Often, an info-leak might exist because comprehensive checks do not exist to guard the data, and attackers can exploit info-leaks by sending crafted requests.
insecure-credentials	Applications and Threats	Detects the use of weak, compromised, and manufacturer default passwords for software, network appliances, and IoT devices.
overflow	Applications and Threats	Detects an overflow vulnerability, where a lack of proper checks on requests could be exploited by an attacker. A successful attack could lead to remote code execution with the privileges of the application, server or operating system.

https://t.me/learningnets

Threat Category	Content Update that Provides These Signatures	Description
phishing	Applications and Threats	Detects when a user attempts to connect to a phishing kit landing page (likely after receiving an email with a link to the malicious site). A phishing website tricks users into submitting credentials that an attacker can steal to gain access to the network. *In addition to blocking access to phishing kit landing pages, enable Multi-Factor Authentication and Credential Phishing Prevention to prevent phishing attacks at all stages.*
protocol-anomaly	Applications and Threats	Detects protocol anomalies, where a protocol behavior deviates from standard and compliant usage. For example, a malformed packet, poorly-written application, or an application running on a non-standard port would all be considered protocol anomalies, and could be used as evasion tools. It is a best practice to block protocol anomalies of any severity.
sql-injection	Applications and Threats	Detects a common hacking technique where an attacker inserts SQL queries into an application's requests, in order to read from or modify a database. This type of technique is often used on websites that do not comprehensively sanitize user input.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Create Threat Exceptions

Palo Alto Networks defines a recommended default action (such as block or alert) for threat signatures. You can use a threat ID to exclude a threat signature from enforcement or modify the action the firewall enforces for that threat signature. For example, you can modify the action for threat signatures that are triggering false positives on your network.

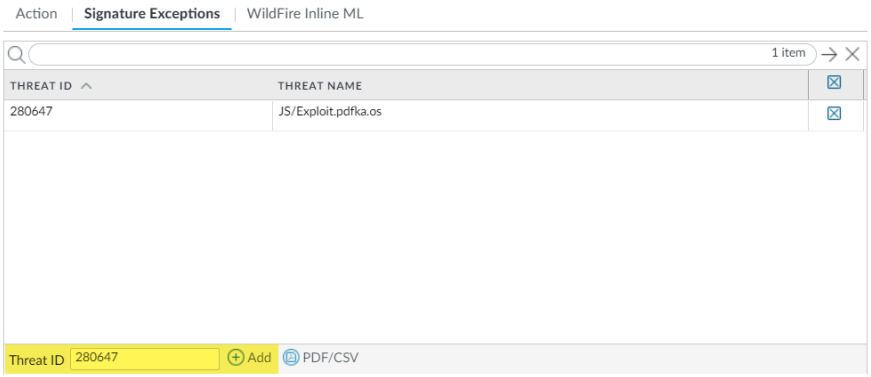
Configure threat exceptions for antivirus, vulnerability, spyware, and DNS signatures to change firewall enforcement for a threat. However, before you begin, make sure the firewall is detecting and enforcing threats based on the default signature settings:

- Get the latest Antivirus, Threats and Applications, and WildFire signature updates.
- Set Up Antivirus, Anti-Spyware, and Vulnerability Protection and apply these security profiles to your security policy.

STEP 1 | Exclude antivirus signatures from enforcement.

*While you can use an Antivirus profile to exclude antivirus signatures from enforcement, you cannot change the action the firewall enforces for a specific antivirus signature. However, you can define the action for the firewall to enforce for viruses found in different types of traffic by editing the Decoders (**Objects** > **Security Profiles** > **Antivirus** > **<antivirus-profile>** > **Antivirus**).*

1. Select **Objects** > **Security Profiles** > **Antivirus**.
2. **Add** or modify an existing Antivirus profile from which you want to exclude a threat signature and select **Signature Exceptions**.
3. **Add** the **Threat ID** for the threat signature you want to exclude from enforcement.



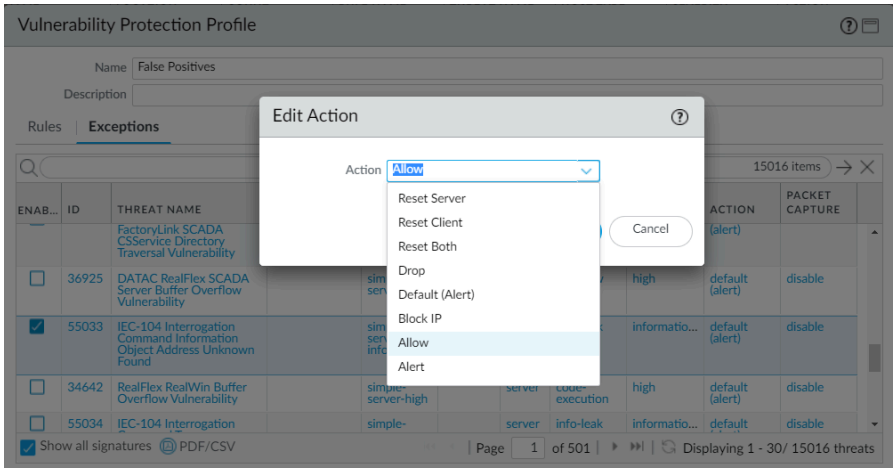
4. Click **OK** to save the Antivirus profile.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 2 | Modify enforcement for vulnerability and spyware signatures (except DNS signatures; skip to the next option to modify enforcement for DNS signatures, which are a type of spyware signature).

1. Select **Objects** > **Security Profiles** > **Anti-Spyware** or **Objects** > **Security Profiles** > **Vulnerability Protection**.
2. **Add** or modify an existing Anti-Spyware or Vulnerability Protection profile from which you want to exclude the threat signature and then select either **Signature Exceptions** for Anti-Spyware Protection profiles or **Exceptions** for Vulnerability Protection profiles.
3. **Show all signatures** and then filter to select the signature for which you want to modify enforcement rules.
4. Check the box under the **Enable** column for the signature whose enforcement you want to modify.
5. Select the **Action** you want the firewall to enforce for this threat signature.



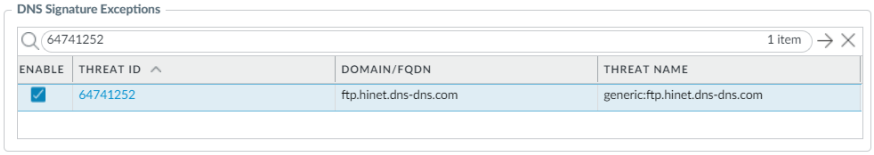
For signatures that you want to exclude from enforcement because they trigger false positives, set the **Action** to **Allow**.

6. Click **OK** to save your new or modified Anti-Spyware or Vulnerability Protection profile.

STEP 3 | Modify enforcement for DNS signatures.

By default, the DNS lookups to malicious hostnames that DNS signatures are detect are sinkholed.

1. Select **Objects** > **Security Profiles** > **Anti-Spyware**.
2. **Add** or modify the Anti-Spyware profile from which you want to exclude the threat signature, and select **DNS Exceptions**.
3. Search for the DNS Threat ID for the DNS signature that you want to exclude from enforcement and select the box of the applicable signature:



4. Click **OK** to save your new or modified Anti-Spyware profile.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Custom Signatures

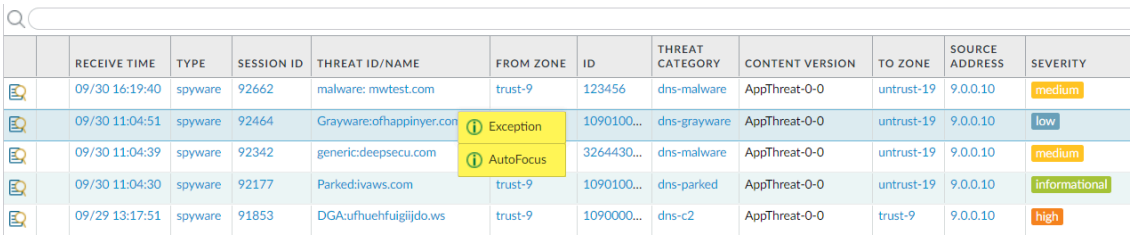
You can create custom threat signatures to detect and block specific traffic. When the firewall is managed by a Panorama management server, the ThreatID is mapped to the corresponding custom threat on the firewall to enable the firewall to generate a threat log populated with the configured custom ThreatID. Learn more by visiting our guide to Custom Application and Threat Signatures.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Monitor and Get Threat Reports

Features of Threat Vault and AutoFocus are integrated into the firewall to provide visibility into the nature of the threats the firewall detects and to give a more complete picture of how an artifact fits into your organization's network traffic (an artifact is property, activity, or behavior associated with a file, email link, or session). You can get immediate, contextual information about a threat or to seamlessly shift your threat investigation from the firewall to the Threat Vault and AutoFocus.



Additionally, you can use Threat Signature Categories—which classify types of threat events—to narrow your view into a certain type of threat activity or to build custom reports.

- Monitor Activity and Create Custom Reports Based on Threat Categories
- Learn More About Threat Signatures
- AutoFocus Threat Intelligence for Network Traffic

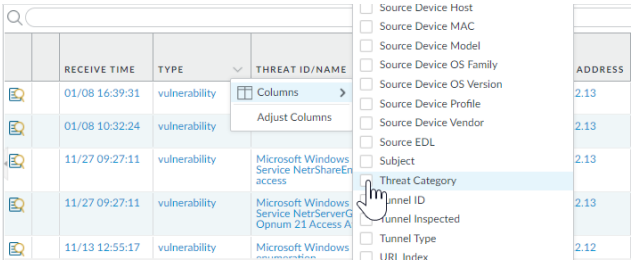
Monitor Activity and Create Custom Reports Based on Threat Categories

Threat categories classify different types of threat signatures to help you understand and draw connections between events threat signatures detect. Threat categories are subsets of the more broad threat signature types: spyware, vulnerability, antivirus, and DNS signatures. Threat log entries display the **Threat Category** for each recorded event.

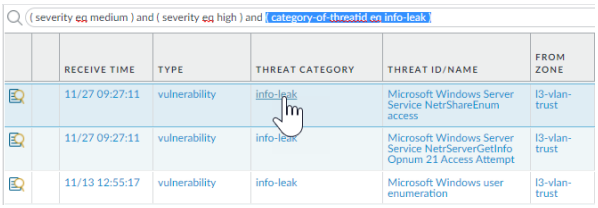
©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

- Filter Threat logs by threat category.
  1. Select **Monitor** > **Logs** > **Threat**.
  2. Add the Threat Category column so you can view the Threat Category for each log entry:



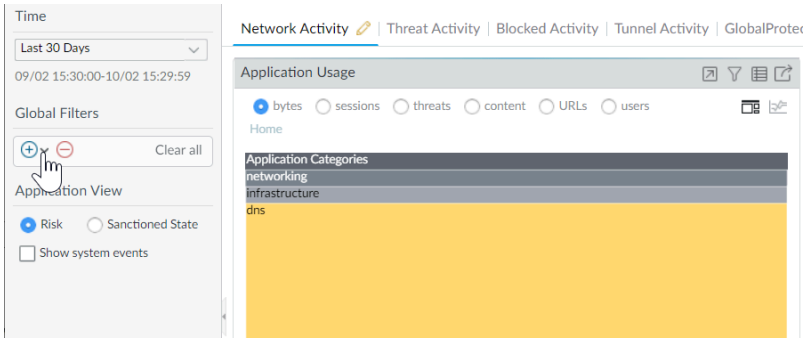
  3. To filter based on Threat Category:
     - Use the log query builder to add a filter with the **Attribute** Threat Category and in the **Value** field, enter a Threat Category.
     - Select the Threat Category of any log entry to add that category to the filter:



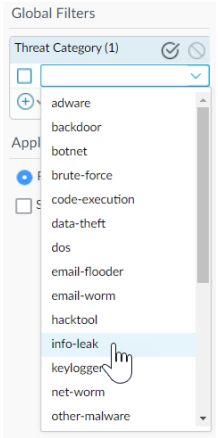
 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

- Filter ACC activity by threat category.
  1. Select **ACC** and add Threat Category as a global filter:



  2. Select the Threat Category to filter all ACC tabs.



- Create custom reports based on threat categories to receive information about specific types of threats that the firewall has detected.
  1. Select **Monitor** > **Manage Custom** reports to add a new custom report or modify an existingone.
  2. Choose the **Database** to use as the source for the custom report—in this case, select **Threat** from either of the two types of database sources, summary databases and Detailed logs. Summary database data is condensed to allow a faster response time when generating reports. Detailed logs take longer to generate but provide an itemized and complete set of data for each log entry.
  3. In the Query Builder, add a report filter with the Attribute **Threat Category** and in the Value field, select a threat category on which to base your report.
  4. To test the new report settings, click **Run Now**.
  5. Click **OK** to save the report.

Learn More About Threat Signatures

Firewall Threat logs record all threats the firewall detects based on threat signatures (Set Up Antivirus, Anti-Spyware, and Vulnerability Protection) and the ACC displays an overview of the top threats on your network. Each event the firewall records includes an ID that identifies the associated threat signature.

https://t.me/learningnets

You can use the threat ID found with a Threat log or ACC entry to:

- Easily check if a threat signature is configured as an exception to your security policy (Create Threat Exceptions).
- Find the latest Threat Vault information about a specific threat. Because the Threat Vault is integrated with the firewall, you can view threat details directly in the firewall context or launch a Threat Vault search in a new browser window for a threat the firewall logged.

If a signature has been disabled, the signature UTID might be reused for a new signature.

Review the content update release notes for notifications regarding new and disabled signatures. Signatures might disabled in cases where: the activity the signature detects has fallen out of use by attackers, the signature generated significant false positives, or the signature was consolidated with other like signatures into a single signature (signature optimization).

STEP 1 | Confirm the firewall is connected to the Threat Vault.

Select **Device** > **Setup** > **Management** and edit the **Logging and Reporting** setting to **Enable Threat Vault Access**. Threat vault access is enabled by default.

STEP 2 | Find the threat ID for threats the firewall detects.

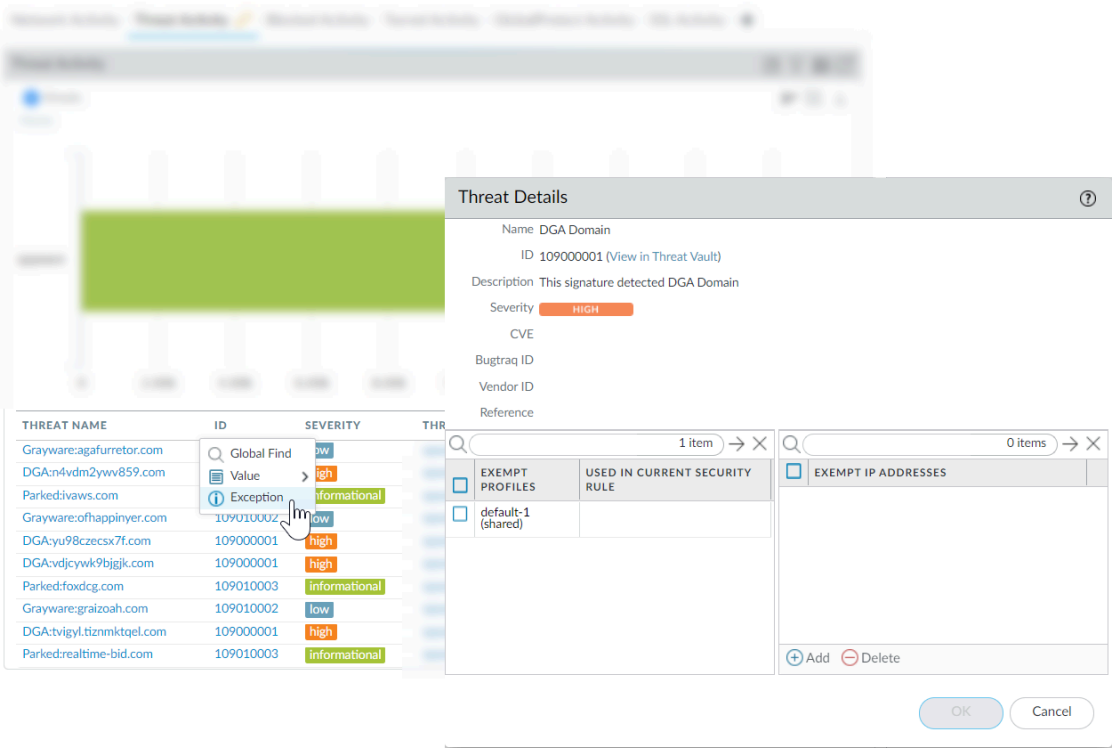
- To see each threat event the firewall detects based on threat signatures, select **Monitor** > **Logs** > **Threat**. You can find the ID for a threat entry listed in the ID column, or select the log entry to view log details, including the Threat ID.
- To see an overview of top threats on the network, select **ACC** > **Threat Activity** and take a look at the Threat Activity widget. The ID column displays the threat ID for each threat displayed.
- To see details for threats that you can configure as threat exceptions (meaning, the firewall enforces the threat differently than the default action defined for the threat signature), select **Objects** > **Security Profiles** > **Anti-Spyware/Vulnerability Protection**. **Add** or modify a profile and click the **Exceptions** tab to view configured exceptions. If no exceptions are configured, you can filter for threat signatures or select **Show all signatures**.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 3 | Hover over a **Threat Name** or the threat **ID** to open the drop-down, and click **Exception** to review both the threat details and how the firewall is configured to enforce the threat.

For example, find out more about a top threat charted on the ACC:



STEP 4 | Review the latest **Threat Details** for the threat and launch a Threat Vault search based on the threat ID.

- Threat details displayed include the latest Threat Vault information for the threat, resources you can use to learn more about the threat, and CVEs associated with the threat.
- Select **View in Threat Vault** to open a Threat Vault search in a new window and look up the latest information the Palo Alto Networks threat database has for this threat signature.

STEP 5 | Check if a threat signature is configured as an exception to your security policy.

- If the **Used in current security rule** column is clear, the firewall is enforcing the threat based on the recommended default signature action (for example, block or alert).
- A checkmark anywhere in the **Used in current security rule** column indicates that a security policy rule is configured to enforce a non-default action for the threat (for example, allow), based on the associated **Exempt Profiles** settings.

*The **Used in security rule column** does not indicate if the Security policy rule is enabled, only if the Security policy rule is configured with the threat exception. Select **Policies** > **Security** to check if an indicated security policy rule is enabled.*

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 6 | **Add** an IP address on which to filter the threat exception or view existing **Exempt IP Addresses**.

Configure an exempt IP address to enforce a threat exception only when the associated session has either a matching source or destination IP address; for all other sessions, the threat is enforced based on the default signature action.

AutoFocus Threat Intelligence for Network Traffic

With a valid AutoFocus subscription, you can compare the activity on your network with the latest threat data available on the AutoFocus portal. Connecting your firewall and AutoFocus unlocks the following features:

- View an AutoFocus intelligence summary for session artifacts recorded in the firewall logs.
- Open an AutoFocus search for log artifacts from the firewall.

The AutoFocus intelligence summary reveals the prevalence of an artifact on your network and on a global scale. The WildFire verdicts and AutoFocus tags listed for the artifact indicate whether the artifact poses a security risk.

- AutoFocus Intelligence Summary
- Enable AutoFocus Threat Intelligence
- View and Act on AutoFocus Intelligence Summary Data

You can also enforce policy based on AutoFocus findings:

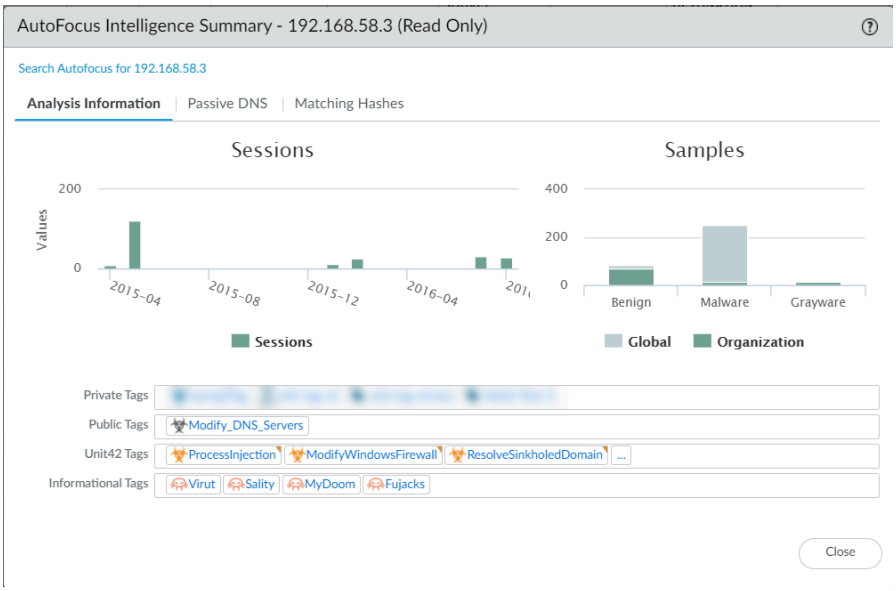
- Export AutoFocus artifacts (IP addresses, URLs, and domains) and use them in an external dynamic list.
- Use an AutoFocus miner as an external dynamic list source.

AutoFocus Intelligence Summary

The AutoFocus Intelligence Summary offers a centralized view of information about an artifact that AutoFocus has extracted from threat intelligence gathered from other AutoFocus users, WildFire, the PAN-DB URL filtering database, Unit 42, and open-source intelligence.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets



AutoFocus Intelligence Summary	
Analysis Information	The Analysis Information tab displays the following information: • Sessions—The number of sessions logged in your firewall(s) in which the firewall detected samples associated with the artifact. • Samples—A comparison of organization and global samples associated with the artifact and grouped by WildFire verdict (benign, malware, or grayware). *Global* refers to samples from all WildFire submissions, while *organization* refers only to samples submitted to WildFire by your organization. • Matching Tags—The AutoFocus tags matched to the artifact. AutoFocus Tags indicate whether an artifact is linked to malware or targeted attacks.
Passive DNS	The Passive DNS tab displays passive DNS history that includes the artifact. This passive DNS history is based on global DNS intelligence in AutoFocus; it is not limited to the DNS activity in your network. Passive DNS history consists of: • The domain request • The DNS request type • The IP address or domain to which the DNS request resolved (private IP addresses are not displayed) • The number of times the request was made • The date and time the request was first seen and last seen
Matching Hashes	The Matching Hashes tab displays the 5 most recently detected matching samples. Sample information includes:

https://t.me/learningnets

AutoFocus Intelligence Summary	
	• The SHA256 hash of the sample • The sample file type • The date and time that WildFire analyzed a sample and assigned a WildFire verdict to it • The WildFire verdict for the sample • The date and time that WildFire updated the WildFire verdict for the sample (if applicable)

Enable AutoFocus Threat Intelligence

Activate the AutoFocus license, and enable the firewall to communicate with AutoFocus. Once you're set up, you'll be able to view the AutoFocus Intelligence Summary for a log or ACC artifact, to assess its pervasiveness in your network and any associated threats.

STEP 1 | Verify that the AutoFocus license is activated on the firewall.

1. Select **Device** > **Licenses** to verify that the AutoFocus Device License is installed and valid (check the expiration date).
2. If the firewall doesn't show the license, Activate Subscription Licenses.

STEP 2 | Connect the firewall to AutoFocus.

1. Select **Device** > **Setup** > **Management** and edit the AutoFocus settings.
2. Enter the **AutoFocus URL**:

   `https://autofocus.paloaltonetworks.com:10443`
3. Use the **Query Timeout** field to set the duration of time for the firewall to attempt to query AutoFocus for threat intelligence data. If the AutoFocus portal does not respond before the end of the specified period, the firewall closes the connection.

   As a best practice, set the query timeout to the default value of 15 seconds. AutoFocus queries are optimized to complete within this duration.
4. Select **Enabled** to allow the firewall to connect to AutoFocus.
5. Click **OK**.
6. **Commit** your changes to retain the AutoFocus settings upon reboot.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 3 | Connect AutoFocus to the firewall.

1. Log in to the AutoFocus portal: https://autofocus.paloaltonetworks.com
2. Select **Settings**.
3. **Add new** remote systems.
4. Enter a descriptive **Name** to identify the firewall.
5. Select **PanOS** as the System Type.
6. Enter the firewall IP **Address**.
7. Click **Save changes** to add the remote system.
8. Click **Save changes** again on the Settings page to ensure the firewall is successfully added.

STEP 4 | Test the connection between the firewall and AutoFocus.

1. On the firewall, select **Monitor** > **Logs** > **Traffic**.
2. Verify that you can Assess Firewall Artifacts with AutoFocus.

View and Act on AutoFocus Intelligence Summary Data

Interact with the AutoFocus Intelligence Summary to display more information about an artifact or extend your artifact research to AutoFocus. AutoFocus tags reveal if the artifact is associated with certain types of malware or malicious behavior.

STEP 1 | Confirm that the firewall is connected to AutoFocus.

Enable AutoFocus Threat Intelligence on the firewall (active AutoFocus subscription required).

STEP 2 | Find artifacts to investigate.

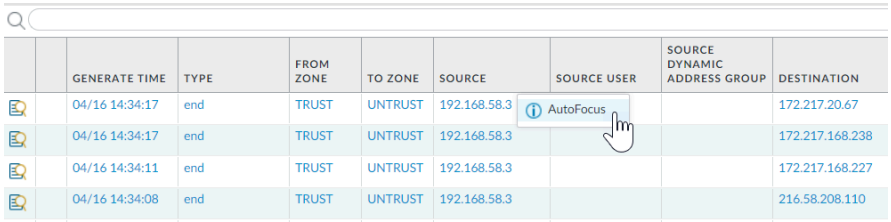
You can view an AutoFocus Intelligence Summary for artifacts when you:

- View Logs (Traffic, Threat, URL Filtering, WildFire Submissions, Data Filtering, and Unified logs only).
- View External Dynamic List Entries.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 3 | Hover over an artifact to open the drop-down, and click **AutoFocus**.



The AutoFocus Intelligence Summary is only available for the following types of artifacts:

IP address

URL

Domain

User agent

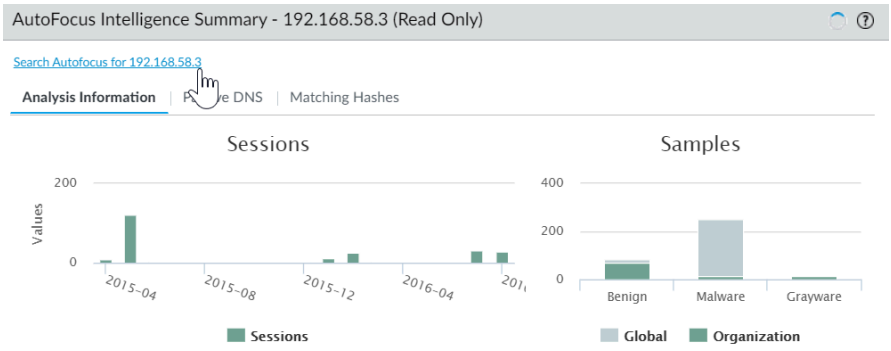
Threat name (only for threats of the subtypes virus and wildfire-virus)

Filename

SHA-256 hash

STEP 4 | Launch an AutoFocus search for the artifact for which you opened the AutoFocus Intelligence Summary.

Click the **Search AutoFocus for...** link at the top of the AutoFocus Intelligence Summary window. The search results include all samples associated with the artifact. Toggle between the **My Samples** and **All Samples** tabs and compare the number of samples to determine the pervasiveness of the artifact in your organization.



STEP 5 | Launch an AutoFocus search for other artifacts in the AutoFocus Intelligence Summary.

Click on the following artifacts to determine their pervasiveness in your organization:

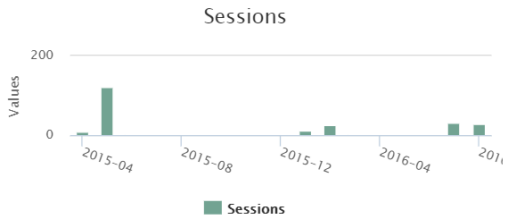
- WildFire verdicts in the Analysis Information tab
- URLs and IP addresses in the Passive DNS tab
- The SHA256 hashes in the Matching Hashes tab

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

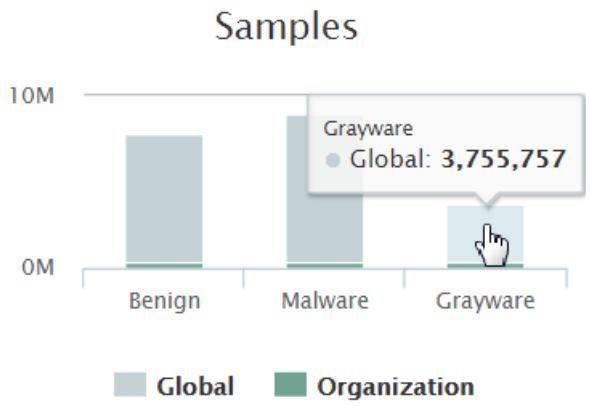
STEP 6 | View the number of sessions associated with the artifact in your organization per month.

Hover over the session bars.



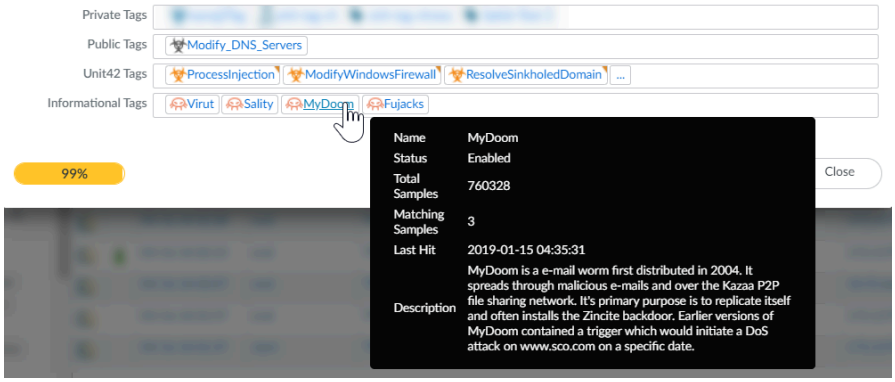
STEP 7 | View the number of samples associated with the artifact by scope and WildFire verdict.

Hover over the samples bars.



STEP 8 | View more details about matching AutoFocus. tags.

Hover over a matching tag to view the tag description and other tag details.



STEP 9 | View other samples associated with a matching tag.

Click a matching tag to launch an AutoFocus search for that tag. The search results include all samples matched to the tag.

Unit 42 tags identify threats and campaigns that pose a direct security risk. Click on a Unit 42 matching tag to see how many samples in your network are associated with the threat the tag identifies.

https://t.me/learningnets

STEP 10 | Find more matching tags for an artifact.

Click the ellipsis (...) to launch an AutoFocus search for the artifact. The Tags column in the search results displays more matching tags for the artifact, which give you an idea of other malware, malicious behavior, threat actors, exploits, or campaigns where the artifact is commonly detected.



 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Share Threat Intelligence with Palo Alto Networks

Telemetry is the process of collecting and transmitting data for analysis. When you enable telemetry on the firewall, the firewall periodically collects and sends information that includes applications, threats, and device health to Palo Alto Networks. Sharing threat intelligence provides the following benefits:

- Enhanced vulnerability and spyware signatures delivered to you and other customers worldwide. For example, when a threat event triggers vulnerability or spyware signatures, the firewall shares the URLs associated with the threat with the Palo Alto Networks threat research team, so they can properly classify the URLs as malicious.
- Rapid testing and evaluation of experimental threat signatures with no impact to your network, so that critical threat prevention signatures can be released to all Palo Alto Networks customers faster.
- Improved accuracy and malware detection abilities within PAN-DB URL filtering, DNS-based command-and-control (C2) signatures, and WildFire.

Palo Alto Networks uses the threat intelligence extracted from telemetry to deliver these benefits to you and other Palo Alto Networks users. All Palo Alto Networks users benefit from the telemetry data shared by each user, making telemetry a community-driven approach to threat prevention. Palo Alto Networks does not share your telemetry data with other customers or third-party organizations.

To read more about telemetry, including its benefits, usages, and configuration, see Device Telemetry.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Threat Prevention Resources

For more information on threat prevention best practices, refer to the following sources:

- Creating Custom Threat Signatures
- Best Practices for Securing Your Network from Layer 4 and Layer 7 Evasions
- URL Filtering Best Practices
- Zero Trust Best Practices
- DoS and Zone Protection Best Practices

To view a list of threats and applications that Palo Alto Networks products can identify, use the following links:

- Applipedia—Provides details on the applications that Palo Alto Networks can identify.
- Threat Vault—Lists threats that Palo Alto Networks products can identify. You can search by Vulnerability, Spyware, or Virus. Click the Details icon next to the ID number for more information about a threat.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Decryption

Palo Alto Networks firewalls can decrypt and inspect traffic to provide visibility into threats and to control protocols, certificate verification, and failure handling. Decryption can enforce policies on encrypted traffic so that the firewall handles encrypted traffic according to your configured security settings. Decrypt traffic to prevent malicious encrypted content from entering your network and sensitive content from leaving your network concealed as encrypted traffic. Enabling decryption can include preparing the keys and certificates required for decryption, creating decryption profiles and policies, and configuring decryption port mirroring.

- Decryption Overview
- Decryption Concepts
- Prepare to Deploy Decryption
- Define Traffic to Decrypt
- Configure SSL Forward Proxy
- Configure SSL Inbound Inspection
- Configure SSH Proxy
- Configure Server Certificate Verification for Undecrypted Traffic
- Decryption Exclusions
- Block Private Key Export
- Enable Users to Opt Out of SSL Decryption
- Temporarily Disable SSL Decryption
- Configure Decryption Port Mirroring
- Verify Decryption
- Troubleshoot and Monitor Decryption
- Activate Free Licenses for Decryption Features

https://t.me/learningnets

Decryption Overview

The Secure Sockets Layer (SSL) and Secure Shell (SSH) encryption protocols secure traffic between two entities, such as a web server and a client. SSL and SSH encapsulate traffic, encrypting data so that it is meaningless to entities other than the client and server with the certificates to affirm trust between the devices and the keys to decode the data. Decrypt SSL and SSH traffic to:

- Prevent malware concealed as encrypted traffic from being introduced into your network. For example, an attacker compromises a website that uses SSL encryption. Employees visit that website and unknowingly download an exploit or malware. The malware then uses the infected employee endpoint to move laterally through the network and compromise other systems.
- Prevent sensitive information from moving outside the network.
- Ensure the appropriate applications are running on a secure network.
- Selectively decrypt traffic; for example, create a Decryption policy and profile to exclude traffic for financial or healthcare sites from decryption.

Palo Alto Networks firewall decryption is policy-based, and can decrypt, inspect, and control inbound and outbound SSL and SSH connections. A Decryption policy enables you to specify traffic to decrypt by destination, source, service, or URL category, and to block, restrict, or forward the specified traffic according to the security settings in the associated Decryption profile. A Decryption profile controls SSL protocols, certificate verification, and failure checks to prevent traffic that uses weak algorithms or unsupported modes from accessing the network. The firewall uses certificates and keys to decrypt traffic to plaintext, and then enforces App-ID and security settings on the plaintext traffic, including Decryption, Antivirus, Vulnerability, Anti-Spyware, URL Filtering, WildFire, and File-Blocking profiles. After decrypting and inspecting traffic, the firewall re-encrypts the plaintext traffic as it exits the firewall to ensure privacy and security.

The firewall provides three types of Decryption policy rules: SSL Forward Proxy to control outbound SSL traffic, SSL Inbound Inspection to control inbound SSL traffic, and SSH Proxy to control tunneled SSH traffic. You can attach a Decryption profile to a policy rule to apply granular access settings to traffic, such as checks for server certificates, unsupported modes, and failures.

SSL decryption (both forward proxy and inbound inspection) requires certificates to establish the firewall as a trusted third party, and to establish trust between a client and a server to secure an SSL/TLS connection. You can also use certificates when excluding servers from SSL decryption for technical reasons (the site breaks decryption for reasons such as certificate pinning, unsupported ciphers, or mutual authentication). SSH decryption does not require certificates.

Use the Decryption Best Practices Checklist to plan, implement, and maintain your decryption deployment.

You can integrate a hardware security module (HSM) with a firewall to enable enhanced security for the private keys used in SSL forward proxy and SSL inbound inspection decryption. To learn more about storing and generating keys using an HSM and integrating an HSM with your firewall, see Secure Keys with a Hardware Security Module.

You can also use Decryption Mirroring to forward decrypted traffic as plaintext to a third party solution for additional analysis and archiving.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

If you enable Decryption mirroring, be aware of local laws and regulations about what traffic you can mirror and where and how you can store the traffic, because all mirrored traffic, including sensitive information, is forwarded in cleartext.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Decryption Concepts

Review the following topics to learn more about decryption features and support:

- Keys and Certificates for Decryption Policies
- SSL Forward Proxy
- SSL Forward Proxy Decryption Profile
- SSL Inbound Inspection
- SSL Inbound Inspection Decryption Profile
- SSL Protocol Settings Decryption Profile
- SSH Proxy
- SSH Proxy Decryption Profile
- SSL Profile for No Decryption
- SSL Decryption for Elliptical Curve Cryptography (ECC) Certificates
- Perfect Forward Secrecy (PFS) Support for SSL Decryption
- SSL Decryption and Subject Alternative Names (SANs)
- TLSv1.3 Decryption
- High Availability Support for Decrypted Sessions
- Decryption Mirroring

Keys and Certificates for Decryption Policies

Keys are strings of numbers typically generated using a mathematical operation involving random numbers and large primes. Keys transform strings—such as passwords and shared secrets—from unencrypted plaintext to encrypted ciphertext and from encrypted ciphertext to unencrypted plaintext. Keys can be symmetric (the same key is used to encrypt and decrypt) or asymmetric (one key is used for encryption and a mathematically related key is used for decryption). Any system can generate a key.

X.509 certificates establish trust between a client and a server to establish an SSL connection. A client attempting to authenticate a server (or a server authenticating a client) knows the structure of the X.509 certificate and therefore knows how to extract identifying information about the server from fields within the certificate, such as the FQDN or IP address (called a *common name* or *CN* within the certificate) or the name of the organization, department, or user to which the certificate was issued. A certificate authority (CA) must issue all certificates. After the CA verifies a client or server, the CA issues the certificate and signs it with a private key.

*If you have two CAs (**Device** > **Certificate Management** > **Device Certificates**) with the same subject and key, and one CA expires, delete (custom) or disable (predefined) the expired CA. If you do not delete or disable an expired CA, the firewall can build a chain to the expired CA if it is enabled in the trusted chain resulting in a Block page.*

When you apply a decryption policy to traffic, a session between the client and the server is established only if the firewall trusts the CA that signed the server certificate. In order to establish

https://t.me/learningnets

trust, the firewall must have the server root CA certificate in its certificate trust list (CTL) and use the public key contained in that root CA certificate to verify the signature. The firewall then presents a copy of the server certificate signed by the Forward Trust certificate for the client to authenticate. You can also configure the firewall to use an enterprise CA as a forward trust certificate for SSL Forward Proxy. If the firewall does not have the server root CA certificate in its CTL, the firewall will present a copy of the server certificate signed by the Forward Untrust certificate to the client. The Forward Untrust certificate ensures that clients are prompted with a certificate warning when attempting to access sites hosted by a server with untrusted certificates.

For detailed information on certificates, see Certificate Management.

*To control the trusted CAs that your firewall trusts, use the **Device** > **Certificate Management** > **Certificates** > **Default Trusted Certificate Authorities** tab on the firewall web interface.*

The following table describes the different certificates Palo Alto Networks firewalls use for decryption.

Certificates Used With Decryption	Description
Forward Trust (Used for SSL Forward Proxy decryption)	The certificate the firewall presents to clients during decryption if the site the client is attempting to connect to has a certificate signed by a CA that the firewall trusts. To configure a Forward Trust certificate on the firewall to present to clients when the server certificate is signed by a trusted CA, see Configure SSL Forward Proxy. By default, the firewall determines the key size to use for the client certificate based on the key size of the destination server. However, you can Configure the Key Size for SSL Proxy Server certificates. For added security, consider storing the private key associated with the Forward Trust certificate on a hardware security module (see Store Private Keys on an HSM). *Back up the private key associated with the firewall's Forward Trust CA certificate (not the firewall's master key) in a secure repository so that if an issue occurs with the firewall, you can still access the Forward Trust CA certificate. For added security, consider storing the private key associated with the Forward Trust certificate on a hardware security module (see Store Private Keys on an HSM).*
Forward Untrust (Used for SSL Forward Proxy decryption)	The certificate the firewall presents to clients during decryption if the site the client is attempting to connect to has a certificate that is signed by a CA that the firewall does not trust. To configure a Forward Untrust certificate on the firewall, see Configure SSL Forward Proxy.

https://t.me/learningnets

Certificates Used With Decryption	Description
SSL Inbound Inspection	The certificates of the servers on your network for which you want to perform SSL Inbound Inspection of traffic destined for those servers. Import the server certificates onto the firewall. *Beginning in PAN-OS 8.0, firewalls use the Elliptic-Curve Diffie-Hellman Ephemeral (ECDHE) algorithm to perform strict certificate checking. This means that if the firewall uses an intermediate certificate, you must reimport the certificate from your web server to the firewall after you upgrade to a PAN-OS 8.0 or later release and combine the server certificate with the intermediate certificate (install a chained certificate). Otherwise, SSL Inbound Inspection sessions that have an intermediate certificate in the chain will fail. To install a chained certificate:* 1. *Open each certificate (.cer) file in a plain-text editor such as Notepad.* 2. *Paste each certificate end-to-end with the Server Certificate at the top with each signer included below.* 3. *Save the file as a text (.txt) or certificate (.cer) file (the name of the file cannot contain blank spaces).* 4. *Import the combined (chained) certificate into the firewall.*

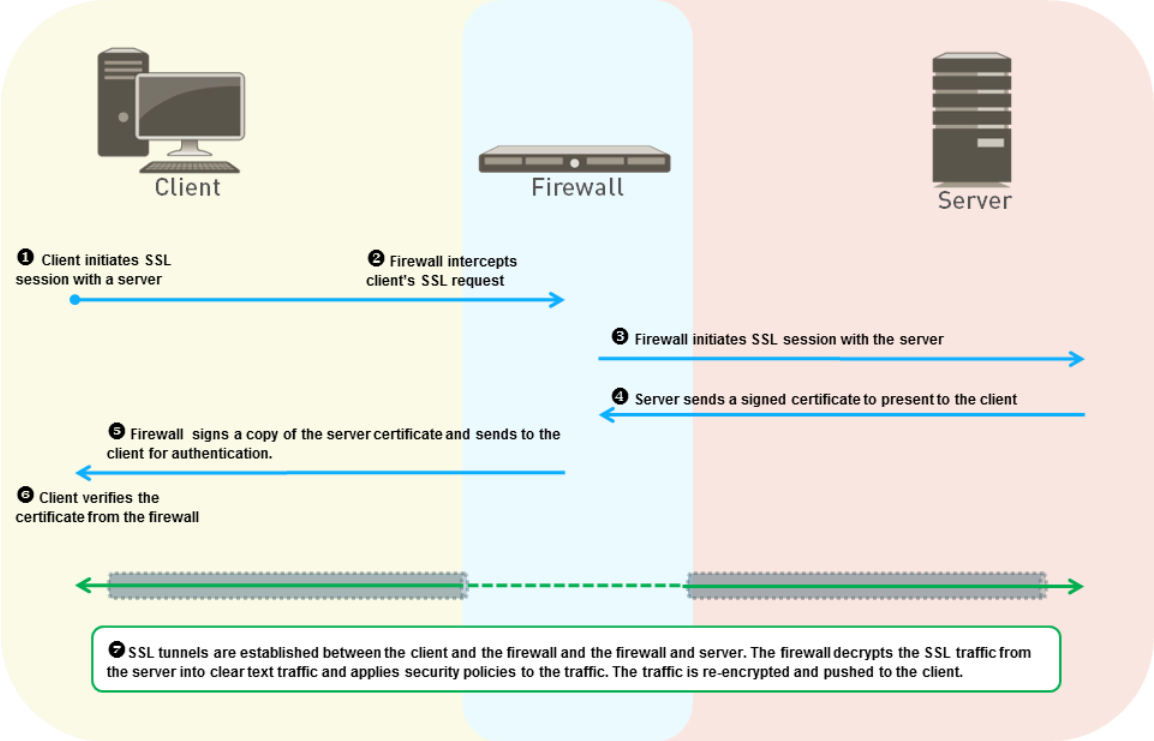
SSL Forward Proxy

When you configure the firewall to decrypt SSL traffic going to external sites, it functions as an SSL forward proxy. Use an SSL Forward Proxy decryption policy to decrypt and inspect SSL/TLS traffic from internal users to the web. SSL Forward Proxy decryption prevents malware concealed as SSL encrypted traffic from being introduced into your corporate network by decrypting the traffic so that the firewall can apply decryption profiles and security policies and profiles to the traffic.

In SSL Forward Proxy decryption, the firewall is a man-in-the-middle between the internal client and the external server. The firewall uses certificates to transparently represent the client to the server and to transparently represent the server to the client, so that the client believes it is communicating directly with the server (even though the client session is with the firewall), and server believes it is communicating directly with the client (even though the server session is also with the firewall). The firewall uses certificates to establish itself as a trusted third party (man-in-the-middle) for the client-server session (for details on certificates, see Keys and Certificates for Decryption Policies).

The following figure shows this process in detail. See Configure SSL Forward Proxy for details on configuring SSL Forward Proxy.

https://t.me/learningnets



1. The internal client on your network attempts to initiate a TLS session with an external server.
2. The firewall intercepts the client's SSL certificate request. For the client, the firewall acts as the external server, even though the secure session being established is with the firewall, not with the actual server.
3. The firewall then forwards the client's SSL certificate request to the server to initiate a separate session with the server. To the server, the firewall looks like the client, the server doesn't know there's a man-in-the-middle, and the server verifies the certificate.
4. The server sends the firewall a signed certificate intended for the client.
5. The firewall analyzes the server certificate. If the server certificate is signed by a CA that the firewall trusts and meets the policies and profiles you configure, the firewall generates an SSL Forward Trust copy of the server certificate and sends it to the client. If the server certificate is signed by a CA that the firewall does not trust, the firewall generates an SSL Forward Untrust copy of the server certificate and sends it to the client. The certificate copy the firewall generates and sends to the client contains extensions from the original server certificate and is called an *impersonation* certificate because it is not the server's actual certificate. If the firewall does not trust the server, the client sees a block page warning message that the site they're attempting to connect to is not trusted, and if you Enable Users to Opt Out of SSL Decryption, the client can choose to proceed or terminate the session.
6. The client verifies the firewall's impersonation certificate. The client then initiates a session key exchange with the server, which the firewall proxies in the same manner as it proxies the certificates. The firewall forwards the client key to the server, and makes an impersonation copy of the server key for the client, so that firewall remains an "invisible" proxy, the client and server believe their session is with each other, but there are still two separate sessions, one between the client and the firewall, and the other between the firewall and the server. Now all parties have the certificates and keys required and the firewall can decrypt the traffic.

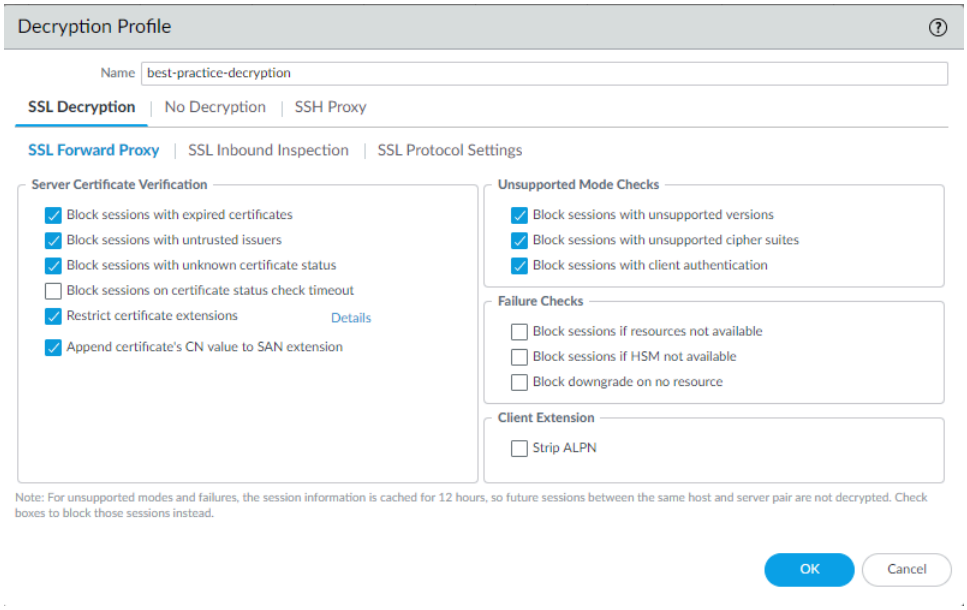
https://t.me/learningnets

7. All SSL session traffic between goes through the firewall transparently between the client and the server. The firewall decrypts the SSL traffic, applies security policies and profiles and decryption profiles to the traffic, re-encrypts the traffic, and then forwards it.

When you configure SSL Forward Proxy, the proxied traffic does not support DSCP code points or QoS.

SSL Forward Proxy Decryption Profile

The SSL Forward Proxy Decryption profile (**Objects** > **Decryption Profile** > **SSL Decryption** > **SSL Forward Proxy**) controls the server verification, session mode checks, and failure checks for outbound SSL/TLS traffic defined in Forward Proxy Decryption policies to which you attach the profile. The following figure shows the general best practice recommendations for Forward Proxy Decryption profile settings, but the settings you use also depend on your company's security compliance rules and local laws and regulations. There are also specific best practices for perimeter internet gateway decryption profiles and for data center decryption profiles.

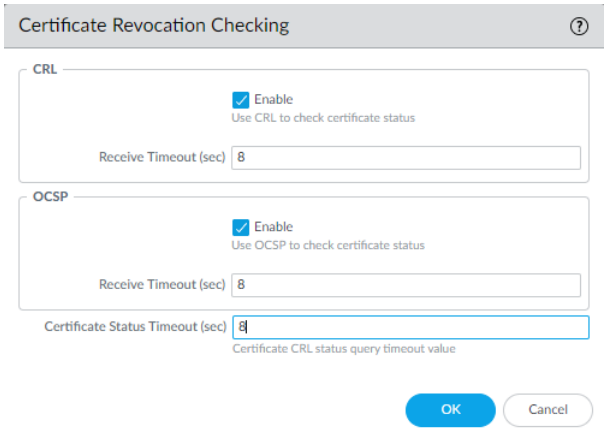


Server Certificate Verification:

- **Block sessions with expired certificates**—Always check this box to block sessions with servers that have expired certificates and prevent access to potentially insecure sites. If you don't check this box, users can connect with and transact with potentially malicious sites and see warning messages when they attempt to connect, but the connection is not prevented.
- **Block sessions with untrusted issuers**—Always check this box to block sessions with servers that have untrusted certificate issuers. An untrusted issuer may indicate a man-in-the-middle attack, a replay attack, or other attack.
- **Block sessions with unknown certificate status**—Blocks the SSL/TLS session when a the certificate revocation status of the server returns with the status "unknown". Because certificate status may be unknown for multiple reasons, for general decryption security, checking this box usually tightens security too much. However, in higher-security areas of the network such as the data center, checking this box makes sense.

https://t.me/learningnets

- **Block sessions on certificate status check timeout**—Whether to block sessions if the status check times out depends on your company's security compliance stance because it's a tradeoff between tighter security and a better user experience. Certificate status verification examines the Certificate Revocation List (CRL) on a revocation server or uses Online Certificate Status Protocol (OCSP) to find out if the issuing CA has revoked the certificate and the certificate should not be trusted. However, revocation servers can be slow to respond, which can cause the session to timeout and the firewall to block the session even though the certificate may be valid. If you **Block sessions on certificate status check timeout** and the revocation server is slow to respond, you can use **Device** > **Setup** > **Session** > **Decryption Settings** and click **Certificate Revocation Checking** to change the default timeout value of five seconds to another value. For example, you could increase the timeout value to eight seconds, as shown in the following figure. Enable both CRL and OCSP certificate revocation checking because server certificates can contain the CRL URL in the CRL Distribution Point (CDP) extension or the OCSP URL in the Authority Information Access (AIA) certificate extension.



- **Restrict certificate extensions**—Checking this box limits the certificate extensions in the server certificate to key usage and extended key usage and blocks certificates with other extensions. However, in certain deployments, some other certificate extensions may be necessary, so only check this box if your deployment requires no other certificate extensions.
- **Append certificate's CN value to SAN extension**—Checking this box ensures that when a browser requires a server certificate to use a Subject Alternative Name (SAN) and doesn't support certificate matching based on the Common Name (CN), if the certificate doesn't have a SAN extension, users can still access the requested web resources because the firewall adds the SAN extension (based on the CN) to the impersonation certificate.

Unsupported Mode Checks. If you don't block sessions with unsupported modes, users receive a warning message if they connect with potentially unsafe servers, and they can click through that message and reach the potentially dangerous site. Blocking these sessions protects you from servers that use weak, risky protocol versions and algorithms:

- **Block sessions with unsupported versions**—When you configure the SSL Protocol Settings Decryption Profile, you specify the minimum version of SSL protocol to allow on your network to reduce the attack surface by blocking weak protocols. Always check this box to block sessions with the weak SSL/TLS protocol versions that you have chosen not to support.
- **Block sessions with unsupported cipher suites**—Always check this box to block sessions if the firewall doesn't support the cipher suite specified in the handshake. You configure which algorithms the firewall supports on the **SSL Protocol Settings** tab of the Decryption profile.

https://t.me/learningnets

- **Block sessions with client authentication**—If you have no critical applications that require client authentication, block it because firewall can't decrypt sessions that require client authentication. The firewall needs both the client and the server certificates to perform bi-directional decryption, but with client authentication, the firewall only knows the server certificate. This breaks decryption for client authentication sessions. When you check this box, the firewall blocks all sessions with client authentication except sessions from sites on the SSL Decryption Exclusion list (**Device** > **Certificate Management** > **SSL Decryption Exclusion**).

  If you don't **Block sessions with client authentication**, when the firewall attempts to decrypt a session that uses client authentication, the firewall allows the session and adds an entry that contains the server URL/IP address, the application, and the Decryption profile to its Local Decryption Exclusion Cache.

  You may need to allow traffic on your network from sites that use client authentication and are not in the Predefined sites on the SSL Decryption Exclusion list. Create a Decryption profile that allows sessions with client authentication. Add it to a Decryption policy rule that applies only to the server(s) that host the application. To increase security even more, you can require Multi-Factor Authentication to complete the user login process.

Failure Checks:

- **Block sessions if resources not available**—If you block sessions when no firewall processing resources are available, the firewall drops traffic when it doesn't have the resources to decrypt the traffic. If you don't block sessions when the firewall can't process decryption due to a lack of resources, then traffic that you want to decrypt enters the network still encrypted and therefore is not inspected. However, blocking sessions when resources aren't available may affect the user experience by making sites that users normally can reach temporarily unreachable. Whether to implement this failure check depends on your company's security compliance stance and the importance of the user experience, weighed against tighter security. Alternatively, consider using firewall models with more processing power so that you can decrypt more traffic.
- **Block sessions if HSM not available**—If you use a Hardware Security Module (HSM) to store your private keys, whether you use one depends on your compliance rules about where the private key must come from and how you want to handle encrypted traffic if the HSM isn't available. For example, if your company mandates the use of an HSM for private key signing, then block sessions if the HSM isn't available. However, if your company is less strict about this, then you can consider not blocking sessions if the HSM isn't available. (If the HSM is down, the firewall can process decryption for sites for which it has cached the response from the HSM, but not for other sites.) The best practice in this case depends on your company's policies. If the HSM is critical to your business, run the HSM in a high-availability (HA) pair (PAN-OS 8.1 supports two members in an HSM HA pair).
- **Block downgrade on no resource**—Prevents the firewall from downgrading TLSv1.3 to TLSv1.2 if the firewall has no available TLSv1.3 processing resources. If you block the downgrade, then when the firewall runs out of TLSv1.3 resources, it drops traffic that uses TLSv1.3 instead of downgrading it to TLSv1.2. If you don't block downgrade, then when the firewall runs out of TLSv1.3 resources, it downgrades to TLSv1.2. However, blocking downgrade when resources aren't available may affect the user experience by making sites that users normally can reach temporarily unreachable. Whether to implement this failure check depends on your company's security compliance stance and the importance of the user experience, weighed against tighter

https://t.me/learningnets

security. You may want to create a separate Decryption policy and profile to govern decryption for sensitive traffic for which you don't want to downgrade the TLS version.

SSL Inbound Inspection

Use SSL Inbound Inspection to decrypt and inspect inbound SSL/TLS traffic from a client to a targeted network server (any server you have the certificate for and can import it onto the firewall) and block suspicious sessions. For example, if an employee connects remotely to a web server hosted on the company network and attempts to add restricted internal documents to a Dropbox folder (which uses SSL for data transmission), SSL Inbound Inspection can ensure that the sensitive data does not move outside the secure company network by blocking or restricting the session.

On the firewall, you must install the certificate and private key for each server for which you want to perform SSL Inbound Inspection. You must also install the public key certificate as well as the private key on each firewall that performs SSL Inbound Inspection.

When you configure the SSL Protocol Settings Decryption Profile for SSL Inbound Inspection traffic, create separate profiles for servers with different security capabilities. For example, if one set of servers supports only RSA, the SSL Protocol Settings only need to support RSA. However, the SSL Protocol Settings for servers that support PFS should support PFS. Configure SSL Protocol Settings for the highest level of security that the server supports, but check performance to ensure that the firewall resources can handle the higher processing load that higher security protocols and algorithms require.

SSL Inbound Inspection works similarly to SSL Forward Proxy except that the firewall decrypts inbound traffic to servers instead of decrypting outbound traffic from internal clients. The firewall acts as a man-in-the-middle proxy between the external client and the internal server and must generate a new session key for each secure session. As a proxy, the firewall creates a secure session between the client and the firewall and another secure session between the firewall and the server so that the firewall can decrypt and inspect the traffic.

If you enabled SSL Inbound Inspection using PFS key exchange algorithms, you must upload a certificate bundle (a single file) to the firewall with your certificates arranged as follows:

1. *End-entity (leaf) certificate*
2. *Intermediate certificates (in issuing order)*
3. *(Optional) Root certificate*

Uploading the file ensures that clients receive the complete certificate chain during SSL handshakes, avoiding client-side server certificate authentication issues.

When you configure SSL Inbound Inspection and use a PFS cipher, session resumption is not supported.

See Configure SSL Inbound Inspection for details on enabling this feature.

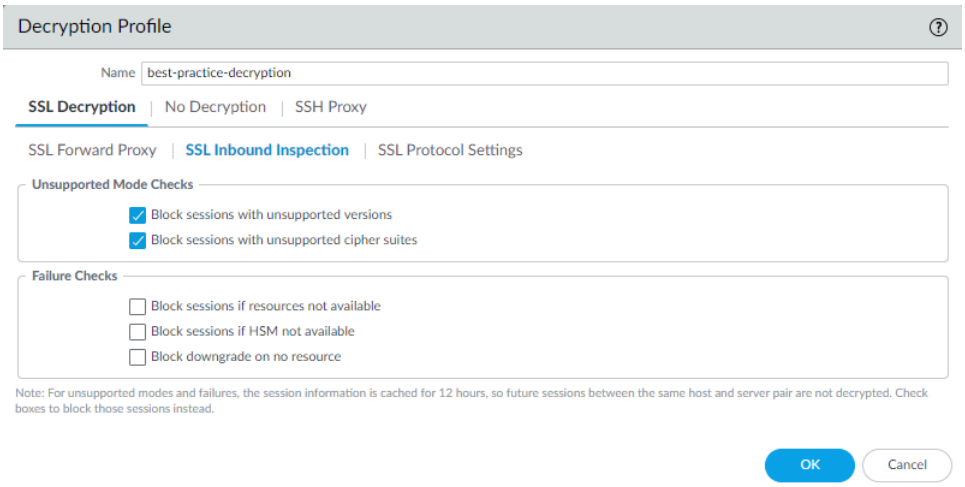
When you configure SSL Inbound Inspection, the proxied traffic does not support DSCP code points or QoS.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

SSL Inbound Inspection Decryption Profile

The SSL Inbound Inspection Decryption profile (**Objects** > **Decryption Profile** > **SSL Decryption** > **SSL Inbound Inspection**) controls the session mode checks and failure checks for inbound SSL/TLS traffic defined in the Inbound Inspection Decryption policies to which you attach the profile. The following figure shows the general best practice recommendations for Inbound Inspection Decryption profile settings, but the settings you use also depend on your company's security compliance rules and local laws and regulations.



Unsupported Mode Checks. If you don't block sessions with unsupported modes, users receive a warning message if they connect with potentially unsafe servers, and they can click through that message and reach the potentially dangerous site. Blocking these sessions protects you from servers that use weak, risky protocol versions and algorithms:

1. **Block sessions with unsupported versions**—When you configure the SSL Protocol Settings Decryption Profile, you specify the minimum version of TLS protocol to allow on your network to reduce the attack surface by blocking weak protocols. Always check this box to block sessions with the weak SSL and TLS protocol versions that you have chosen not to support.
2. **Block sessions with unsupported cipher suites**—Always check this box to block sessions if the firewall doesn't support the cipher suite specified in the handshake. You configure which algorithms the firewall supports on the **SSL Protocol Settings** tab of the Decryption profile.

Failure Checks:

- **Block sessions if resources not available**—If you block sessions when no firewall processing resources are available, the firewall drops traffic when it doesn't have the resources to decrypt the traffic. If you don't block sessions when the firewall can't process decryption due to a lack of resources, then traffic that you want to decrypt enters the network still encrypted and therefore is not inspected. However, blocking sessions when resources aren't available may affect the user experience by making sites that users normally can reach temporarily unreachable. Whether to implement this failure check depends on your company's security compliance stance and the importance of the user experience, weighed against tighter security. Alternatively, consider using firewall models with more processing power so that you can decrypt more traffic.
- **Block sessions if HSM not available**—If you use a Hardware Security Module (HSM) to store your private keys, whether you use one depends on your compliance rules about where the

https://t.me/learningnets

private key must come from and how you want to handle encrypted traffic if the HSM isn't available. For example, if your company mandates the use of an HSM for private key signing, then block sessions if the HSM isn't available. However, if your company is less strict about this, then you can consider not blocking sessions if the HSM isn't available. (If the HSM is down, the firewall can process decryption for sites for which it has cached the response from the HSM, but not for other sites.) The best practice in this case depends on your company's policies. If the HSM is critical to your business, run the HSM in a high-availability (HA) pair (PAN-OS 8.1 supports two members in an HSM HA pair).

- **Block downgrade on no resource**—Prevents the firewall from downgrading TLSv1.3 to TLSv1.2 if the firewall has no available TLSv1.3 processing resources. If you block the downgrade, then when the firewall runs out of TLSv1.3 resources, it drops traffic that uses TLSv1.3 instead of downgrading it to TLSv1.2. If you don't block downgrade, then when the firewall runs out of TLSv1.3 resources, it downgrades to TLSv1.2. However, blocking downgrade when resources aren't available may affect the user experience by making sites that users normally can reach temporarily unreachable. Whether to implement this failure check depends on your company's security compliance stance and the importance of the user experience, weighed against tighter security. You may want to create a separate Decryption policy and profile to govern decryption for sensitive traffic for which you don't want to downgrade the TLS version.

SSL Protocol Settings Decryption Profile

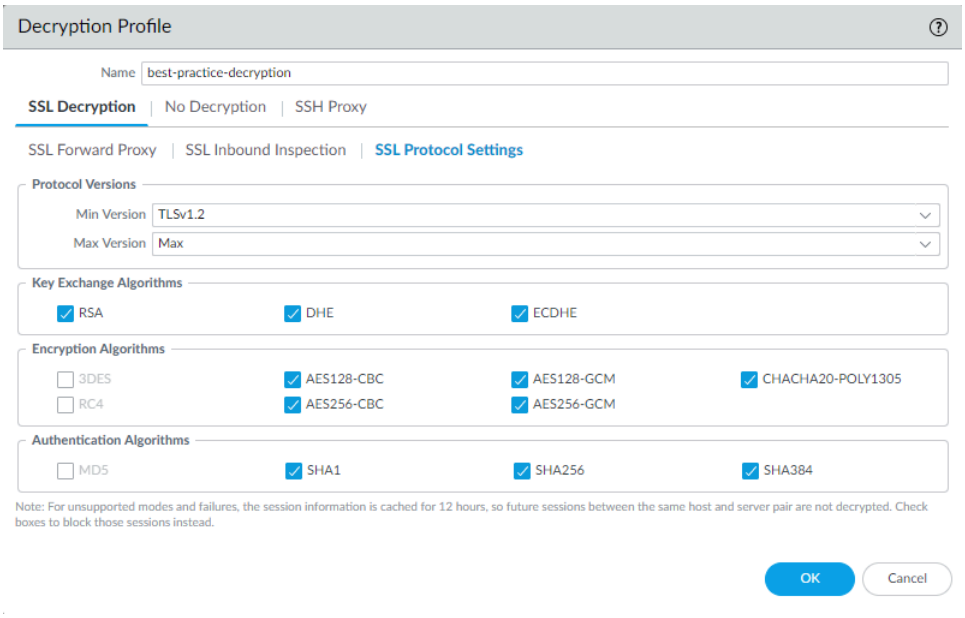
The SSL Protocol Settings (**Objects** > **Decryption Profile** > **SSL Decryption** > **SSL Protocol Settings**) control whether you allow vulnerable SSL/TLS protocol versions, weak encryption algorithms, and weak authentication algorithms. SSL Protocol Settings apply to outbound SSL Forward Proxy and inbound SSL Inbound Inspection traffic. These settings don't apply to SSH Proxy traffic or to traffic that you don't decrypt.

The following figure shows the general best practice recommendations for SSL Protocol Settings. There are also specific best practices for perimeter internet gateway decryption profiles and for data center decryption profiles.

When you configure SSL Protocol Settings for SSL Inbound Inspection traffic, create separate profiles for servers with different security capabilities. For example, if one set of servers supports only RSA, the SSL Protocol Settings only need to support RSA. However, the SSL Protocol Settings for servers that support PFS should support PFS. Configure SSL Protocol Settings for the highest level of security that the target server you are protecting supports, but check performance to ensure that the firewall resources can handle the higher processing load that higher security protocols and algorithms require.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets



Protocol Versions:

- Set the **Min Version** to **TLSv1.2** to provide the strongest security—business sites that value security support TLSv1.2. If a site (or a category of sites) only supports weaker ciphers, review the site and determine if it hosts a legitimate business application. If it does, make an exception for only that site by configuring a Decryption profile with a **Min Version** that matches the strongest cipher the site supports and then applying the profile to a Decryption policy rule that limits allowing the weak cipher to only the site or sites in question. If the site doesn't host a legitimate business application, don't weaken your security posture to support the site—weak protocols (and ciphers) contain known vulnerabilities that attackers can exploit.

  If the site belongs to a category of sites that you don't need for business purposes, use URL Filtering to block access to the entire category. Don't support weak encryption or authentication algorithms unless you must to support important legacy sites, and when you make exceptions, create a separate Decryption profile that allows the weaker protocol just for those sites. Don't downgrade the main Decryption profile that you apply to most sites to TLSv1.1 just to accommodate a few exceptions.

  Qualys SSL Labs SSL Pulse web page provides up-to-date statistics on the percentages of different ciphers and protocols in use on the 150,000 most popular sites in the world so you can see trends and understand how widespread worldwide support is for more secure ciphers and protocols.

- Set the **Max Version** to **Max** rather than to a particular version so that as the protocols improve, the firewall automatically supports the newest and best protocols. Whether you

https://t.me/learningnets

intend to attach a Decryption profile to a Decryption policy rule that governs inbound (SSL Inbound Inspection) or outbound (SSL Forward Proxy) traffic, avoid allowing weak algorithms.

*If your Decryption policy supports mobile applications, many of which use pinned certificates, set the **Max Version** to **TLSv1.2**. Because TLSv1.3 encrypts certificate information that was not encrypted in previous TLS versions, the firewall can't automatically add decryption exclusions based on certificate information, which affects some mobile applications. Therefore, if you enable TLSv1.3, the firewall may drop some mobile application traffic unless you create a No Decryption policy for that traffic.*

If you know the mobile applications you use for business, consider creating a separate Decryption policy and profile for those applications so that you can enable TLSv1.3 for all other application traffic.

Key Exchange Algorithms: Leave all three boxes checked (default) to support both RSA and PFS (DHE and ECDHE) key exchanges unless the minimum version is set to TLSv1.3, which only supports ECDHE.

To support HTTP/2 traffic, you must leave the ECDHE box checked.

Encryption Algorithms: When you set the minimum protocol version to TLSv1.2, the older, weaker 3DES and RC4 algorithms are automatically unchecked (blocked). When you set the minimum protocol version to TLSv1.3, the 3DES, RC4, AES128-CBC, and AES256-CBC algorithms are automatically blocked. For any traffic for which you must allow a weaker TLS protocol, create a separate Decryption profile and apply it only to traffic for that site, and deselect the appropriate boxes to allow the algorithm. Allowing traffic that uses the 3DES or RC4 algorithms exposes your network to excessive risk. If blocking 3DES or RC4 prevents you from accessing a site that you must use for business, create a separate Decryption profile and policy for that site. Don't weaken decryption for any other sites.

Authentication Algorithms: The firewall automatically blocks the older, weaker MD5 algorithm. When TLSv1.3 is the minimum version, the firewall also blocks SHA1. Do not allow MD5 authenticated traffic on your network; SHA1 is the weakest authentication algorithm you should allow. If no necessary sites use SHA1, block SHA1 traffic to further reduce the attack surface.

SSH Proxy

In an SSH Proxy configuration, the firewall resides between a client and a server. SSH Proxy enables the firewall to decrypt inbound and outbound SSH connections and ensures that attackers don't use SSH to tunnel unwanted applications and content. SSH decryption does not require certificates and the firewall automatically generates the key used for SSH decryption when the firewall boots up. During the boot up process, the firewall checks if there is an existing key. If not, the firewall generates a key. The firewall uses the key to decrypt SSH sessions for all virtual systems configured on the firewall and all SSH v2 sessions.

SSH allows tunneling, which can hide malicious traffic from decryption. The firewall can't decrypt traffic inside an SSH tunnel. You can block all SSH tunnel traffic by configuring a Security policy rule for the application **ssh-tunnel** with the **Action** set to **Deny** (along with a Security policy rule to allow traffic from the **ssh** application).

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

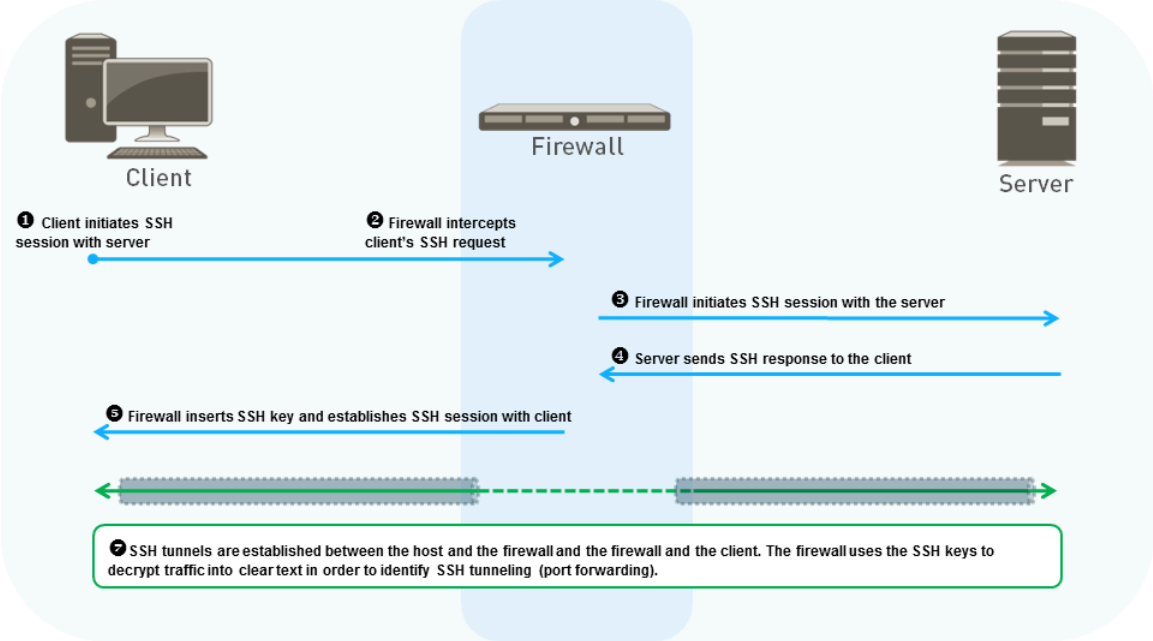
SSH tunneling sessions can tunnel X11 Windows packets and TCP packets. One SSH connection may contain multiple channels. When you apply an SSH Decryption profile to traffic, for each channel in the connection, the firewall examines the App-ID of the traffic and identifies the channel type. The channel type can be:

- session
- X11
- forwarded-tcpip
- direct-tcpip

When the channel type is session, the firewall identifies the traffic as allowed SSH traffic such as SFTP or SCP. When the channel type is X11, forwarded-tcpip, or direct-tcpip, the firewall identifies the traffic as SSH tunneling traffic and blocks it.

Limit SSH use to administrators who need to manage network devices, log all SSH traffic, and consider configuring Multi-Factor Authentication to help ensure that only legitimate users can use SSH to access devices, which reduces the attack surface.

The following figure shows how SSH Proxy decryption works. See Configure SSH Proxy for how to enable SSH Proxy decryption.



1. The client sends an SSH request to the server to initiate a session.
2. The firewall intercepts the client's SSH request .
3. The firewall forwards the request to the server and initiates an SSH session with the server. This establishes the first of two separate sessions that the firewall creates. Each session establishes a separate SSH tunnel.
4. The server responds to the request, which the firewall intercepts.
5. The firewall inserts the SSH key into the server's response and forwards it to the client. This establishes the second separate session (and separate SSH tunnel) that the firewall creates.

https://t.me/learningnets

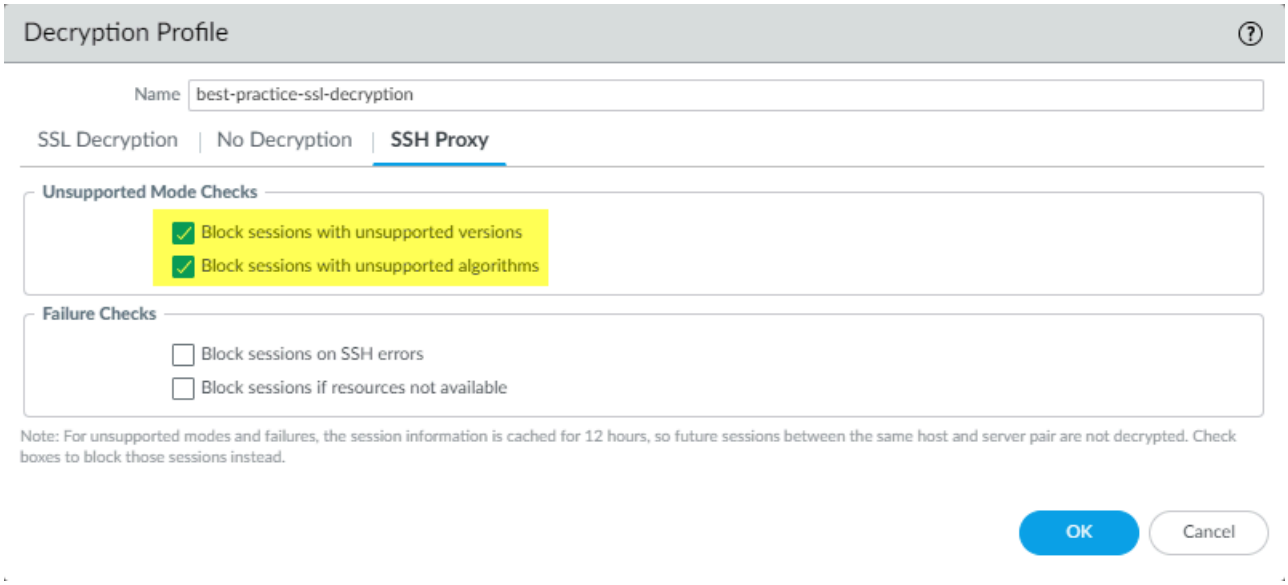
6. (First part of “7” in the diagram) After the firewall establishes separate sessions with the server and the client, the firewall acts as a proxy between them.
7. The firewall checks the traffic between the client and server to see if it is routed normally or if it uses SSH port forwarding (SSH tunneling). If the firewall identifies SSH port forwarding, the firewall blocks the tunneled traffic and restricts it according to the configured Security policy. The firewall only looks for SSH port forwarding, it does not perform content and threat inspection on SSH tunnels.

When you configure SSH Proxy, the proxied traffic does not support DSCP code points or QoS.

SSH Proxy Decryption Profile

The SSH Proxy Decryption profile (**Objects** > **Decryption Profile** > **SSH Proxy**) controls the session mode checks and failure checks for SSH traffic defined in the SSH Proxy Decryption policies to which you attach the profile. The following figure shows the general best practice recommendations for SSH Proxy Decryption profile settings, but the settings you use also depend on your company’s security compliance rules and local laws and regulations.

The firewall doesn’t perform content and threat inspection on SSH tunnels (port forwarding). However, the firewall distinguishes between the SSH application and the SSH-tunnel application. If the firewall identifies SSH tunnels, it blocks the SSH tunneled traffic and restricts the traffic according to configured security policies.



Unsupported Mode Checks. The firewall supports SSHv2. If you don’t block sessions with unsupported modes, users receive a warning message if they connect with potentially unsafe servers, and they can click through that message and reach the potentially dangerous site. Blocking these sessions protects you from servers that use weak, risky protocol versions and algorithms:

1. **Block sessions with unsupported versions**—The firewall has a set of predefined supported versions. Checking this box blocks traffic with weak versions. Always check this box to block sessions with the weak protocol versions to reduce the attack surface.

https://t.me/learningnets

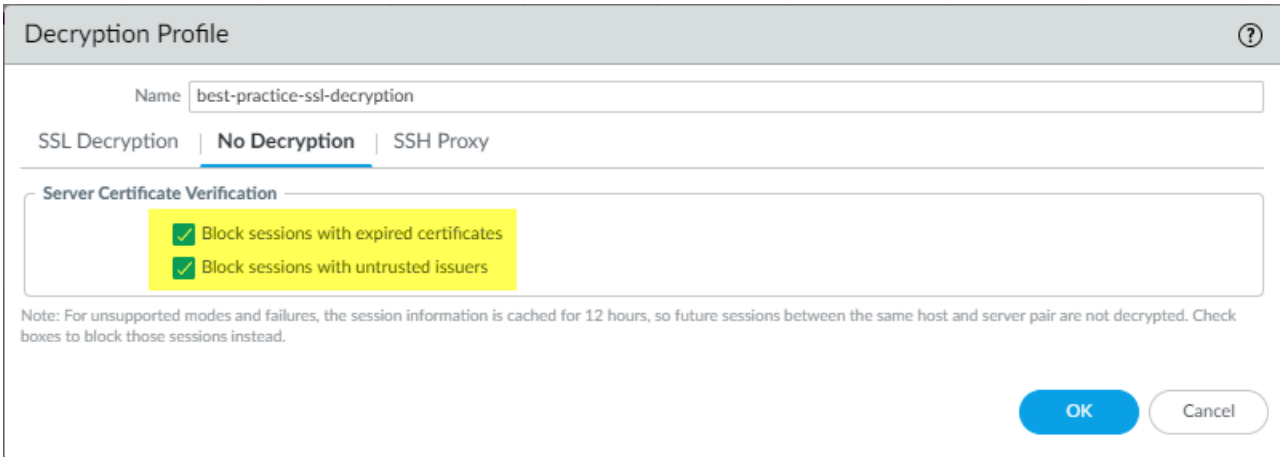
2. **Block sessions with unsupported algorithms**—The firewall has a set of predefined supported algorithms. Checking this box blocks traffic with weak algorithms. Always check this box to block sessions with unsupported algorithms to reduce the attack surface.

Failure Checks:

- **Block sessions on SSH errors**—Checking this box terminates the session if SSH errors occur.
- **Block sessions if resources not available**—If you don't block sessions when firewall processing resources aren't available, then encrypted traffic that you want to decrypt enters the network still encrypted, risking allowing potentially dangerous connections. However, blocking sessions when firewall processing resources aren't available may affect the user experience by making sites that users normally can reach temporarily unreachable. Whether to implement failure checks depends on your company's security compliance stance and the importance to your business of the user experience, weighed against tighter security. Alternatively, consider using firewall models with more processing power so that you can decrypt more traffic.

Profile for No Decryption

No Decryption profiles (**Objects** > **Decryption Profile** > **No Decryption**) perform server verification checks for traffic that you choose not to decrypt. You attach a No Decryption profile to a "No Decryption" Decryption policy that defines the traffic to exclude from decryption. (Don't use policy to exclude traffic that you can't decrypt because a site breaks decryption for technical reasons such as a pinned certificate or mutual authentication. Instead, add the hostname to the Decryption Exclusion List.) The following figure shows the general best practice recommendations for the No Decryption profile settings, but the settings you use also depend on your company's security compliance rules and local laws and regulations.



- **Block sessions with expired certificates**—Check this box to block sessions with servers that have expired certificates and prevent access to potentially insecure sites. If you don't check this box, users can connect with and transact with potentially malicious sites and see warning messages when they attempt to connect, but the connection is not prevented.
- **Block sessions with untrusted issuers**—Check this box to block sessions with servers that have untrusted certificate issuers. An untrusted issuer may indicate a man-in-the-middle attack, a replay attack, or other attack.

https://t.me/learningnets

Do not attach a No Decryption profile to Decryption policies for TLSv1.3 traffic that you don't decrypt. Unlike previous versions, TLSv1.3 encrypts certificate information, so the firewall has no visibility into certificate data and therefore cannot block sessions with expired certificates or untrusted issuers, so the profile has no effect. (The firewall can perform certificate checks with TLSv1.2 and earlier because those protocols do not encrypt certificate information and you should apply a No Decryption profile to their traffic.) However, you should create a Decryption policy for TLSv1.3 traffic that you don't decrypt because the firewall doesn't log undecrypted traffic unless a Decryption policy controls that traffic.

*(Applies to TLSv1.2 and earlier) If you choose to allow sessions with untrusted issuers (not recommended) and only **Block sessions with expired certificates**, there is a scenario in which a session with a trusted, expired issuer may be blocked inadvertently. When the firewall's certificate store contains a valid, self-signed Trusted CA and the server sends an expired CA in the certificate chain, the firewall does not check its certificate store. Instead, the firewall blocks the session based on the expired CA when it should find the trusted, valid alternative trust anchor and allow the session based on that trusted self-signed certificate.*

*To avoid this scenario, in addition to **Block sessions with expired certificates**, enable **Block sessions with untrusted issuers**. This forces the firewall to check its certificate store, find the self-signed Trusted CA, and allow the session.*

SSL Decryption for Elliptical Curve Cryptography (ECC) Certificates

The firewall automatically decrypts SSL traffic from websites and applications using ECC certificates, including Elliptical Curve Digital Signature Algorithm (ECDSA) certificates. As organizations transition to using ECC certificates to benefit from the strong keys and small certificate size, you can continue to maintain visibility into and safely enable ECC-secured application and website traffic.

Decryption for websites and applications using ECC certificates is not supported for traffic that is mirrored to the firewall; encrypted traffic using ECC certificates must pass through the firewall directly for the firewall to decrypt it.

You can use a hardware security module (HSM) to store the private keys associated with ECDSA certificates. For TLSv1.3 traffic, PAN-OS supports HSMs only for SSL Forward Proxy. It does not support HSMs for SSL Inbound Inspection.

Perfect Forward Secrecy (PFS) Support for SSL Decryption

PFS is a secure communication protocol that prevents the compromise of one encrypted session from leading to the compromise of multiple encrypted sessions. With PFS, a server generates unique private keys for each secure session it establishes with a client. If a server private key is compromised, only the single session established with that key is vulnerable—an attacker cannot retrieve data from past and future sessions because the server establishes each connected with a uniquely generated key. The firewall decrypts SSL sessions established with PFS key exchange algorithms, and preserves PFS protection for past and future sessions.

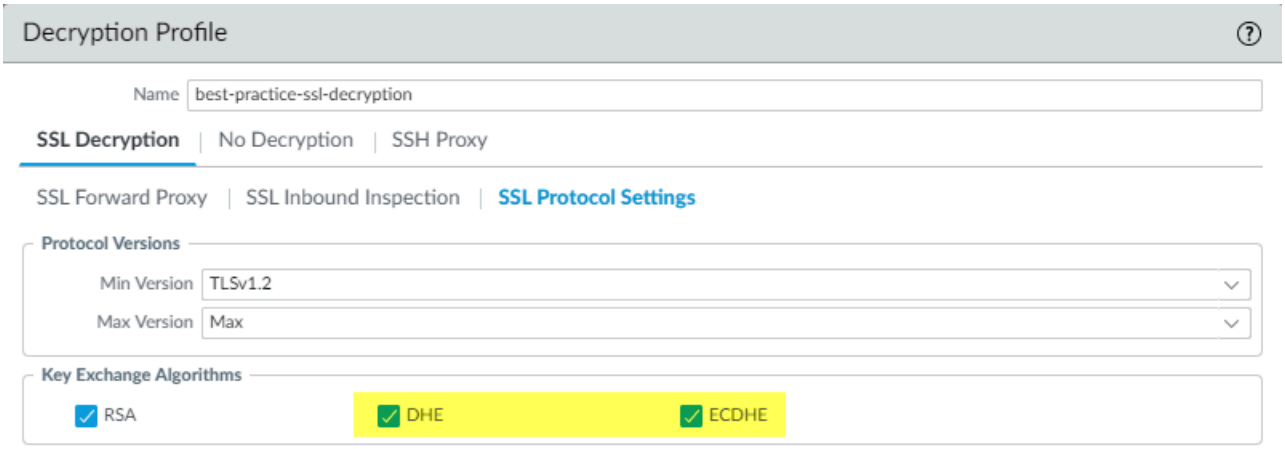
 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Support for Diffie-Hellman (DHE)-based PFS and elliptical curve Diffie-Hellman (ECDHE)-based PFS is enabled by default (**Objects** > **Decryption Profile** > **SSL Decryption** > **SSL Protocol Settings**).

If you use the DHE or ECDHE key exchange algorithms to enable PFS support for SSL decryption, you can use a hardware security module (HSM) to store the private keys for SSL Inbound Inspection.

When you configure SSL Inbound Inspection and use a PFS cipher, session resumption is not supported.

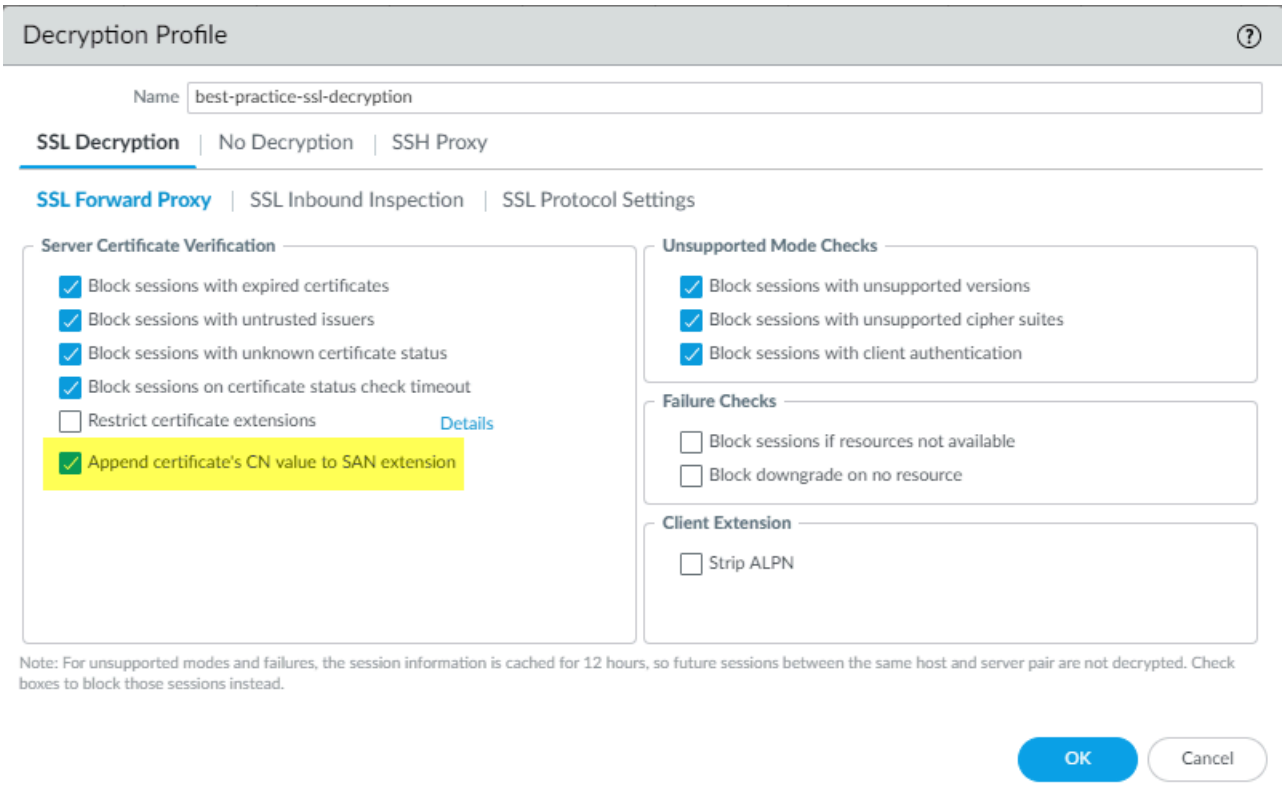


SSL Decryption and Subject Alternative Names (SANs)

Some browsers require server certificates to use a Subject Alternative Name (SAN) to specify the domains the certificate protects, and no longer support certificate matching based on a server certificate Common Name (CN). SANs enable a single server certificate to protect multiple names; CNs are less well-defined than SANs and can protect only a single domain or all first-level subdomains on a domain. However, if a server certificates contains only a CN, browsers that require a SAN will not allow end users to connect to the requested web resource.The firewall can add a SAN to the impersonation certificate it generates to establish itself as a trusted third-party during SSL decryption. When a server certificate contains only a CN, a firewall performing SSL decryption copies the server certificate CN to the impersonation certificate SAN. The firewall presents the impersonation certificate with the SAN to the client, and the browser is able to support the connection. End users can continue to access the resources they need, and the firewall can decrypt the sessions.

To enable SAN support for decrypted SSL traffic, update the decryption profile attached to the relevant decryption policy: select **Objects** > **Decryption Profile** > **SSL Decryption** > **SSL Forward Proxy** > **Append certificate's CN value to SAN extension**).

https://t.me/learningnets



TLSv1.3 Decryption

You can decrypt, gain full visibility into, and prevent known and unknown threats in TLSv1.3 traffic. TLSv1.3 is the latest version of the TLS protocol, which provides application security and performance improvements. Your existing Decryption policies work with TLSv1.3 when you configure the associated Decryption profile to use TLSv1.3 as the minimum protocol version or to use TLSv1.3 or Max as the maximum protocol version. The firewall supports TLSv1.3 decryption for Forward Proxy, Inbound Inspection, decrypted Network Packet Broker traffic, and Decryption Port Mirroring.

To use TLSv1.3, the client and server must be able to negotiate TLSv1.3 ciphers. For websites that don't support TLSv1.3, the firewall selects an older version of the TLS protocol that the server supports.

The firewall supports the following decryption algorithms for TLSv1.3:

- TLS13-AES-128-GCM-SHA256
- TLS13-AES-256-GCM-SHA384
- TLS13-CHACHA20-POLY1305-SHA256

If the Decryption profile you apply to decrypted traffic specifies the protocol's **Max Version** as **Max**, then the profile supports TLSv1.3 and automatically uses TLSv1.3 with sites that support TLSv1.3. (You could set the **Max Version** to **TLSv1.3** to support TLSv1.3, but when the next version of TLS is released, you will need to update the profile. Setting the **Max Version** to **Max** future-proofs the profile to automatically support new TLS versions as they are released.) When you upgrade to PAN-OS 10.0, all Decryption profiles with the **Max Version** set to **Max** are reset

https://t.me/learningnets

to **TLSv1.2** to provide automatic support for mobile applications that use pinned certificates and prevent that traffic from dropping.

Not all applications support the TLSv1.3 protocol. Follow decryption best practices, set the **Min Version** of the TLS protocol to **TLSv1.2**, and leave the **Max Version** setting as **Max**. If business needs require allowing a weaker TLS protocol, create a separate SSL Decryption profile with a **Min Version** that allows the weaker protocol and attach it to a Decryption policy that defines the traffic you need to allow with the weaker TLS protocol.

If your Decryption policy supports mobile applications, many of which use pinned certificates, set the **Max Version** to **TLSv1.2**. Because TLSv1.3 encrypts certificate information that was not encrypted in previous TLS versions, the firewall can't automatically add decryption exclusions based on certificate information, which affects some mobile applications. Therefore, if you enable TLSv1.3, the firewall may drop some mobile application traffic unless you create a No Decryption policy for that traffic. If you know the mobile applications you use for business, consider creating a separate Decryption policy and profile for those applications so that you can enable TLSv1.3 for all other traffic.



Do not attach a No Decryption profile to Decryption policies for TLSv1.3 traffic that you don't decrypt if you know that a particular policy controls only TLSv1.3 traffic. A change from previous TLS versions is that TLSv1.3 encrypts certificate information, so the firewall no longer has visibility into that data and therefore cannot block sessions with expired certificates or untrusted issuers, so the profile has no effect. (The firewall can perform certificate checks with TLSv1.2 and earlier because those protocols do not encrypt certificate information and you should apply a No Decryption profile to their traffic.) However, you can log undecrypted traffic of all types by enabling logging successful and unsuccessful TLS handshakes in the Decryption policy (logging unsuccessful TLS handshakes is enabled by default).

When you allow unsupported modes in the SSL Protocol Settings Decryption Profile, the firewall automatically adds the traffic to the Local Decryption Exclusion Cache. The firewall still decrypts and inspects traffic that is downgraded from TLSv1.3 to TLSv1.2 and the **Reason** shown in the cache for adding the server to the cache is TLS13_UNSUPPORTED.

If you downgrade from PAN-OS 10.1 to a previous version, any Decryption profile that specifies TLSv1.3 as the **Min Version** or the **Max Version** changes to the highest supported version. For example, downgrading from PAN-OS 10.1 to PAN-OS 9.1 would replace TLSv1.3 with TLSv1.2. If a Panorama device on PAN-OS 10.1 pushes the configuration to devices that run older versions of PAN-OS, any Decryption profile that specified TLSv1.3 as the **Min Version** or the **Max Version** also changes to highest supported version.



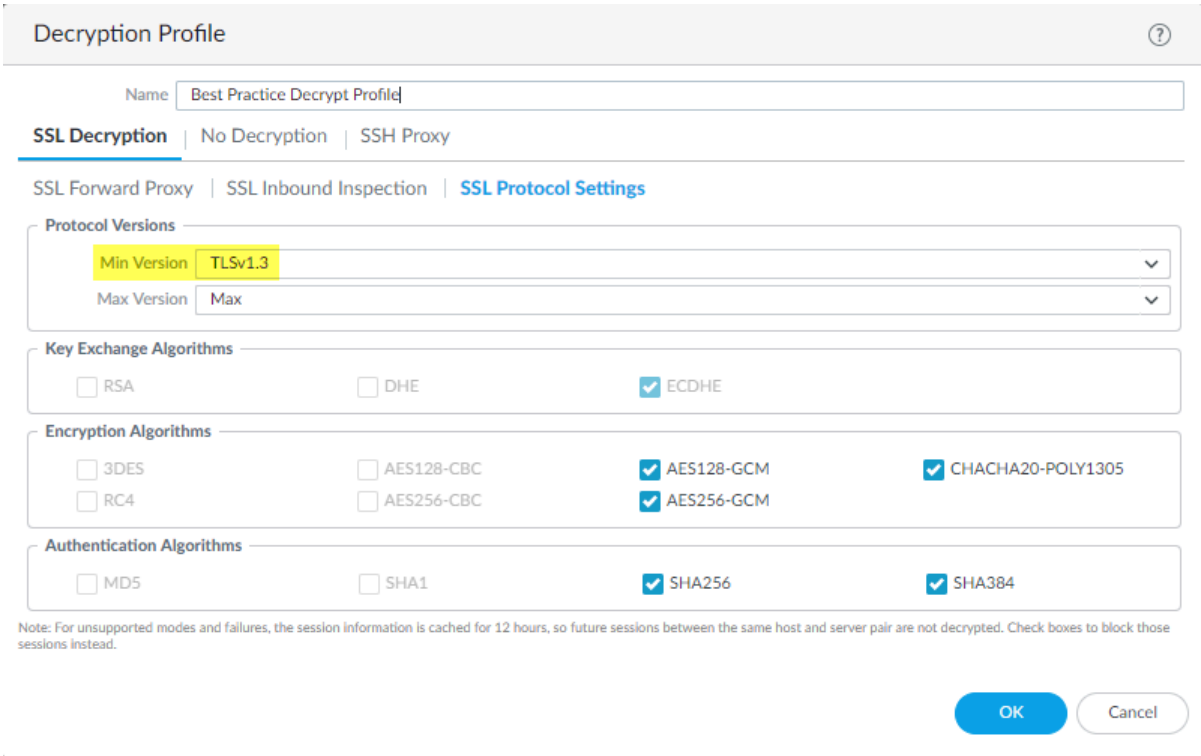
For customers who use Hardware Security Modules (HSMs), PAN-OS supports TLSv1.3 only for SSL Forward Proxy. It does not support HSMs for SSL Inbound Inspection.

You can configure an SSL Decryption profile that sets TLSv1.3 as the minimum allowed protocol version to achieve the tightest security. However, some applications don't support TLSv1.3 and may not work if TLSv1.3 is the minimum allowed protocol. Apply a profile that sets TLSv1.3 as the minimum version only to application traffic that only supports TLSv1.3.

https://t.me/learningnets

1. Create a new SSL Decryption profile or edit an existing profile (**Objects** > **Decryption** > **Decryption Profile**).

   If the profile is new, specify a profile **Name**.
2. Select **SSL Protocol Settings**.
3. Change the **Min Version** to **TLSv1.3**.



   Using **Max** for the **Max Version** ensures that the traffic which the profile controls can use the strongest available protocol version. **Min Version** sets the weakest version of the protocol that the traffic can use. Setting the minimum version to **TLSv1.3** means that the traffic must use TLSv1.3 (or greater) and that weaker protocol versions are blocked. (The Decryption Policy rule defines the traffic the profile controls.)

   When you configure TLSv1.3 as the **Min Version**, you must use Perfect Forward Secrecy (PFS) and the weaker key exchange, encryption, and authentication algorithms are not available.
4. Configure any other Decryption profile settings you need to set or change.
5. Click **OK** to save the profile.
6. Attach the profile to the appropriate Decryption Policy rule to apply it to the appropriate traffic.

High Availability Not Supported for Decrypted Sessions

After a failover, firewalls do not support High Availability (HA) sync for decrypted SSL sessions. The firewall does not resume decrypted SSL Forward Proxy, SSL Inbound Inspection, or SSH Proxy sessions. The firewall decrypts new sessions that start after the failover based on Decryption policy.

https://t.me/learningnets

Decryption Mirroring

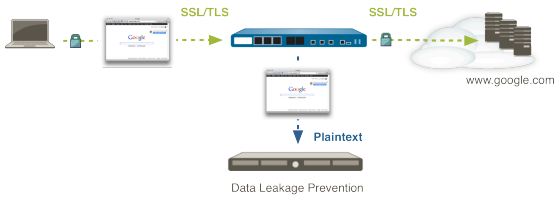
Decryption mirroring creates a copy of decrypted traffic from a firewall and sends it to a traffic collection tool such as NetWitness or Solera, which can receive raw packet captures for archiving and analysis. Organizations that require comprehensive data capture for forensic and historical purposes or for data leak prevention (DLP) can install a free license to enable the feature.

After you install the license, connect the traffic collection tool directly to an Ethernet interface on the firewall and set the **Interface Type** to **Decrypt Mirror**. The firewall simulates a TCP handshake with the collection tool and then sends every data packet through that interface, decrypted (as cleartext).

Decryption port mirroring is not available on the VM-Series for public cloud platforms (AWS, Azure, Google Cloud Platform) and VMware NSX.

Keep in mind that the decryption, storage, inspection, and/or use of SSL traffic is governed in certain countries and user consent might be required in order to use the decryption mirror feature. Additionally, use of this feature could enable malicious users with administrative access to the firewall to harvest usernames, passwords, social security numbers, credit card numbers, or other sensitive information submitted using an encrypted channel. Palo Alto Networks recommends that you consult with your corporate counsel before activating and using this feature in a production environment.

The following graphic shows the process for mirroring decrypted traffic and the section Configure Decryption Port Mirroring describes how to license and enable this feature.



 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Prepare to Deploy Decryption

The most time-consuming part of deploying decryption isn't configuring the decryption policies and profiles, it's preparing for the deployment by working with stakeholders to decide what traffic to decrypt and not to decrypt, educating your user population about changes to website access, developing a private key infrastructure (PKI) strategy, and planning a staged, prioritized rollout.

Set goals for decryption and review Decryption planning best practices checklist to ensure that you understand the recommended best practices. The best practice goal is to decrypt as much traffic as your firewall resources permit and decrypt the most important traffic first.



Migrate from port-based to application-based Security policy rules before you create and deploy Decryption policy rules. If you create Decryption rules based on port-based Security policy and then migrate to application-based Security policy, the change could cause the Decryption rules to block traffic that you intend to allow because Security policy rules are likely to use application default ports to prevent application traffic from using non-standard ports. For example, traffic identified as web-browsing application traffic (default port 80) may have underlying applications that have different default ports, such as HTTPS traffic (default port 443). The application-default rule blocks the HTTPS traffic because it sees the decrypted traffic using a "non-standard" port (443 instead of 80). Migrating to App-ID based rules before deploying decryption means that when you test your decryption deployment in POCs, you'll discover Security policy misconfiguration and fix it before rolling it out to the general user population.

To prepare to deploy Decryption:

- Work with Stakeholders to Develop a Decryption Deployment Strategy
- Develop a PKI Rollout Plan
- Size the Decryption Firewall Deployment
- Plan a Staged, Prioritized Deployment

Work with Stakeholders to Develop a Decryption Deployment Strategy

Work with stakeholders such as legal, finance, HR, executives, security, and IT/support to develop a decryption deployment strategy. Start by getting the required approvals to decrypt traffic to secure the corporation. Decrypting traffic involves understanding how legal regulations and business needs affect what you can and can't decrypt.

Identify and prioritize the traffic you want to decrypt. The best practice is to decrypt as much traffic as you can to gain visibility into potential threats in encrypted traffic and prevent those threats. If incorrect firewall sizing prevents you from decrypting all of the traffic you want to decrypt, prioritize the most critical servers, the highest-risk traffic categories, and less trusted segments and IP subnets. To help prioritize, ask yourself questions such as, "What happens if this server is compromised?" and "How much risk am I willing to take in relation to the level of performance I want to achieve?"

https://t.me/learningnets

Next, identify traffic that you can't decrypt because the traffic breaks decryption for technical reasons such as a pinned certificate, an incomplete certificate chain, unsupported ciphers, or mutual authentication. Decrypting sites that break decryption technically results in blocking that traffic. Evaluate the websites that break decryption technically and ask yourself if you need access to those sites for business reasons. If you don't need access to those sites, allow decryption to block them. If you need access to any of those sites for business purposes, add them to the SSL Decryption Exclusion list to except them from decryption. The SSL Decryption Exclusion list is exclusively for sites that break decryption technically.

Identify sensitive traffic that you *choose* not to decrypt for legal, regulatory, personal, or other reasons, such as financial, health, or government traffic, or the traffic of certain executives. This is not traffic that breaks decryption technically, so you don't use the SSL Decryption Exclusion list to except this traffic from decryption. Instead, you Create a Policy-Based Decryption Exclusion to identify and control traffic you choose not to decrypt and apply the No Decryption decryption profile to the policy to prevent servers with certificate issues from accessing the network. Policy-based decryption exclusions are only for traffic you choose not to decrypt.

When you plan decryption policy, consider your company's security compliance rules, computer usage policy, and your business goals. Extremely strict controls can impact the user experience by preventing access to non-business sites the user used to access, but may be required for government or financial institutions. There is always a tradeoff between usability, management overhead, and security. The tighter the decryption policy, the greater the chance that a website will become unreachable, which may result in user complaints and possibly modifying the rulebase.



Although a tight decryption policy may initially cause a few user complaints, those complaints can draw your attention to unsanctioned or undesirable websites that are blocked because they use weak algorithms or have certificate issues. Use complaints as a tool to better understand the traffic on your network.

Different groups of users and even individual users may require different decryption policies, or you may want to apply the same decryption policy to all users. For example, executives may be exempted from decryption policies that apply to other employees. And you may want to apply different decryption policies to employee groups, contracts, partners, and guests. Prepare updated legal and HR computer usage policies to distribute to all employees, contractors, partners, guests, and any other network users so that when you roll out decryption, users understand their data can be decrypted and scanned for threats.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

How you handle guest users depends on the access they require. Isolate guests from the rest of your network by placing them on a separate VLAN and on a separate SSID for wireless access. If guests don't need to access your corporate network, don't let them on it and there will be no need to decrypt their traffic. If guests need to access your corporate network, decrypt their traffic:

- *Enterprises don't control guest devices. Decrypt guest traffic and subject it to your guest Security policy so the firewall can inspect the traffic and prevent threats. To do this, redirect guest users through an Authentication Portal, instruct them how to download and install the CA certificate, and clearly notify guests that their traffic will be decrypted. Include the process in your company's privacy and computer usage policy.*
- *Create separate Decryption policy rules and Security policy rules to tightly control guest access so that guests can only access the areas of your network that they need to access.*

Similarly to different groups of users, decide which devices to decrypt and which applications to decrypt. Today's networks support not only corporate devices, but BYOD, mobile, remote-user and other devices, including contractor, partner, and guest devices. Today's users attempt to access many sites, both sanctioned and unsanctioned, and you should decide how much of that traffic you want to decrypt.

Enterprises don't control BYOD devices. If you allow BYOD devices on your network, decrypt their traffic and subject it to the same Security policy that you apply to other network traffic so the firewall can inspect the traffic and prevent threats. To do this, redirect BYOD users through an Authentication Portal, instruct them how to download and install the CA certificate, and clearly notify users that their traffic will be decrypted. Educate BYOD users about the process and include it in your company's privacy and computer usage policy.

Decide what traffic you want to log and investigate what traffic you can log. Be aware of local laws regarding what types of data you can log and store, and where you can log and store the data. For example, local laws may prevent logging and storing personal information such as health and financial data.

Decide how to handle bad certificates. For example, will you block or allow sessions for which the certificate status is unknown? Understanding how you want to handle bad certificates determines how you configure the decryption profiles that you attach to decryption policies to control which sessions you allow based on the server certificate verification status.

Develop a PKI Rollout Plan

Plan how to roll out your public key infrastructure (PKI). Network devices need an SSL Forward Trust CA certificate for trusted sites and an SSL Forward Untrust CA certificate for untrusted sites. Generate separate Forward Trust and Forward Untrust certificates (do not sign the Forward Untrust certificate with the Enterprise Root CA because you want the Untrust certificate to warn users that they are trying to access potentially unsafe sites). Palo Alto Networks next-generation firewalls have two methods of generating CA certificates for SSL decryption:

- **Generate the SSL CA certificates from your Enterprise Root CA as subordinate certificates**—If you have an existing Enterprise PKI, this is the best practice. Generating a subordinate

https://t.me/learningnets

certificate from your Enterprise Root CA makes the rollout easier and smoother because network devices already trust the Enterprise Root CA, so you avoid any certificate issues when you begin the deployment phase. If you don't have an Enterprise Root CA, consider getting one.

- **Generate a self-signed Root CA certificate on the firewall and create subordinate CA certificates on that firewall**—If you don't have an Enterprise Root CA, this method provides a self-signed Root CA certificate and the subordinate Forward Trust and Untrust CA certificates. With this method, you need to install the self-signed certificates on all of your network devices so that those devices recognize the firewall's self-signed certificates. Because the certificates must be deployed to all devices, this method is better for small deployments and proof-of-concept (POC) trials than for large deployments.



Do not export the Forward Untrust certificate to the Certificate Trust Lists of your network devices! This is critical because installing the Untrust certificate in the Trust List results in devices trusting websites that the firewall does not trust. In addition, users won't see certificate warnings for untrusted sites, so they won't know the sites are untrusted and may access those sites, which could expose your network to threats.



Regardless of whether you generate Forward Trust certificates from your Enterprise Root CA or use a self-signed certificate generated on the firewall, generate a separate subordinate Forward Trust CA certificate for each firewall. The flexibility of using separate subordinate CAs enables you to revoke one certificate when you decommission a device (or device pair) without affecting the rest of the deployment and reduces the impact in any situation in which you need to revoke a certificate. Separate Forward Trust CAs on each firewall also helps troubleshoot issues because the CA error message the user sees includes information about the firewall the traffic is traversing. If you use the same Forward Trust CA on every firewall, you lose the granularity of that information.

There is no benefit to using different Forward Untrust certificates on different firewalls, so you can use the same Forward Untrust certificate on all firewalls. If you need additional security for your private keys, consider storing them on an HSM.

You may need to make special accommodations for guest users. If guest users don't need access to your corporate network, don't allow access, and then you won't have to decrypt their traffic or create infrastructure to support guest access. If you need to support guest users, discuss with your legal department whether you can decrypt guest traffic.

If you can decrypt guest traffic, treat guests similarly to the way you treat BYOD devices. Decrypt guest traffic and subject it to the same Security policy that you apply to other network traffic. Do this by redirecting guest users through an Authentication Portal, instruct them how to download and install the CA certificate, and clearly notify users that their traffic will be decrypted. Include the process in your company's privacy and computer usage policy. In addition, restrict guest traffic to only the areas guests need to access.

If you can't decrypt guest traffic for legal reasons, then isolate guest traffic and prevent it from moving laterally in your network:

- Create a separate zone for guests and restrict guest access to that zone. To prevent lateral movement, don't allow guest access to other zones.
- Allow only sanctioned applications, use URL filtering to prevent access to risky URL categories, and apply the best practice Security profiles.

https://t.me/learningnets

- Apply a No Decrypt decryption policy and profile to prevent guests from accessing websites with unknown or expired CAs.

All employees, contractors, partners, and other users should use your normal corporate infrastructure and you should decrypt and inspect their traffic.

Size the Decryption Firewall Deployment

Decrypting encrypted traffic consumes firewall CPU resources and can affect throughput. In general, the tighter the security (the more SSL traffic you decrypt combined with the more stringent your protocol settings), the more firewall resources decryption consumes. Work with your Palo Alto Networks SE/CE to size your firewall deployment and avoid sizing mistakes. Factors that affect decryption resource consumption and therefore how much traffic the firewall can decrypt include:

- The amount of SSL traffic you want to decrypt. This varies from network to network. For example, some applications must be decrypted to prevent the injection of malware or exploits into the network or unauthorized data transfers, some applications can't be decrypted due to local laws and regulations or business reasons, and other applications are cleartext (unencrypted) and don't need to be decrypted. The more traffic you want to decrypt, the more resources you need.
- The TLS protocol version. Higher versions are more secure but consume more resources. Use the highest TLS protocol version you can to maximize security.
- The key size. The larger the key size, the better the security, but also the more resources the key processing consumes.
- The key exchange algorithm. Perfect Forward Secrecy (PFS) ephemeral key exchange algorithms such as Diffie-Hellman Ephemeral (DHE) Elliptic-Curve Diffie-Hellman Exchange (ECDHE) consume more processing resources than Rivest-Shamir-Adleman (RSA) algorithms. PFS key exchange algorithms provide greater security than RSA key exchange algorithms because the firewall has to generate a new cipher key for each session—but generating the new key consumes more firewall resources. However, if an attacker compromises a session key, PFS prevents the attacker from using it to decrypt any other sessions between the same client and server and RSA does not.
- The encryption algorithm. The key exchange algorithm determines whether the encryption algorithm is PFS or RSA.
- The certificate authentication method. RSA (not the RSA key exchange algorithm) consumes less resources than Elliptic Curve Digital Signature Algorithm (ECDSA) but ECDSA is more secure.

  The combination of the key exchange algorithm and the certificate authentication method affect throughput performance as shown in RSA and ECDSA benchmark tests. The performance cost of PFS trades off against the higher security that PFS achieves, but PFS may not be needed for all types of traffic. You can save firewall CPU cycles by using RSA for traffic that you want to decrypt and inspect for threats but that isn't sensitive.

- Average transaction sizes. For example, small average transaction sizes consume more processing power to decrypt. Measure the average transaction size of all traffic, then measure the average transaction size of traffic on port 443 (the default port for HTTPS encrypted traffic)

https://t.me/learningnets

to understand the proportion of encrypted traffic going to the firewall in relation to your total traffic and the average transaction sizes. Eliminate anomalous outliers such as unusually large transactions to get a truer measurement of average transaction size.

- The firewall model and resources. Newer firewall models have more processing power than older models.

The combination of these factors determines how decryption consumes firewall processing resources. To best utilize the firewall's resources, understand the risks of the data you're protecting. If firewall resources are an issue, use stronger decryption for higher-priority traffic and use less processor-intensive decryption to decrypt and inspect lower-priority traffic until you can increase the available resources. For example, you could use RSA instead of ECDHE and ECDSA for traffic that isn't sensitive or high-priority to preserve firewall resources for using PFS-based decryption for higher priority, sensitive traffic. (You're still decrypting and inspecting the lower-priority traffic, but trading off consuming fewer computational resources with using algorithms that aren't as secure as PFS.) The key is to understand the risks of different traffic types and treat them accordingly.

Measure firewall performance so that you understand the currently available resources, which helps you understand whether you need more firewall resources to decrypt the traffic you want to decrypt. Measuring firewall performance also sets a baseline for performance comparisons after deploying decryption.

When you size the firewall deployment, base it not only on your current needs, but also on your future needs. Include headroom for the growth of decryption traffic because Gartner predicts that through 2019, more than 80 percent of enterprise web traffic will be encrypted, and more than 50 percent of new malware campaigns will use various forms of encryption. Work with your Palo Alto Networks representatives and take advantage of their experience in sizing firewalls to help you size your firewall decryption deployment.

Plan a Staged, Prioritized Deployment

Plan to roll out decryption in a controlled manner, piece by piece. Don't roll out your entire decryption deployment at one time. Test and ensure that decryption is working as planned and that users understand what you are doing and why. Rolling out decryption in this manner makes it easier to troubleshoot in case anything doesn't work as expected and helps users adjust to the changes.

Educating stakeholders, employees, and other users such as contractors and partners is critical because decryption settings may change their ability to access some websites. Users should understand how to respond to situations in which previously reachable websites become unreachable and what information to give technical support. Support should understand what is being rolled out when and how to help users who encounter issues. Before you roll out decryption to the general population:

- Identify early adopters to help champion decryption and who will be able to help other employees who have questions during the full rollout. Enlist the help of department managers and help them understand the benefits of decrypting traffic.
- Set up proof-of-concept (POC) trials in each department with early adopters and other employees who understand why decrypting traffic is important. Educate POC participants about the changes and how to contact technical support if they run into issues. In this way, decryption POCs become an opportunity to work with technical support to POC how to

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

support decryption and to develop the most painless method for supporting the general rollout. The interaction between POC users and technical support also allows you to fine-tune policies and how to communicate with users.

POCs enable you to experiment with prioritizing what to decrypt first, so that when you phase in decryption in the general population, your POC experience helps you understand how to phase in decrypting different URL Categories. Measure the way decryption affects firewall CPU and memory utilization to help understand if the firewall sizing is correct or if you need to upgrade. POCs can also reveal applications that break decryption technically (decrypting them blocks their traffic) and need to be added to the Decryption Exclusion list.

When you set up POCs, also set up a user group that can certify the operational readiness and procedures prior to the general rollout.

- Educate the user population before the general rollout, and plan to educate new users as they join the company. This is a critical phase of deploying decryption because the deployment may affect websites that users previously visited but are not safe, so those sites are no longer reachable. The POC experience helps identify the most important points to communicate.
- Phase in decryption. You can accomplish this several ways. You can decrypt the highest priority traffic first (for example, the URL Categories most likely to harbor malicious traffic, such as gaming) and then decrypt more as you gain experience. Alternatively, you can take a more conservative approach and decrypt the URL Categories that don't affect your business first (so if something goes wrong, no issues occur that affect business), for example, news feeds. In all cases, the best way to phase in decryption is to decrypt a few URL Categories, take user feedback into account, run reports to ensure that decryption is working as expected, and then gradually decrypt a few more URL Categories and verify, and so on. Plan to make Decryption Exclusions to exclude sites from decryption if you can't decrypt them for technical reasons or because you choose not to decrypt them.

  If you Enable Users to Opt Out of SSL Decryption (users see a response page that allows them either to opt out of decryption and end the session without going to the site or to proceed to the site and agree to have the traffic decrypted), educate them about what it is, why they're seeing it, and what their options are.
- Create realistic deployment schedules that allow time to evaluate each stage of the rollout.



Place firewalls in positions where they can see all of the network traffic so that no encrypted traffic inadvertently gains access to your network because it bypasses the firewall.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Define Traffic to Decrypt

A Decryption policy rule allows you to define traffic that you want the firewall to decrypt and to define traffic that you choose to exclude from decryption because the traffic is personal or because of local regulations, for example.

Attach a Decryption profile to each Decryption policy rule to enable certificate checks, session mode checks, failure checks, and protocol and algorithm checks, depending on the profile. These checks prevent risky connections, such as sessions with untrusted certificate issuers, weak protocols, ciphers, and algorithms, and servers that have certificate issues.

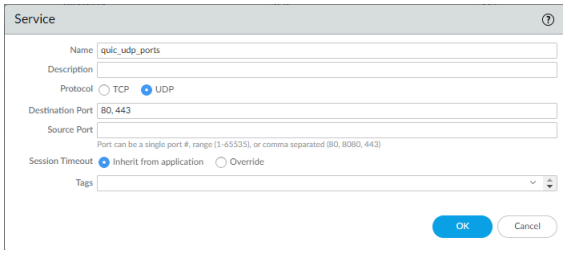
Review the Decryption deployment best practices checklist to ensure that you understand the recommended best practices.

Block known dangerous URL Filtering categories such as malware, phishing, dynamic-dns, unknown, command-and-control, proxy-avoidance-and-anonymizers, copyright-infringement, extremism, newly-registered-domain, grayware, and parked. If you must allow any of these categories for business reasons, decrypt them and apply strict Security profiles to the traffic.

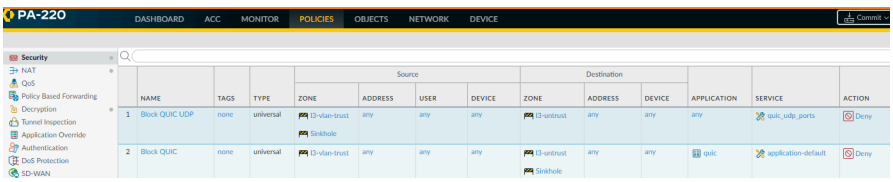
URL categories that you should always decrypt if you allow them include: online-storage-and-backup, web-based-email, web-hosting, personal-sites-and-blogs, and content-delivery-networks.

In Security policy, block Quick UDP Internet Connections (QUIC) protocol unless for business reasons, you want to allow encrypted browser traffic. Chrome and some other browsers establish sessions using QUIC instead of TLS, but QUIC uses proprietary encryption that the firewall can't decrypt, so potentially dangerous traffic may enter the network as encrypted traffic. Blocking QUIC forces the browser to fall back to TLS and enables the firewall to decrypt the traffic.

*Create a Security policy rule to block QUIC on its UDP service ports (80 and 443) and create a separate rule to block the QUIC application. For the rule that blocks UDP ports 80 and 443, create a Service (**Objects** > **Services**) that includes UDP ports 80 and 443:*



Use the Service to specify the UDP ports to block for QUIC. In the second rule, block the QUIC application:



- Create a Decryption Profile

https://t.me/learningnets

- Create a Decryption Policy Rule

Create a Decryption Profile

A decryption profile allows you to perform checks on both decrypted traffic and SSL traffic that you *choose* to exclude from decryption. (If a server breaks SSL decryption technically due to certificate pinning or other reasons, add the server to the Decryption Exclusion list.) Depending on your needs, create Decryption profiles to:

- Block sessions based on certificate status, including blocking sessions with expired certificates, untrusted issuers, unknown certificate status, certificate status check timeouts, and certificate extensions.
- Block sessions with unsupported versions and cipher suites, and that require using client authentication.
- Block sessions if the resources to perform decryption are not available or if a hardware security module is not available to sign certificates.
- Define the protocol versions and key exchange, encryption, and authentication algorithms allowed for SSL Forward Proxy and SSL Inbound Inspection traffic in the SSL Protocol Settings.

Don't weaken the main Decryption profile that you apply to most sites to accommodate weaker sites. Instead, create one or more separate Decryption profiles for sites that you need to support but that don't support strong ciphers and algorithms. You can also create different Decryption profiles for different URL Categories to fine tune security vs. performance for traffic that contains no sensitive material; however, you should always decrypt and inspect all the traffic you can.

After you create a decryption profile, attach it to a decryption policy rule; the firewall then enforces the decryption profile settings on traffic that matches the decryption policy rule.

Palo Alto Networks firewalls include a default decryption profile that you can use to enforce the basic recommended protocol versions and cipher suites for decrypted traffic. However, the best practice is to enable tighter decryption controls as described in SSL Forward Proxy Decryption Profile, SSL Inbound Inspection Decryption Profile, and SSL Protocol Settings Decryption Profile.

Avoid supporting weak protocols or algorithms because they contain known vulnerabilities that attackers can exploit. If you must allow a weaker protocol or algorithm to support a key partner or contractor who uses legacy systems with weak protocols, create a separate Decryption profile for the exception and attach it to a Decryption policy rule that applies the profile only to the relevant traffic (for example, the source IP address of the partner). Don't allow the weak protocol for all traffic.

STEP 1 | Create a new decryption profile.

Select **Objects** > **Decryption Profile**, **Add** or modify a decryption profile rule, and give the rule a descriptive **Name**.

STEP 2 | (Optional) Allow the profile rule to be **Shared** across every virtual system on a firewall or every Panorama device group.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 3 | (Decryption Mirroring Only) Enable an Ethernet Interface for the firewall to use to copy and forward decrypted traffic.

Separate from this task, follow the steps to Configure Decryption Port Mirroring. Be aware of local privacy regulations that may prohibit mirroring or control the type of traffic that you can mirror. Decryption port mirroring requires a decryption port mirror license.

STEP 4 | (Optional) Block and control SSL tunneled and/or inbound traffic:

Although applying a Decryption profile to decrypted traffic is optional, it is a best practice to always apply a Decryption profile to the policy rules to protect your network against encrypted threats. You can't protect yourself against threats you can't see.

Select **SSL Decryption**:

- Select **SSL Forward Proxy** to configure the settings to verify certificates, enforce protocol versions and cipher suites, and perform failure checks on SSL decrypted traffic. These settings are active only when this profile is attached to a decryption policy rule configured to perform SSL Forward Proxy decryption.
- Select **SSL Inbound Inspection** to configure the settings to enforce protocol versions and cipher suites and to perform failure checks on inbound SSL traffic. These settings are active only when this profile is attached to a decryption policy rule that performs SSL Inbound Inspection.
- Select **SSL Protocol Settings** to configure the settings that control minimum and maximum protocol versions and key exchange, encryption, and authentication algorithms to enforce on decrypted SSL traffic. These settings are active when this profile is attached to decryption policy rules that are set to perform either SSL Forward Proxy decryption or SSL Inbound Inspection.

*If the firewall is in FIPS-CC mode and managed by a Panorama™ management server in standard mode, a decryption profile must be created locally on the firewall. Decryption profiles created on Panorama in standard mode contain references to **3DES** and **RC4** encryption algorithms and **MD5** authentication algorithm that are not supported and cause pushes to the managed firewall to fail.*

STEP 5 | (Optional) Block and control traffic (for example, a URL category) for which you choose to Create a Policy-Based Decryption Exclusion.

Although applying a Decryption profile to traffic that you choose not to decrypt is optional, it is a best practice to always apply a Decryption profile to the policy rules to protect your network against sessions with expired certificates or untrusted issuers.

Select **No Decryption** to configure the Profile for No Decryption and check the **Block sessions with expired certificates** and **Block sessions with untrusted issuers** boxes to validate certificates for traffic that is excluded from decryption. Create policy-based exclusions only for traffic that you choose not to decrypt. If a server breaks decryption for technical reasons, don't create a policy-based exclusion, add the server to the SSL Decryption Exclusion list (**Device** > **Certificate Management** > **SSL Decryption Exclusion**).

These setting are active only when the decryption profile is attached to a decryption policy rule that disables decryption for certain traffic.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 6 | (Optional) Block and control decrypted SSH traffic.

Select **SSH Proxy** to configure the SSH Proxy Decryption Profile and configure settings to enforce supported protocol versions and to block sessions if system resources are not available to perform decryption.

These settings are active only when the decryption profile is attached to a decryption policy rule that decrypts SSH traffic.

STEP 7 | Add the decryption profile when you Create a Decryption Policy Rule.

The firewall applies the decryption profile to and enforces the profile's settings on the traffic that matches the decryption policy rule.

STEP 8 | **Commit** the configuration.

Create a Decryption Policy Rule

Create a decryption policy rule to define traffic for the firewall to decrypt and the type of decryption you want the firewall to perform: SSL Forward Proxy, SSL Inbound Inspection, or SSH Proxy decryption. You can also use a decryption policy rule to define Decryption Mirroring.

STEP 1 | Add a new decryption policy rule.

Select **Policies** > **Decryption**, **Add** a new decryption policy rule, and give the policy rule a descriptive **Name**.

STEP 2 | Configure the decryption rule to match to traffic based on network and policy objects:

- **Firewall security zones**—Select **Source** and/or **Destination** and match to traffic based on the **Source Zone** and/or the **Destination Zone**.
- **IP addresses, address objects, and/or address groups**—Select **Source** and/or **Destination** to match to traffic based on **Source Address** and/or the **Destination Address**. Alternatively, select **Negate** to exclude the source address list from decryption.
- **Users**—Select **Source** and set the **Source User** for whom to decrypt traffic. You can decrypt specific user or group traffic, or decrypt traffic for certain types of users, such as unknown users or pre-logon users (users that are connected to GlobalProtect but are not yet logged in).
- **Ports and protocols**—Select **Service/URL Category** to set the rule to match to traffic based on service. By default, the policy rule is set to decrypt **Any** traffic on TCP and UDP ports.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

You can **Add** a service or a service group, and optionally set the rule to **application-default** to match to applications only on the application default ports.



The application-default setting can be useful when you Create a Policy-Based Decryption Exclusion. You can exclude applications running on their default ports from decryption, while continuing to decrypt the same applications when they are detected on non-standard ports.

- **URLs and URL categories**—Select Service/URL Category and decrypt traffic based on:
  - An externally-hosted list of URLs that the firewall retrieves for policy-enforcement (see **Objects** > **External Dynamic Lists**).
  - Palo Alto Networks predefined URL categories, which make it easy to decrypt entire categories of allowed traffic. This option is also useful when you create policy-based decryption exclusions because you can exclude sensitive sites by category instead of individually. For example, although you can create a custom URL category to group sites that you do not want to decrypt, you can also exclude financial or healthcare-related sites from decryption based on the predefined Palo Alto Networks URL categories. In addition, you can block risky URL Categories and create comfort pages to communicate the reason the sites are blocked or Enable Users to Opt Out of SSL Decryption.

    You can use the predefined high-risk and medium-risk URL categories to create a Decryption policy rule that decrypts all high-risk and medium-risk URL traffic. Place the rule at the bottom of the rulebase (all decryption exceptions must be above this rule so that you don't decrypt sensitive information) as a safety net to ensure that you decrypt and inspect all risky traffic. However, if high-risk or medium-risk sites to which you allow access contain personally identifiable information (PII) or other sensitive information that you don't want to decrypt, either block those sites to avoid allowing encrypted risky traffic while also avoiding privacy issues, or create a No Decryption rule to handle the sensitive traffic.
  - Custom URL categories (see **Objects** > **Custom Objects** > **URL Category**). For example, you can create a custom URL Category to specify a group of sites you need to access for business purposes but that don't support the safest protocols and algorithms, and then apply a customized Decryption profile to allow the looser protocols and algorithms for just those sites (that way, you don't decrease security by downgrading the Decryption profile you use for most sites).

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 3 | Set the rule to either decrypt matching traffic or to exclude matching traffic from decryption.

Select **Options** and set the policy rule **Action**:

To decrypt matching traffic:

1. Set the **Action** to **Decrypt** .
2. Set the **Type** of decryption for the firewall to perform on matching traffic:
   - SSL Forward Proxy
   - SSL Inbound Inspection. If you want to enable SSL Inbound Inspection, also select the **Certificate** for the destination internal server for the inbound SSL traffic.
   - SSH Proxy

To exclude matching traffic from decryption:

Set the **Action** to **No Decrypt**.

STEP 4 | (Optional) Select a **Decryption Profile** to perform additional checks on traffic that matches the policy rule.

Although applying a Decryption profile to decrypted traffic is optional, it is a best practice to always apply a Decryption profile to the policy rules to protect your network against encrypted threats. You can't protect yourself against threats you can't see.

For example, attach a decryption profile to a policy rule to ensure that server certificates are valid and to block sessions using unsupported protocols or ciphers. To Create a Decryption Profile, select **Objects** > **Decryption Profile**.

1. Create a decryption policy rule or open an existing rule to modify it.
2. Select **Options** and select a **Decryption Profile** to block and control various aspects of the traffic matched to the rule.

   The profile rule settings the firewall applies to matching traffic depends on the policy rule **Action** (Decrypt or No Decrypt) and the policy rule **Type** (SSL Forward Proxy, SSL Inbound Inspection, or SSH Proxy). This allows you to use the different decryption profiles with different types of decryption policy rules that apply to different types of traffic and users.
3. Click **OK**.

STEP 5 | Configure Decryption Logging (configure whether to log both successful and unsuccessful TLS handshakes and configure Decryption log forwarding).

STEP 6 | Click **OK** to save the policy.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 7 | Choose your next step to fully enable the firewall to decrypt traffic...

- ConfigureSSL Forward Proxy
- ConfigureSSL Inbound Inspection
- ConfigureSSH Proxy
- Create policy-based Decryption Exclusions for traffic you *choose* not to decrypt and add sites that break decryption for technical reasons such as pinned certificates or mutual authentication to the SSL Decryption Exclusion list.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Configure SSL Forward Proxy

To enable the firewall to perform SSL Forward Proxy decryption, you must set up the certificates required to establish the firewall as a trusted third party (proxy) to the session between the client and the server. The firewall can use certificates signed by an enterprise certificate authority (CA) or self-signed certificates generated on the firewall as *Forward Trust certificates* to authenticate the SSL session with the client.

- (Best Practice) **Enterprise CA-signed Certificates**—An enterprise CA can issue a signing certificate that the firewall can use to sign the certificates for sites which require SSL decryption. When the firewall trusts the CA that signed the certificate of the destination server, the firewall can send a copy of the destination server certificate to the client, signed by the enterprise CA. This is a best practice because usually all network devices already trust the Enterprise CA (it is usually already installed in the devices' CA Trust storage), so you don't need to deploy the certificate on the endpoints, so the rollout process is smoother.
- **Self-signed Certificates**—The firewall can act as a CA and generate self-signed certificates that the firewall can use to sign the certificates for sites which require SSL decryption. The firewall can sign a copy of the server certificate to present to the client and establish the SSL session. This method requires that you need to install the self-signed certificates on all of your network devices so that those devices recognize the firewall's self-signed certificates. Because the certificates must be deployed to all devices, this method is better for small deployments and proof-of-concept (POC) trials than for large deployments.

Additionally, set up a *Forward Untrust certificate* for the firewall to present to clients when the server certificate is signed by a CA that the firewall does not trust. This ensures that clients are prompted with a certificate warning when attempting to access sites with untrusted certificates.



Regardless of whether you generate Forward Trust certificates from your Enterprise Root CA or use a self-signed certificate generated on the firewall, generate a separate subordinate Forward Trust CA certificate for each firewall. The flexibility of using separate subordinate CAs enables you to revoke one certificate when you decommission a device (or device pair) without affecting the rest of the deployment and reduces the impact in any situation in which you need to revoke a certificate. Separate Forward Trust CAs on each firewall also helps troubleshoot issues because the CA error message the user sees includes information about the firewall the traffic is traversing. If you use the same Forward Trust CA on every firewall, you lose the granularity of that information.

After setting up the Forward Trust and Forward Untrust certificates required for SSL Forward Proxy decryption, create a Decryption policy rule to define the traffic you want the firewall to decrypt and create a Decryption profile to apply SSL controls and checks to the traffic. The Decryption policy decrypts SSL tunneled traffic that matches the rule into clear text traffic. The firewall blocks and restricts traffic based on the Decryption profile attached to the Decryption policy and on the firewall Security policy. The firewall re-encrypts traffic as it exits the firewall.



When you configure SSL Forward Proxy, the proxied traffic does not support DSCP code points or QoS.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 1 | Ensure that the appropriate interfaces are configured as either virtual wire, Layer 2, or Layer 3 interfaces.

View configured interfaces on the **Network** > **Interfaces** > **Ethernet** tab. The **Interface Type** column displays if an interface is configured to be a **Virtual Wire** or **Layer 2**, or **Layer 3** interface. You can select an interface to modify its configuration, including what type of interface it is.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 2 | Configure the Forward Trust certificate for the firewall to present to clients when a trusted CA has signed the server certificate. You can use an enterprise CA-signed certificate or a self-signed certificate as the forward trust certificate.

(Recommended Best Practice) Use an enterprise CA-signed certificate as the Forward Trust certificate. Create a uniquely named Forward Trust certificate on each firewall:

1. Generate a Certificate Signing Request (CSR) for the enterprise CA to sign and validate:
   1. Select **Device** > **Certificate Management** > **Certificates** and click **Generate**.
   2. Enter a **Certificate Name**. Use a unique name for each firewall.
   3. In the **Signed By** drop-down, select **External Authority (CSR)**.
   4. (Optional) If your enterprise CA requires it, add **Certificate Attributes** to further identify the firewall details, such as Country or Department.
   5. Click **Generate** to save the CSR. The pending certificate is now displayed on the **Device Certificates** tab.
2. Export the CSR:
   1. Select the pending certificate displayed on the **Device Certificates** tab.
   2. Click **Export** to download and save the certificate file.

      *Leave **Export private key** unselected in order to ensure that the private key remains securely on the firewall.*
   3. Click **OK**.
3. Provide the certificate file to your enterprise CA. When you receive the enterprise CA-signed certificate from your enterprise CA, save the enterprise CA-signed certificate to import onto the firewall.
4. Import the enterprise CA-signed certificate onto the firewall:
   1. Select **Device** > **Certificate Management** > **Certificates** and click **Import**.
   2. Enter the pending **Certificate Name** exactly. The **Certificate Name** that you enter must exactly match the pending certificate name in order for the pending certificate to be validated.
   3. Select the signed **Certificate File** that you received from your enterprise CA.
   4. Click **OK**. The certificate is displayed as valid with the Key and CA check boxes selected.
5. Select the validated certificate to enable it as a **Forward Trust Certificate** to be used for SSL Forward Proxy decryption.
6. Click **OK** to save the enterprise CA-signed forward trust certificate.

Use a self-signed certificate as the Forward Trust certificate:

1. Create a self-signed Root CA certificate.
2. Click the self-signed root CA certificate (**Device** > **Certificate Management** > **Certificates** > **Device Certificates**) to open **Certificate information** and then click the **Trusted Root CA** checkbox.
3. Click **OK**.
4. Generate new subordinate CA certificates for each firewall:
   1. Select **Device** > **Certificate Management** > **Certificates**.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

2. Click **Generate** at the bottom of the window.
3. Enter a **Certificate Name**.
4. Enter a **Common Name**, such as 192.168.2.1. This should be the IP or FQDN that will appear in the certificate. In this case, we are using the IP of the trust interface. Avoid using spaces in this field.
5. In the **Signed By** field, select the self-signed Root CA certificate that you created.
6. Click the **Certificate Authority** check box to enable the firewall to issue the certificate. Selecting this check box creates a certificate authority (CA) on the firewall that is imported to the client browsers, so clients trust the firewall as a CA.
7. **Generate** the certificate.

5. Click the new certificate to modify it and click the **Forward Trust Certificate** checkbox to configure the certificate as the Forward Trust Certificate.
6. Click **OK** to save the self-signed forward trust certificate.
7. Repeat this procedure to generate a unique subordinate CA certificate on each firewall.

STEP 3 | Distribute the forward trust certificate to client system certificate stores.

If you are using an enterprise-CA signed certificate as the forward trust certificate for SSL Forward Proxy decryption, and the client systems already have the enterprise CA installed in the local trusted root CA list, you can skip this step. (The client systems trust the subordinate CA certificates you generate on the firewall because the Enterprise Trusted Root CA has signed them.)

If you do not install the forward trust certificate on client systems, users see certificate warnings for each SSL site they visit.

On a firewall configured as a GlobalProtect portal:

This option is supported with Windows and Mac client OS versions, and requires GlobalProtect agent 3.0.0 or later to be installed on the client systems.

1. Select **Network** > **GlobalProtect** > **Portals** and then select an existing portal configuration or **Add** a new one.
2. Select **Agent** and then select an existing agent configuration or **Add** a new one.
3. **Add** the self-signed firewall Trusted Root CA certificate to the Trusted Root CA section. After GlobalProtect distributes the firewall's Trusted Root CA certificate to client

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

systems, the client systems trust the firewall's subordinate CA certificates because the clients trust the firewall's Root CA certificate.

4. **Install in Local Root Certificate Store** so that the GlobalProtect portal automatically distributes the certificate and installs it in the certificate store on GlobalProtect client systems.
5. Click **OK** twice.

Without GlobalProtect:

Export the firewall Trusted Root CA certificate so that you can import it into client systems. Highlight the certificate and click **Export** at the bottom of the window. Choose PEM format.

*Do not select the **Export private key** checkbox! The private key should remain on the firewall and should not be exported to client systems.*

Import the firewall's Trusted Root CA certificate into the browser Trusted Root CA list on the client systems in order for the clients to trust it. When importing into the client browser, ensure that you add the certificate to the Trusted Root Certification Authorities certificate store. On Windows systems, the default import location is the Personal certificate store. You can also simplify this process by using a centralized deployment option, such as an Active Directory Group Policy Object (GPO).

STEP 4 | Configure the Forward Untrust certificate (use the same Forward Untrust certificate for all firewalls).

1. Click **Generate** at the bottom of the certificates page.
2. Enter a **Certificate Name**, such as my-ssl-fwd-untrust.
3. Set the **Common Name**, for example 192.168.2.1. Leave **Signed By** blank.
4. Click the **Certificate Authority** check box to enable the firewall to issue the certificate.
5. Click **Generate** to generate the certificate.
6. Click **OK** to save.
7. Click the new my-ssl-fwd-untrust certificate to modify it and enable the **Forward Untrust Certificate** option.

   Do not export the Forward Untrust certificate to the Certificate Trust Lists of your network devices! Do not install the Forward Untrust certificate on client systems. This is critical because installing the Untrust certificate in the Trust List results in devices trusting websites that the firewall does not trust. In addition, users won't see certificate warnings for untrusted sites, so they won't know the sites are untrusted and may access those sites, which could expose your network to threats.
8. Click **OK** to save.

STEP 5 | (Optional) Configure the Key Size for SSL Forward Proxy Server Certificates that the firewall presents to clients. By default, the firewall determines the key size to use based on the key size of the destination server certificate.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 6 | Create a Decryption Policy Rule to define traffic for the firewall to decrypt and Create a Decryption Profile to apply SSL controls to the traffic.

Although Decryption profiles are optional, it is a best practice to include a Decryption profile with each Decryption policy rule to prevent weak, vulnerable protocols and algorithms from allowing questionable traffic on your network.

1. Select **Policies** > **Decryption**, Add or modify an existing rule, and define traffic to be decrypted.
2. Select **Options** and:
   - Set the rule **Action** to **Decrypt** matching traffic.
   - Set the rule **Type** to **SSL Forward Proxy**.
   - (Optional but a best practice) Configure or select an existing **Decryption Profile** to block and control various aspects of the decrypted traffic (for example, create a decryption profile to perform certificate checks and enforce strong cipher suites and protocol versions).
3. Click **OK** to save.

STEP 7 | Enable the firewall to forward decrypted SSL traffic for WildFire analysis.

This option requires an active WildFire license and is a WildFire best practice.

STEP 8 | **Commit** the configuration.

STEP 9 | Choose your next step:

- Enable Users to Opt Out of SSL Decryption.
- Configure Decryption Exclusions to disable decryption for certain types of traffic.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Configure SSL Inbound Inspection

Use SSL Inbound Inspection to decrypt and inspect inbound SSL traffic destined for a network server (you can perform SSL Inbound Inspection for any server if you load the server certificate onto the firewall). With an SSL Inbound Inspection Decryption policy enabled, the firewall decrypts all SSL traffic identified by the policy to clear text traffic and inspects it. The firewall blocks, restricts, or allows the traffic based on the Decryption profile attached to the policy and the Security policy that applies to the traffic, including and any configured Antivirus, Vulnerability Protection, Anti-Spyware, URL Filtering, and File Blocking profiles. As a best practice, enable the firewall to forward decrypted SSL traffic for WildFire analysis and signature generation.

Configuring SSL Inbound Inspection includes installing the targeted server certificate on the firewall, creating an SSL Inbound Inspection Decryption policy, and applying a Decryption profile to the policy.

When you configure SSL Inbound Inspection, the proxied traffic does not support DSCP code points or QoS.

SSL Inbound Inspection does not support Authentication Portal redirect. To use Authentication Portal redirect and decryption, you must use SSL Forward Proxy.

STEP 1 | Ensure that the appropriate interfaces are configured as either Tap, Virtual Wire, Layer 2, or Layer 3 interfaces.

You cannot use a Tap mode interface for SSL Inbound Inspection if the negotiated ciphers include PFS key exchange algorithms (DHE and ECDHE).

View configured interfaces on the **Network** > **Interfaces** > **Ethernet** tab. The **Interface Type** column displays if an interface is configured to be a **Virtual Wire** or **Layer 2**, or **Layer 3** interface. You can select an interface to modify its configuration, including the interface type.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 2 | Ensure that the targeted server certificate is installed on the firewall.

If you have enabled SSL Inbound Inspection using PFS key exchange algorithms, you must upload a certificate bundle (a single file) to the firewall with your certificates arranged as follows:

1. *End-entity (leaf) certificate*
2. *Intermediate certificates (in issuing order)*
3. *(Optional) Root certificate*

Uploading the file ensures that clients receive the complete certificate chain during SSL/TLS handshakes, avoiding client-side server certificate authentication issues.

On the web interface, select **Device** > **Certificate Management** > **Certificates** > **Device Certificates** to view certificates installed on the firewall.

To import the targeted server certificate onto the firewall:

1. On the **Device Certificates** tab, select **Import**.
2. Enter a descriptive **Certificate Name**.
3. Browse for and select the targeted server **Certificate File**.
4. Click **OK**.

STEP 3 | Create a Decryption Policy Rule to define traffic for the firewall to decrypt and Create a Decryption Profile to apply SSL controls to the traffic.

Although Decryption profiles are optional, it is a best practice to include a Decryption profile with each Decryption policy rule to prevent weak, vulnerable protocols and algorithms from allowing questionable traffic on your network.

1. Select **Policies** > **Decryption**, **Add** or modify an existing rule, and define traffic to be decrypted.
2. Select **Options** and:
   - Set the rule **Action** to **Decrypt** matching traffic.
   - Set the rule **Type** to **SSL Inbound Inspection**.
   - Select the **Certificate** for the internal server that is the destination of the inbound SSL traffic.
   - (Optional but a best practice) Configure or select an existing **Decryption Profile** to block and control various aspects of the decrypted traffic (for example, create

https://t.me/learningnets

a Decryption profile to terminate sessions with unsupported algorithms and unsupported cipher suites).

When you configure the SSL Protocol Settings Decryption Profile for SSL Inbound Inspection traffic, create separate profiles for servers with different security capabilities. For example, if one set of servers supports only RSA, the SSL Protocol Settings only need to support RSA. However, the SSL Protocol Settings for servers that support PFS should support PFS. Configure SSL Protocol Settings for the highest level of security that the server supports, but check performance to ensure that the firewall resources can handle the higher processing load that higher security protocols and algorithms require.

3. Click **OK** to save.

STEP 4 | Enable the firewall to forward decrypted SSL traffic for WildFire analysis.

This option requires an active WildFire license and is a WildFire best practice.

STEP 5 | **Commit** the configuration.

STEP 6 | Choose your next step...

- Enable Users to Opt Out of SSL Decryption.
- Configure Decryption Exclusions to disable decryption for certain types of traffic.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Configure SSH Proxy

Configuring SSH Proxy does not require certificates and the key used to decrypt SSH sessions is generated automatically on the firewall during boot up. With SSH decryption enabled, the firewall decrypts SSH traffic and blocks and or restricts the SSH traffic based on your decryption policy and decryption profile settings. Traffic is re-encrypted as it exits the firewall.

When you configure SSH Proxy, the proxied traffic does not support DSCP code points or QoS.

STEP 1 | Ensure that the appropriate interfaces are configured as either virtual wire, Layer 2, or Layer 3 interfaces. Decryption can only be performed on virtual wire, Layer 2, or Layer 3 interfaces.

View configured interfaces on the **Network** > **Interfaces** > **Ethernet** tab. The **Interface Type** column displays if an interface is configured to be a **Virtual Wire** or **Layer 2**, or **Layer 3** interface. You can select an interface to modify its configuration, including what type of interface it is.

STEP 2 | Create a Decryption Policy Rule to define traffic for the firewall to decrypt and Create a Decryption Profile to apply checks to the SSH traffic.

Although Decryption profiles are optional, it is a best practice to include a Decryption profile with each Decryption policy rule to prevent weak, vulnerable protocols and algorithms from allowing questionable traffic on your network.

1. Select **Policies** > **Decryption**, Add or modify an existing rule, and define traffic to be decrypted.
2. Select **Options** and:
   - Set the rule **Action** to **Decrypt** matching traffic.
   - Set the rule **Type** to **SSH Proxy**.
   - (Optional but a best practice) Configure or select an existing **Decryption Profile** to block and control various aspects of the decrypted traffic (for example, create a Decryption profile to terminate sessions with unsupported versions and unsupported algorithms).
3. Click **OK** to save.

STEP 3 | **Commit** the configuration.

STEP 4 | (Optional) Continue to Decryption Exclusions to disable decryption for certain types of traffic.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Configure Server Certificate Verification for Undecrypted Traffic

You create no-decryption policies for traffic that you *choose* not to decrypt because the traffic is personal, sensitive, or subject to local laws and regulations. For example, you may choose not to decrypt the traffic of certain executives or traffic between finance users and finance servers that contain personal information. (Don't exclude traffic that you can't decrypt because a site breaks decryption for technical reasons such as a pinned certificate or mutual authentication by policy. Instead, add the hostname to the Decryption Exclusion List.)

However, just because you don't decrypt the traffic doesn't mean you should let any and all undecrypted traffic on your network. It is a best practice to apply a No Decryption profile to undecrypted traffic to block sessions with expired certificates and untrusted issuers.

STEP 1 | Create a Decryption Policy Rule to identify the undecrypted traffic and Create a Decryption Profile to block bad sessions.

1. Select **Policies** > **Decryption** and Add or modify an existing rule to identify the undecrypted traffic.
2. Select **Options** and:
   - Set the rule **Action** to **No Decrypt** so that the firewall doesn't decrypt traffic that matches the rule.
   - Ignore the rule **Type** because the traffic is not decrypted.
   - (Optional but a best practice) Configure or select an existing Decryption profile for undecrypted traffic to block sessions with expired certificates and untrusted certificate issuers.



Do not attach a No Decryption profile to Decryption policies for TLSv1.3 traffic that you don't decrypt because the firewall can't read the encrypted certificate information so it can't perform certificate checks. However, you should still create a Decryption policy for TLSv1.3 traffic that you don't decrypt because undecrypted traffic isn't logged unless a Decryption policy controls that traffic.

STEP 2 | **Commit** the configuration.

STEP 3 | Choose your next step:

- Enable Users to Opt Out of SSL Decryption.
- Configure Decryption Exclusions to disable decryption for certain types of traffic.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Decryption Exclusions

You can exclude two types of traffic from decryption:

- Traffic that breaks decryption for *technical reasons*, such as using a pinned certificate, an incomplete certificate chain, unsupported ciphers, or mutual authentication (attempting to decrypt the traffic results in blocking the traffic). Palo Alto Networks provides a predefined SSL Decryption Exclusion list (**Device** > **Certificate Management** > **SSL Decryption Exclusion**) that excludes hosts with applications and services that are known to break decryption technically from SSL Decryption by default. If you encounter sites that break decryption technically and are not on the SSL Decryption Exclusion list, you can add them to list manually by server hostname. The firewall blocks sites whose applications and services break decryption technically unless you add them to the SSL Decryption Exclusion list.

  If the Decryption profile allows **Unsupported Modes** (sessions with client authentication, unsupported versions, or unsupported cipher suites), the firewall automatically adds servers and applications that use the allowed unsupported modes to the its Local SSL Decryption Exclusion Cache (**Device** > **Certificate Management** > **SSL Decryption Exclusion** > **Show Local Exclusion Cache**). When you block unsupported modes, you increase security but you also block communication with applications that use those modes.
- Traffic that you *choose* not to decrypt because of business, regulatory, personal, or other reasons, such as financial-services, health-and-medicine, or government traffic. You can choose to exclude traffic based on source, destination, URL category, and service.

You can use asterisks (*) as wildcards to create decryption exclusions for multiple hostnames associated with a domain. Asterisks behave the same way that carets (^) behave for URL category exceptions—each asterisk controls one variable subdomain (label) in the hostname. This enables you to create both very specific and very general exclusions. For example:

- mail.*.com matches mail.company.com but does not match mail.company.sso.com.
- *.company.com matches tools.company.com but does not match eng.tools.company.com.
- *.*.company.com matches eng.tools.company.com but does not match eng.company.com.
- *.*.*.company.com matches corp.exec.mail.company.com, but does not match corp.mail.company.com.
- mail.google.* matches mail.google.com, but does not match mail.google.uk.com.
- mail.google.*.* matches mail.google.co.uk, but does not match mail.google.com.

For example, to use wildcards to exclude video-stats.video.google.com from decryption but not to exclude video.google.com from decryption, exclude *.*.google.com.

*Regardless of the number of asterisk wildcards that precede a hostname (without a non-wildcard label preceding the hostname), the hostname matches the entry. For example, *.google.com, *.*.google.com, and *.*.*.google.com all match google.com. However, *.dev.*.google.com does not match google.com because one label (dev) is not a wildcard.*

To increase visibility into traffic and reduce the attack surface as much as possible, don't make decryption exceptions unless you must.

- Palo Alto Networks Predefined Decryption Exclusions

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

- Exclude a Server from Decryption for Technical Reasons
- Local Decryption Exclusion Cache
- Create a Policy-Based Decryption Exclusion

Palo Alto Networks Predefined Decryption Exclusions

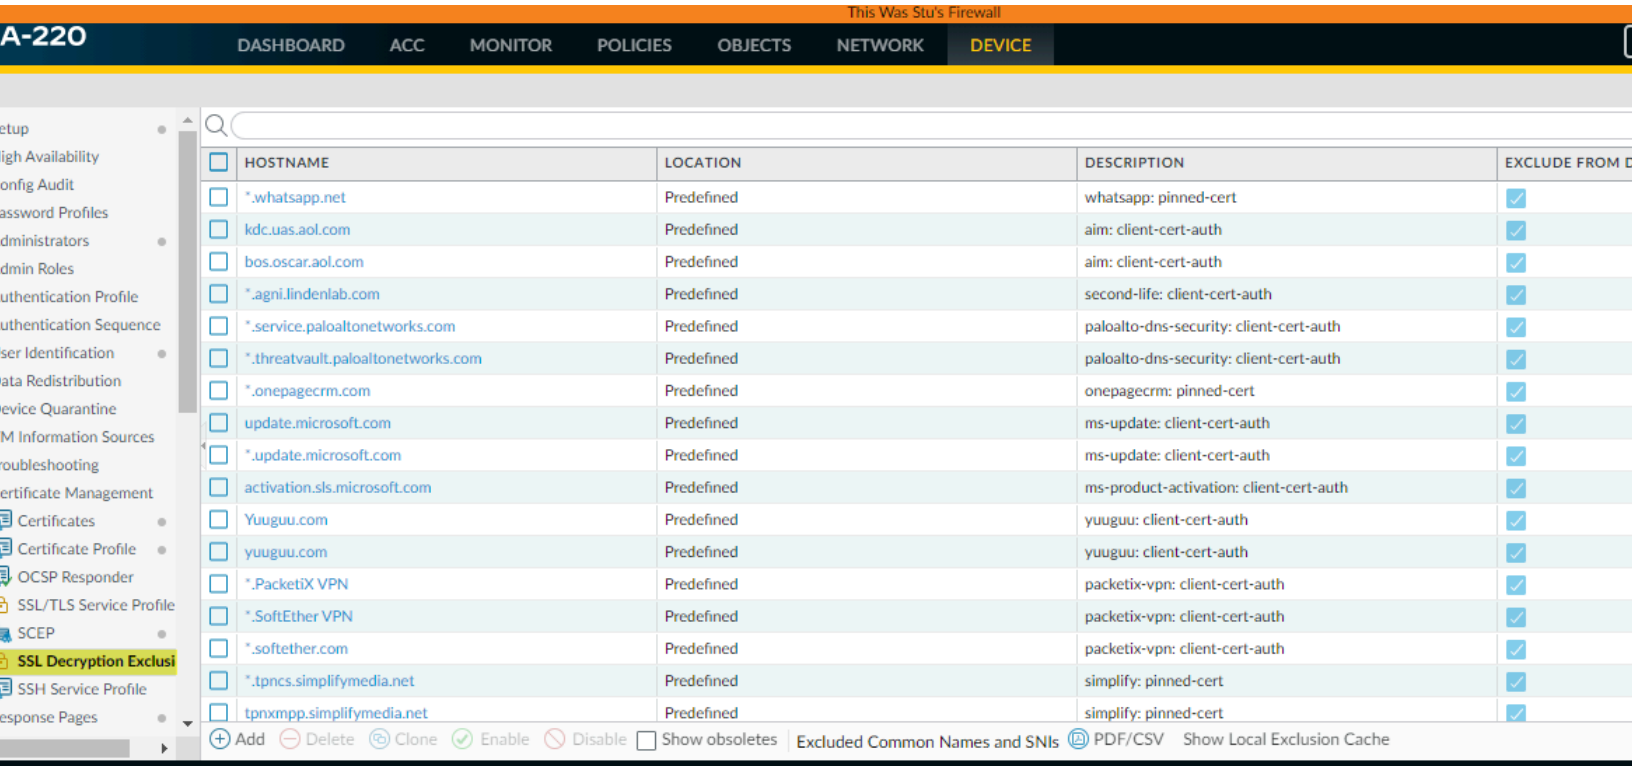
The firewall provides a predefined SSL Decryption Exclusion list to exclude from decryption commonly used sites that break decryption because of technical reasons such as pinned certificates and mutual authentication. The predefined decryption exclusions are enabled by default and Palo Alto Networks delivers new and updated predefined decryption exclusions to the firewall as part of the Applications and Threats content update (or the Applications content update, if you do not have a Threat Prevention license). The firewall does not decrypt traffic that matches predefined exclusions and allows the encrypted traffic based on the Security policy that governs that traffic. However, the firewall can't inspect the encrypted traffic or enforce Security policy on it.

*The SSL Decryption Exclusion list is **not** for sites that you choose not to decrypt for legal, regulatory, business, privacy, or other volitional reasons, it is only for sites that break decryption technically (decrypting these sites blocks their traffic). For traffic such as IP addresses, users, URL categories, services, and even entire zones that you choose not to decrypt,* Create a Policy-Based Decryption Exclusion.

Because the traffic of sites on the SSL Decryption Exclusion list remains encrypted, the firewall does not inspect or provide further security enforcement the traffic. You can disable a predefined exclusion. For example, you may choose to disable predefined exclusions to enforce a strict security policy that allows only applications and services that the firewall can inspect and on which the firewall can enforce Security policy. However, the firewall blocks sites whose applications and services break decryption technically if they are not enabled on the SSL Decryption Exclusion list.

You can view and manage all Palo Alto Networks predefined SSL decryption exclusions directly on the firewall (**Device** > **Certificate Management** > **SSL Decryption Exclusions**).

https://t.me/learningnets



The **Hostname** displays the name of the host that houses the application or service that breaks decryption technically. You can also **Add** hosts to Exclude a Server from Decryption for Technical Reasons if it is not on the predfined list.

The **Description** displays the reason the firewall can't decrypt the site's traffic, for example, **pinned-cert** (a pinned certificate) or **client-cert-auth** (client authentication).

The firewall automatically removes enabled predefined SSL decryption exclusions from the list when they become obsolete (the firewall removes an application that decryption previously caused to break when the application becomes supported with decryption). **Show Obsoletes** checks if any disabled predefined exclusions remain on the list and are no longer needed. The firewall does not remove disabled predefined decryption exclusions from the list automatically, but you can select and **Delete** obsolete entries.

You can select a hostname's checkbox and then click **Disable** to remove predefined sites from the list. Use the SSL Decryption Exclusion list only for sites that break decryption for technical reasons, don't use it for sites that you choose not to decrypt.

Exclude a Server from Decryption for Technical Reasons

If decryption breaks an important application or service technically (decrypting the traffic blocks it), you can add the hostname of the site that hosts to the application or service to the Palo Alto Networks predefined SSL Decryption Exclusion list to create a custom decryption exception. The firewall doesn't decrypt, inspect, and enforce Security policy on traffic that the SSL Decryption Exclusion list allows because the traffic remains encrypted, so be sure that the sites you add to the list really are sites with applications or services you need for business. For example, some business-critical internal custom applications may break decryption and you can add them to the list so that the firewall allows the encrypted custom application traffic.

https://t.me/learningnets

*The SSL Decryption Exclusion list is **not** for sites that you choose not to decrypt for legal, regulatory, business, privacy, or other volitional reasons, it is only for sites that break decryption technically. For traffic (IP addresses, users, URL categories, services, and even entire zones) that you choose not to decrypt,* Create a Policy-Based Decryption Exclusion.

Reasons that sites break decryption technically include pinned certificates, client authentication, incomplete certificate chains, and unsupported ciphers. For HTTP public key pinning (HPKP), most browsers that use HPKP permit Forward Proxy decryption as long as you install the enterprise CA certificate (or the certificate chain) on the client.

If the technical reason for excluding a site from decryption is an incomplete certificate chain, the next-generation firewall doesn't automatically fix the chain as a browser would. If you need to add a site to the SSL Decryption Exclusion list, manually review the site to ensure it's a legitimate business site, then download the missing sub-CA certificates and load and deploy *them onto the firewall.*

After you add a server to the SSL Decryption Exclusion list, the firewall compares the server hostname that you use to define the decryption exclusion against both the Server Name Indication (SNI) in the client hello message and the Common Name (CN) in the server certificate. If either the SNI or CN match the entry in the SSL Decryption Exclusion list, the firewall excludes the traffic from decryption.

STEP 1 | Select **Device** > **Certificate Management** > **SSL Decryption Exclusions**.

STEP 2 | **Add** a new decryption exclusion, or select an existing custom entry to modify it.

STEP 3 | Enter the **hostname** of the website or application you want to exclude from decryption.

The hostname is case-sensitive.

You can use wildcards to exclude multiple hostnames associated with a domain. The firewall excludes all sessions where the server presents a CN that matches the domain from decryption.

Make sure that the hostname field is unique for each custom entry. If a predefined exclusion matches a custom entry, the custom entry takes precedence.

STEP 4 | (Optional) Select **Shared** to share the exclusion across all virtual systems in a multiple virtual system firewall.

STEP 5 | **Exclude** the application from decryption. Alternatively, if you are modifying an existing decryption exclusion, you can clear this checkbox to start decrypting an entry that was previously excluded from decryption.

STEP 6 | Click **OK** to save the new exclusion entry.

Local Decryption Exclusion Cache

The firewall can add servers to the Local Decryption Exclusion cache (**Device** > **Certificate Management** > **SSL Decryption Exclusion** > **Show Local Exclusion Cache**) and exclude their traffic

https://t.me/learningnets

from decryption automatically for 12 hours if that traffic breaks decryption for technical reasons such as a pinned certificate or an unsupported certificate. When the Decryption profile allows unsupported modes—sessions with client authentication, unsupported versions, or unsupported cipher suites—and the allowed traffic uses an unsupported mode, then the device automatically adds the server to the local exclusion cache and bypasses decryption. The firewall doesn't decrypt, inspect, and enforce Security policy on traffic that the Local Decryption Exclusion cache allows because the traffic remains encrypted. Ensure that the sites you exclude from decryption (by applying a Decryption profile that allows unsupported modes) are sites with applications or services you need for business.

Blocking unsupported modes blocks communication with applications that use those modes to increase security. Client authentication is a common reason for excluding applications from decryption, which is why the best practice is to block unsupported versions and unsupported ciphers and to allow client authentication in the Decryption profile. If the Decryption profile allows client authentication, then when a client starts a session with a server that requires the client to authenticate, instead of blocking the traffic because the firewall can't decrypt it, the firewall adds the application and server to the local exclusion cache and allows the traffic.



If you allow traffic from sites that use client authentication and are not in the predefined sites on the SSL Decryption Exclusion list, create a Decryption profile that allows sessions with client authentication. Add the profile to a Decryption policy rule that applies only to the server(s) that host the application. To increase security even more, you can require Multi-Factor Authentication to complete the user login process. Alternatively, you can add the site to the SSL Decryption Exclusion list to skip decryption without using an explicit Decryption policy.

The firewall adds Local SSL Decryption Exclusion cache entries based on the Decryption policy and profile that controls the application traffic. If you don't block **Unsupported Mode Checks** in the Decryption profile, the firewall adds entries to the Local SSL Decryption Exclusion cache when:

- The client supports only TLSv1.2 and the server supports only TLSv1.3. In the local cache, the Reason shown for this exclusion is SSL_UNSUPPORTED.
- The client supports TLSv1.3 and TLSv1.2, and the server supports only TLSv1.2. In this case, the **Reason** column shows TLS13_UNSUPPORTED.

  *When the **Reason** for adding a server to the Local SSL Decryption Exclusion cache is TLS13_UNSUPPORTED, the firewall downgrades the protocol to TLSv1.2 and the firewall decrypts and inspects the traffic.*
- The client advertises a specific cipher that the server doesn't support.
- The client advertises a specific curve that the server doesn't support.

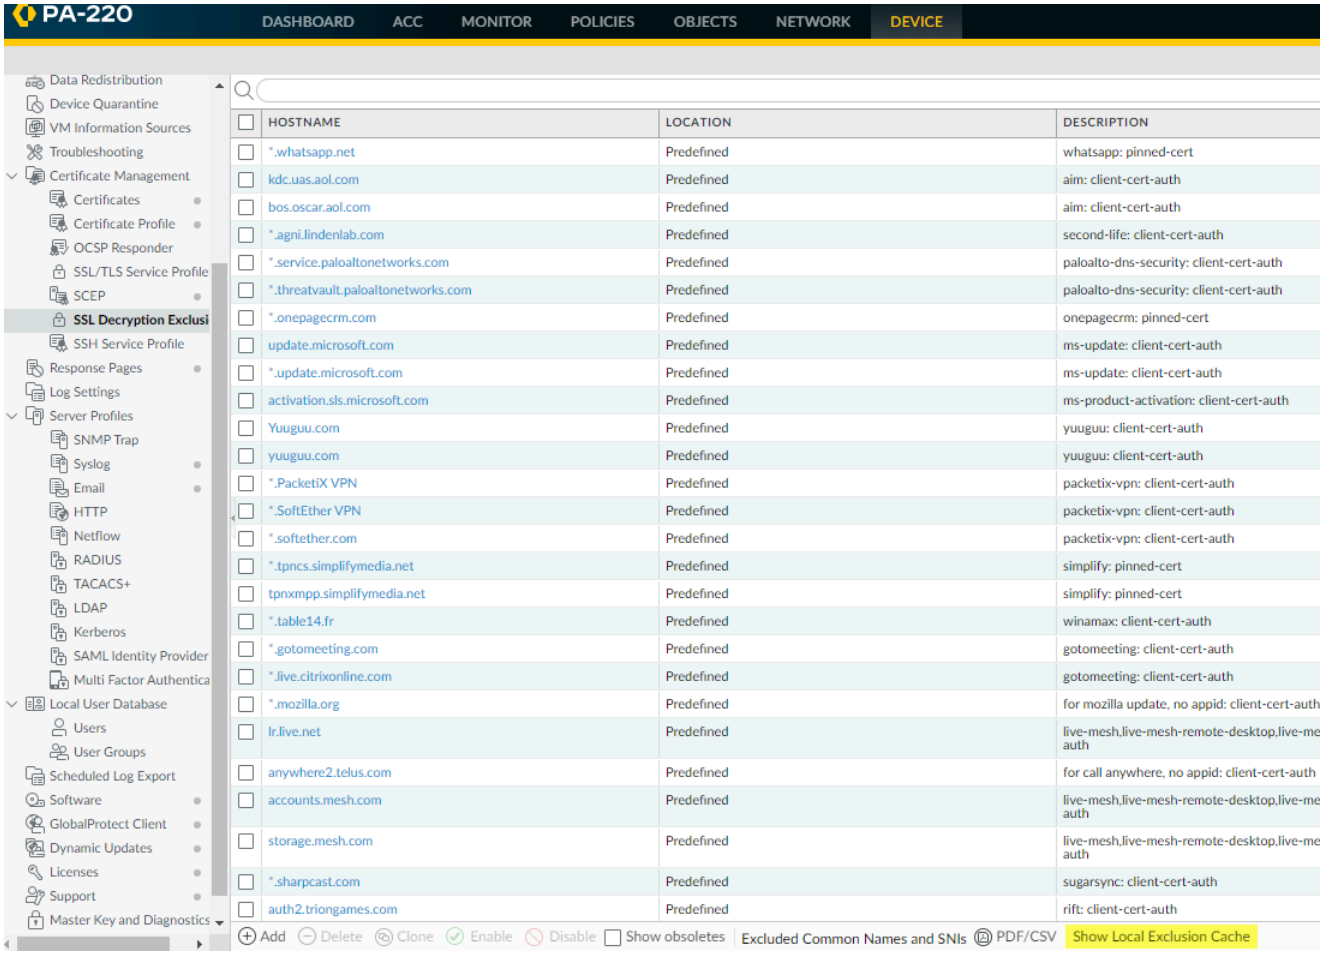
The local cache contains a maximum of 1,024 entries. You can't add local exclusions to the Local SSL Decryption Exclusion cache manually (but you can add decryption exclusions to the SSL Decryption Exclusion list manually).

You must have superuser or Certificate Management administrative access to view the Local SSL Decryption Exclusion cache. To view it, navigate to **Device** > **Certificate Management** > **SSL Decryption Exclusion** and then click **Show Local Exclusion Cache** near the bottom of the screen. The local exclusion cache displays the application, the server, the reason for inclusion in the cache,

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

the Decryption profile that controls the traffic, and more for each entry. You can select and delete entries from the local cache manually.



You can also delete cached entries using the CLI:

```
clear ssl-decrypt exclude-cache [server <value>] [application
 <value>]
```

If anyone attempts to access the same server before the local cache entry ages out (12 hours), the firewall matches the session to the cache entry, bypasses decryption, and allows the traffic. The firewall flushes the local exclusion cache if you change the Decryption policy or profile because those changes may affect the classification of the session. If the cache becomes full, the firewall purges the oldest entries as new entries arrive.

Create a Policy-Based Decryption Exclusion

Policy-based decryption exclusions are for excluding traffic that you *choose* not to decrypt. You can create a policy-based decryption exclusion based on any combination of the traffic's source, destination, service, or URL Category. Examples of traffic you may choose not to decrypt include:

- Traffic that you should never decrypt because it contains personally identifiable information (PII) or other sensitive information, such as the URL Filtering categories financial-services, health-and-medicine, and government.

https://t.me/learningnets

- Traffic that originates or is destined for executives or other users whose traffic shouldn't be decrypted.
- Some devices such as finance servers may need to be excepted from decryption.
- Depending on the business, some companies may value privacy and the user experience more than security for some applications.
- Laws or local regulations that prohibit decryption of some traffic.

  An example of not decrypting traffic for regulatory and legal compliance is the European Union (EU) General Data Protection Regulation (GDPR). The EU GDPR will require strong protection of all personal data for all individuals. The GDPR affects all companies, including foreign companies, that collect or process the personal data of EU residents.

  Different regulations and compliance rules may mean that you treat the same data differently in different countries or regions. Businesses usually can decrypt personal information in their corporate data centers because the business owns the information. The best practice is to decrypt as much traffic as possible so that you can see it and apply security protection to it.

You can use predefined URL Categories to except entire categories of websites from decryption, you can create custom URL Categories to define a customized list of URLs that you don't want to decrypt, or you can create an External Dynamic List (EDL) to define a customized list of URLs that you don't want to decrypt.

In environments such as Office 365 that have dynamically changing IP addresses or in environments where you make frequent changes to the list of URLs that you want to exclude from decryption, it's often preferable to use an EDL instead of a URL Category to specify the excluded URLs. Using an EDL is less disruptive in dynamic environments because editing an EDL changes the URL categories dynamically, without a **Commit**, while editing a custom URL Category requires a **Commit** to take effect.

Create an EDL or a custom URL Category that contains all the categories you choose not to decrypt so that one Decryption policy rule governs the encrypted traffic you choose to allow. Apply a No Decryption profile to the rule. The ability to add categories to an EDL or a custom URL Category makes it easy to exclude traffic from decryption and helps keep the rulebase clean.

Similar to Security policy rules, the firewall compares incoming traffic to Decryption policy rules in the policy rulebase's sequence. Place Decryption exclusion rules at the top of the rulebase to prevent inadvertently decrypting sensitive traffic or traffic that laws and regulations prevent you from decrypting.

If you create policy-based decryption exclusions, the best practice is to place the following exclusion rules at the top of the decryption rulebase, in the following order:

1. IP-address based exceptions for sensitive destination servers.
2. Source-user based exceptions for executives and other users or groups.
3. Custom URL or EDL based exceptions for destination URLs.
4. Sensitive predefined URL Category based exceptions for destination URLs of entire categories such as financial-services, health-and-medicine, and government.

Place rules that decrypt traffic after these rules in the decryption rulebase.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

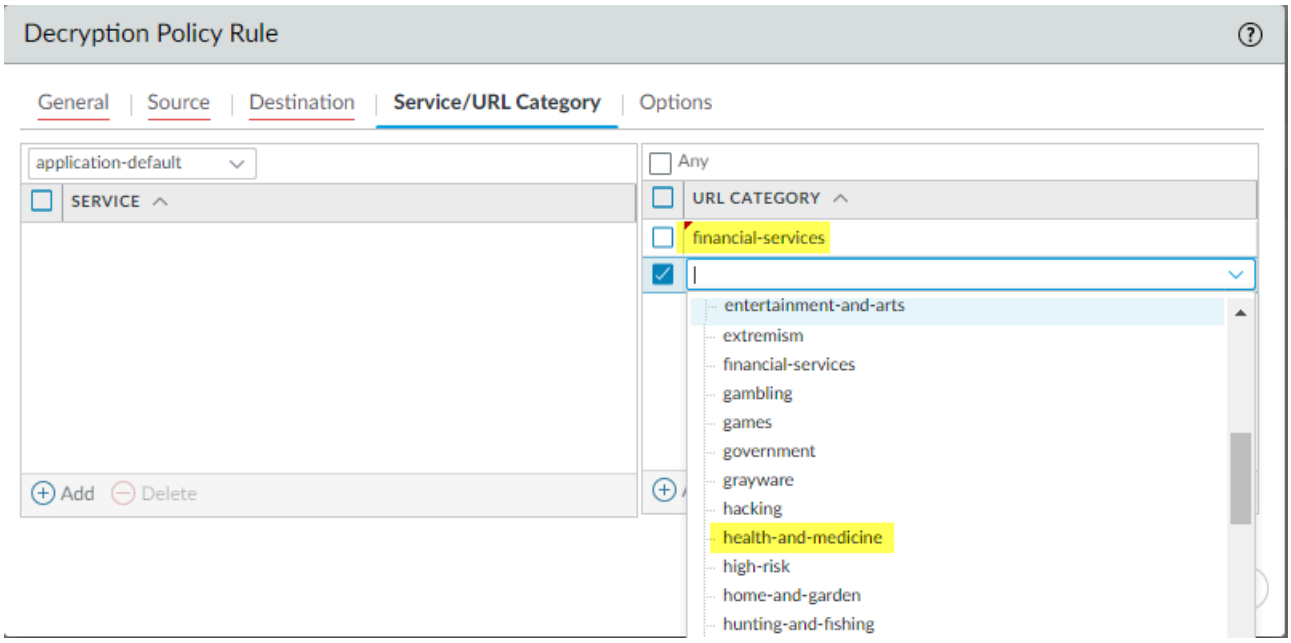
STEP 1 | Exclude traffic from decryption based on match criteria.

This example shows how to exclude traffic categorized as financial or health-related from SSL Forward Proxy decryption.

1. Select **Policies** > **Decryption** and **Add** or modify a decryption policy rule.
2. Define the traffic that you want to exclude from decryption.

   In this example:

   1. Give the rule a descriptive **Name**, such as No-Decrypt-Finance-Health.
   2. Set the **Source** and **Destination** to **Any** to apply the No-Decrypt-Finance-Health rule to all SSL traffic destined for an external server.
   3. Select **URL Category** and **Add** the URL categories financial-services and health-and-medicine.



3. Select **Options** and set the rule to **No Decrypt**.
4. (Optional but a best practice) Create and attach a No Decryption profile to the rule to validate certificates for sessions the firewall does not decrypt. Configure the profile to **Block sessions with expired certificates** and **Block sessions with untrusted issuers**.

   Exception: Do not attach a No Decryption profile to Decryption policies for TLSv1.3 traffic that you don't decrypt because the firewall can't read the encrypted certificate information so it can't perform certificate checks. However, you should still create a Decryption policy for TLSv1.3 traffic that you don't decrypt because undecrypted traffic isn't logged unless a Decryption policy controls that traffic.

5. Click **OK** to save the No-Decrypt-Finance-Health decryption rule.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 2 | Place the decryption exclusion rule at the top of your decryption policy rulebase.

The firewall enforces decryption rules against incoming traffic in the rulebase sequence and enforces the first rule that match the traffic.

Select the **No-Decrypt-Finance-Health** policy (**Decryption** > **Policies**), and click **Move Up** until it appears at the top of the list, or drag and drop the rule.

STEP 3 | Save the configuration.

Click **Commit**.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Block Private Key Export

You can permanently block the export of private keys for certificates when you generate them in or import them into PAN-OS or Panorama. Blocking the export of private keys from your PAN-OS devices hardens your security posture because it prevents rogue administrators or other bad actors from misusing keys. Administrators with roles that include certificate management can block the export of private keys. You can't block keys that already exist on a device; you can only block keys at the time that you generate them in or import them into PAN-OS.

When an administrator blocks the export of a private key, no administrator can export that key, not even Superuser administrators. If you need to export a private key from a PAN-OS appliance, regenerate the certificate and the key without selecting the option to block private key export.

To downgrade to an earlier version of PAN-OS, you must first delete the certificates whose private keys are blocked. If you don't delete the certificates whose private keys are blocked before you attempt to downgrade, an error message asks you to delete those certificates. You can't downgrade until you delete them. After you downgrade, reimport or regenerate the deleted certificates if you need them.

If you use an enterprise Public Key Infrastructure (PKI) to generate certificates and private keys, block the export of private keys because you can install them on new firewalls and Panoramas from your enterprise certificate authority (CA), so there is no reason to export them from PAN-OS.

If you generate self-signed certificates on the firewall or Panorama and apply the block private key export option, you can't export the certificate and key to other PAN-OS appliances.

You can export and import the device state (**Device** > **Setup** > **Operations**) even if you block the export of private keys. We include the private keys in device state imports and exports, but administrators can't read or decode them.

You can import or load the configuration of one firewall on another firewall if the master key is the same on both firewalls. If the master key is different on the firewalls, then importing or loading the configuration doesn't work and the commit fails while reading the certificates.

- Generate a Private Key and Block It
- Import a Private Key and Block It
- Import a Private Key for IKE Gateway and Block It
- Verify Private Key Blocking

Generate a Private Key and Block It

Block the export of a private key to prevent its misuse after generating a certificate.

STEP 1 | Select **Device** > **Certificate Management** > **Certificates** > **Device Certificates**.

If there is more than one virtual system, select a **Location** or **Shared** for the certificate.

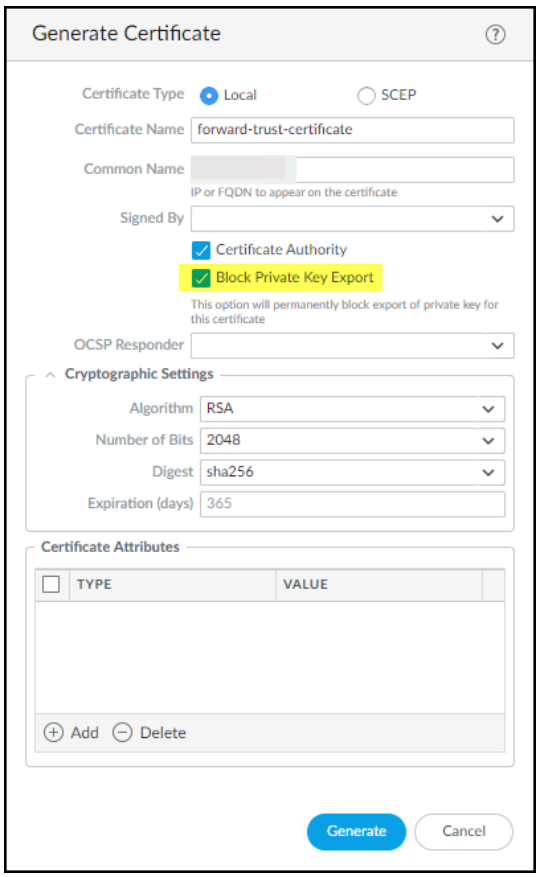
 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 2 | **Generate** the certificate.

STEP 3 | Select **Block Private Key Export** to prevent anyone from exporting the certificate.

See Generate a Certificate for information about the other certificate fields.



STEP 4 | Click **Generate** to generate the new certificate.

You can also generate a certificate and block its private key from export using the operational CLI command:

```
admin@pa-220> request certificate generate block-private-keys yes
```

The preceding CLI command can also include the certificate and other parameters that are not shown.

Import a Private Key and Block It

Block the export of a private key to prevent its misuse after importing a certificate.

STEP 1 | Select **Device** > **Certificate Management** > **Certificates** > **Device Certificates**.

If there is more than one virtual system, select a **Location** or **Shared** for the certificate.

STEP 2 | **Import** the certificate.

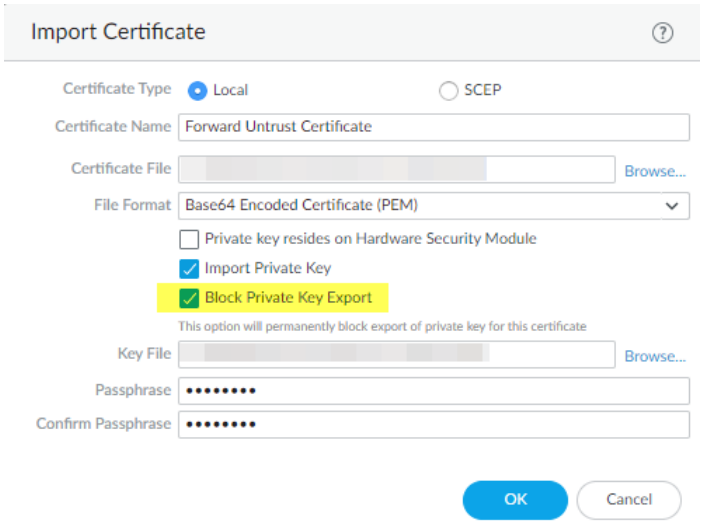
 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 3 | Select **Import Private Key** to activate the option to block private key export.

STEP 4 | Select **Block Private Key Export** to prevent anyone from exporting the certificate.

See Import a Certificate and Private Key for information about the other certificate import fields.



STEP 5 | Click **OK** to import the certificate.

If you use the SCP operational CLI command to import a certificate or to import a private key for a certificate, you can still block export of the private key:

- ***admin@pa-220> scp import private-key block-private-key ...***

Each of the preceding CLI commands can also include keywords to specify the source, the certificate name, and other parameters that are not shown.

*If you use the SCP operational CLI command to export a certificate and include its private key (**scp export certificate passphrase <phrase> remote-port <1-65536> to <destination> certificate-name <name> include-key <yes | no> format <der | pem | pkcs10 | pkcs12>**), and if the certificate's private key is blocked, the command fails and returns an error message because you cannot export a blocked private key.*

Import a Private Key for IKE Gateway and Block It

Block the export of a private key to prevent its misuse after generating a certificate for IKE Gateway authentication.

STEP 1 | Select **Network** > **Network Profiles** > **IKE Gateways**.

STEP 2 | **Add** a new IKE Gateway.

STEP 3 | On the **General** tab, for **Authentication**, select **Certificate**.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 4 | For **Local Certificate** select **Import** or **Generate** depending on whether you want to import an existing certificate or create a certificate.

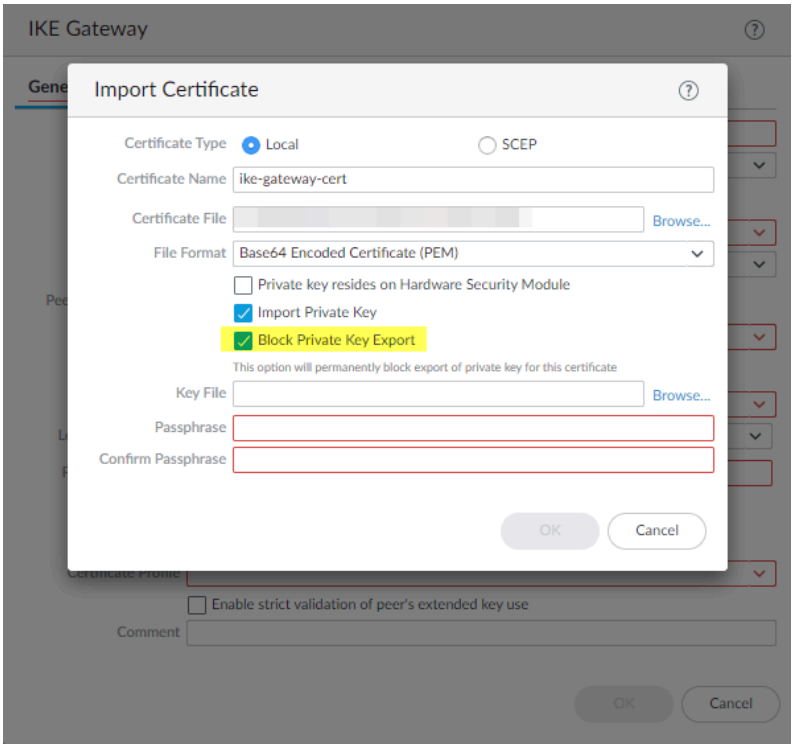
STEP 5 | Enter the certificate information. If you are importing the certificate, select **Import Private Key** to activate the **Block Private Key Export** checkbox.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 6 | Select **Block Private Key Export** to prevent anyone from exporting the key.

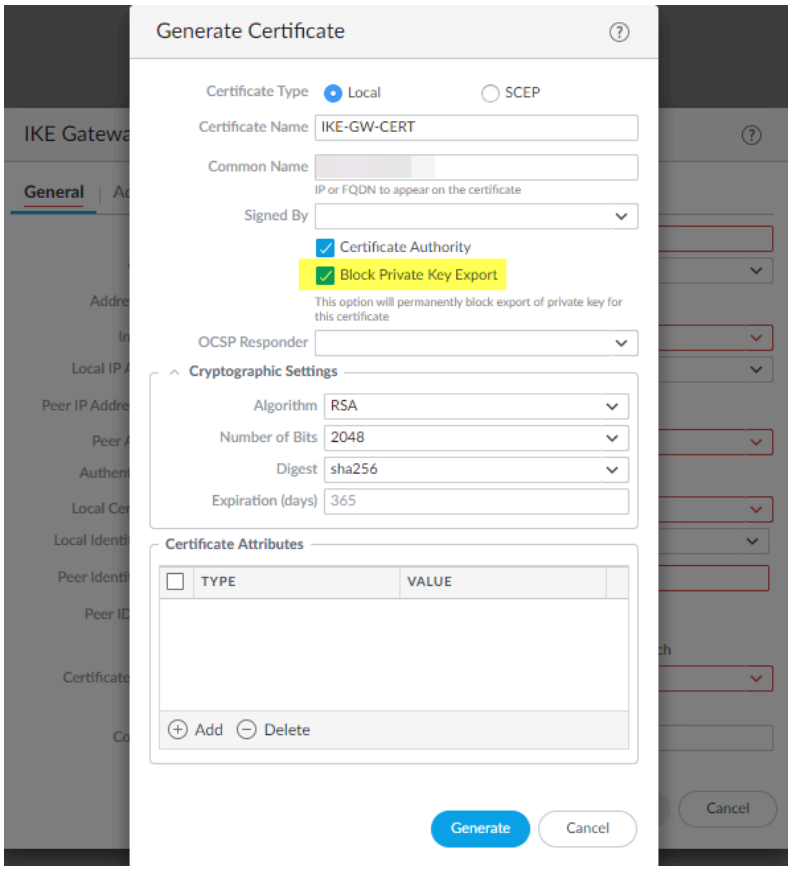
For importing a certificate, enter and confirm the **Passphrase** and then click **OK**



For generating a certificate, click **Generate**.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets



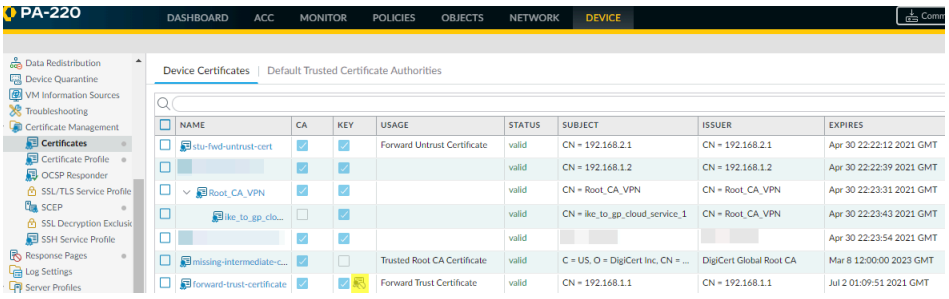
STEP 7 | Enter the **Passphrase**, confirm it, and then click **OK**.

Verify Private Key Blocking

You can verify whether a private key is blocked from export in several ways.

- Check the **Key** column in **Device** > **Certificate Management** > **Certificates** > **Device Certificates**.

  In this example, the forward-trust-certificate is blocked:



- When you attempt to export a certificate whose private key is blocked from export, the **Export Private Key** checkbox is not available and you can't export the key, you can only export the certificate.

https://t.me/learningnets

- Use the following operational CLI command to list all certificates on the device or in a particular Vsys that have private keys blocked from export:

```
admin@pa-220> request certificate show-blocked <shared | vsys>
```

- Use the following operational CLI command to check whether a particular certificate's private key is blocked from export:

```
admin@pa-220> request certificate is-blocked certificate-name
 <name>
```

If the certificate is blocked from export, the command returns **yes** and if the certificate is not blocked the command returns **no**.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Enable Users to Opt Out of SSL Decryption

In privacy-sensitive situations, you may want to alert your users that the firewall is decrypting certain web traffic and allow them either to continue to the site with the understanding that their traffic is decrypted or to terminate the session and be block from going to the site. (There is no option to go to the site and also avoid decryption.)

The first time a user attempts to browse to an HTTPS site or application that matches the decryption policy, the firewall displays a response page notifying users that it will decrypt the session. Users can either click **Yes** to allow decryption and continue to the site or click **No** to opt out of decryption and terminate the session. The choice to allow decryption applies to all HTTPS sites that users try to access for the next 24 hours, after which the firewall redisplays the response page. Users who opt out of SSL decryption cannot access the requested web page, or any other HTTPS site, for the next minute. After the minute elapses, the firewall redisplays the response page the next time the users attempt to access an HTTPS site.

The firewall includes a predefined SSL Decryption Opt-out Page that you can enable. You can optionally customize the page with your own text and/or images. However, the best practice is to not allow users to opt out of decryption.



Custom response pages larger than the maximum supported size are not decrypted or displayed to users. In PAN-OS 8.1.2 and earlier PAN-OS 8.1 releases, custom response pages on a decrypted site cannot exceed 8,191 bytes; the maximum size is increased to 17,999 bytes in PAN-OS 8.1.3 and later releases.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 1 | (Optional) Customize the SSL Decryption Opt-out Page.

1. Select **Device** > **Response Pages**.
2. Select the **SSL Decryption Opt-out Page** link.
3. Select the **Predefined** page and click **Export**.
4. Using the HTML text editor of your choice, edit the page.
5. If you want to add an image, host the image on a web server that is accessible from your end user systems.
6. Add a line to the HTML to point to the image. For example:

```
<img src="http://cdn.slidesharecdn.com/ Acme-logo-96x96.jpg?
1382722588"/>
```

7. Save the edited page with a new filename. Make sure that the page retains its UTF-8 encoding.
8. Back on the firewall, select **Device** > **Response Pages**.
9. Select the **SSL Decryption Opt-out Page** link.
10. Click **Import** and then enter the path and filename in the **Import File** field or **Browse** to locate the file.
11. (Optional) Select the virtual system on which this login page will be used from the **Destination** drop-down or select shared to make it available to all virtual systems.
12. Click **OK** to import the file.
13. Select the response page you just imported and click **Close**.

STEP 2 | Enable SSL Decryption Opt Out.

1. On the **Device** > **Response Pages** page, click the **Disabled** link.
2. Select the **Enable SSL Opt-out Page** and click **OK**.
3. **Commit** the changes.

STEP 3 | Verify that the Opt Out page displays when you attempt to browse to a site.

From a browser, go to an encrypted site that matches your decryption policy.

Verify that the SSL Decryption Opt-out response page displays.



https://t.me/learningnets

Temporarily Disable SSL Decryption

In some cases you may want to temporarily disable SSL decryption. For example, if you deployed SSL decryption too hastily and something doesn't work correctly but you're not sure what it is, and you have a lot of rules to examine, you can use the CLI to temporarily turn off decryption and give yourself time to analyze and solve the issue. After solving the issue, you can use the CLI to turn SSL decryption back on again. Because temporarily disabling and then re-enabling decryption using the CLI doesn't require a Commit operation, you can do it without disrupting network traffic.

The following CLI commands temporarily disable SSL decryption without a Commit and re-enable decryption without a Commit.

The command to disable SSL decryption doesn't persist in the configuration after a reboot. If you turn off decryption temporarily and then reboot the firewall, regardless of whether the issue has been fixed, decryption is turned on again.

- Disable SSL Decryption

```
set system setting
ssl-decrypt skip-ssl-decrypt yes
```

- Re-enable SSL Decryption

```
set system setting
ssl-decrypt skip-ssl-decrypt no
```

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

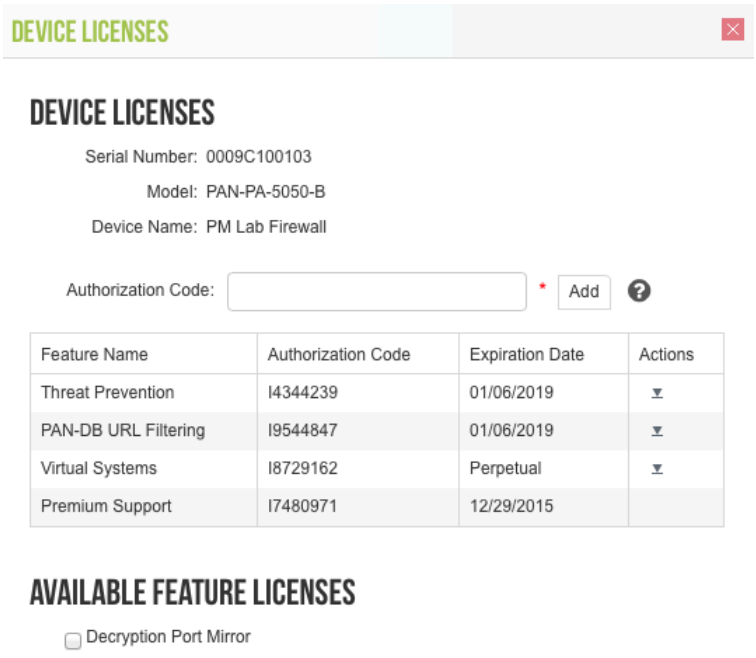
Configure Decryption Port Mirroring

Before you can enable Decryption Mirroring, you must obtain and install a Decryption Port Mirror license. The license is free of charge and can be activated through the support portal as described in the following procedure. After you install the Decryption Port Mirror license and reboot the firewall, you can enable decryption port mirroring.

Keep in mind that the decryption, storage, inspection, and/or use of SSL traffic is regulated in certain countries and user consent may be required in order to use the decryption mirror feature. Additionally, use of this feature could enable malicious users with administrative access to the firewall to harvest usernames, passwords, social security numbers, credit card numbers, or other sensitive information submitted using an encrypted channel. Palo Alto Networks recommends that you consult with your corporate counsel before activating and using this feature in a production environment.

STEP 1 | Request a license for each firewall on which you want to enable decryption port mirroring.

1. Log in to the Palo Alto Networks Customer Support website and navigate to the **Assets** tab.
2. Select the entry for the firewall you want to license and select **Actions**.
3. Select **Decryption Port Mirror**. A legal notice displays.
4. If you are clear about the potential legal implications and requirements and still want to set up decryption port mirroring, click **I understand and wish to proceed**.
5. Click **Activate**.



 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 2 | Install the Decryption Port Mirror license on the firewall.

1. From the firewall web interface, select **Device** > **Licenses**.
2. Click **Retrieve license keys from license server**.
3. Verify that the license has been activated on the firewall.



4. Reboot the firewall (**Device** > **Setup** > **Operations**). This feature is not available for configuration until PAN-OS reloads.

STEP 3 | Enable the firewall to forward decrypted traffic. Superuser permission is required to perform this step.

On a firewall with a single virtual system:

1. Select **Device** > **Setup** > **Content - ID**.
2. Select the **Allow forwarding of decrypted content** check box.
3. Click **OK** to save.

On a firewall with multiple virtual systems:

1. Select **Device** > **Virtual System.**
2. Select a Virtual System to edit or create a new Virtual System by selecting **Add**.
3. Select the **Allow forwarding of decrypted content** check box.
4. Click **OK** to save.

STEP 4 | Enable an Ethernet interface to be used for decryption mirroring.

1. Select **Network** > **Interfaces** > **Ethernet**.
2. Select the Ethernet interface that you want to configure for decryption port mirroring.
3. Select **Decrypt Mirror** as the **Interface Type**.

   This interface type will appear only if the Decryption Port Mirror license is installed.
4. Click **OK** to save.

STEP 5 | Enable mirroring of decrypted traffic.

1. Select **Objects** > **Decryption Profile**.
2. Select an **Interface** to be used for **Decryption Mirroring**.

   The **Interface** drop-down contains all Ethernet interfaces that have been defined as the type: **Decrypt Mirror**.
3. Specify whether to mirror decrypted traffic before or after policy enforcement.

   By default, the firewall will mirror all decrypted traffic to the interface before security policies lookup, which allows you to replay events and analyze traffic that generates a threat or triggers a drop action. If you want to only mirror decrypted traffic after security policy enforcement, select the **Forwarded Only** check box. With this option, only traffic that is forwarded through the firewall is mirrored. This option is useful if you are

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

forwarding the decrypted traffic to other threat detection devices, such as a DLP device or another intrusion prevention system (IPS).

4. Click **OK** to save the decryption profile.

STEP 6 | Attach the decryption profile rule (with decryption port mirroring enabled) to a decryption policy rule. All traffic decrypted based on the policy rule is mirrored.

1. Select **Policies** > **Decryption**.
2. Click **Add** to configure a decryption policy or select an existing decryption policy to edit.
3. In the **Options** tab, select **Decrypt** and the **Decryption Profile** created in step 4.
4. Click **OK** to save the policy.

STEP 7 | Save the configuration.

Click **Commit**.

 ©2022 Palo Alto Networks, Inc.

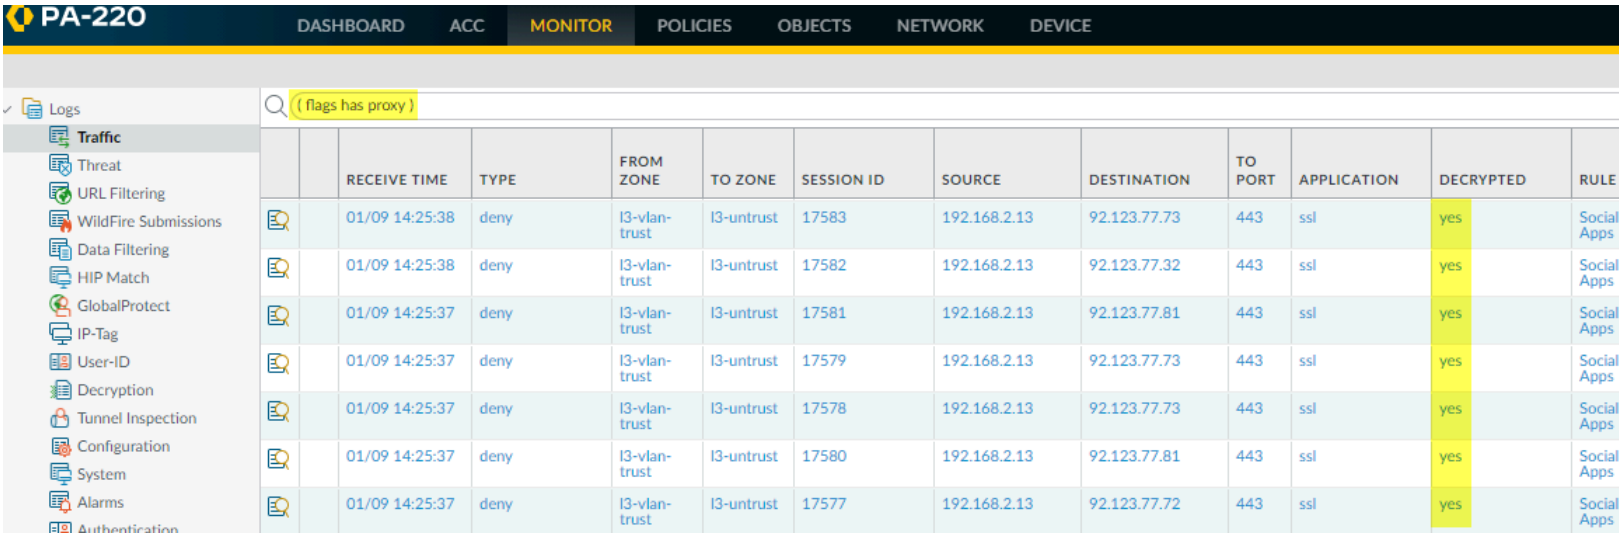
https://t.me/learningnets

Verify Decryption

After you configure a best practice decryption profile and apply it to traffic, you can check both the Decryption logs (introduced in PAN-OS 10.0) and the Traffic logs to verify that the firewall is decrypting the traffic that you intend to decrypt and that the firewall is not decrypting the traffic that you don't want to decrypt. This topic shows you how to check decryption using Traffic logs. In addition, follow post-deployment decryption best practices to maintain the deployment.

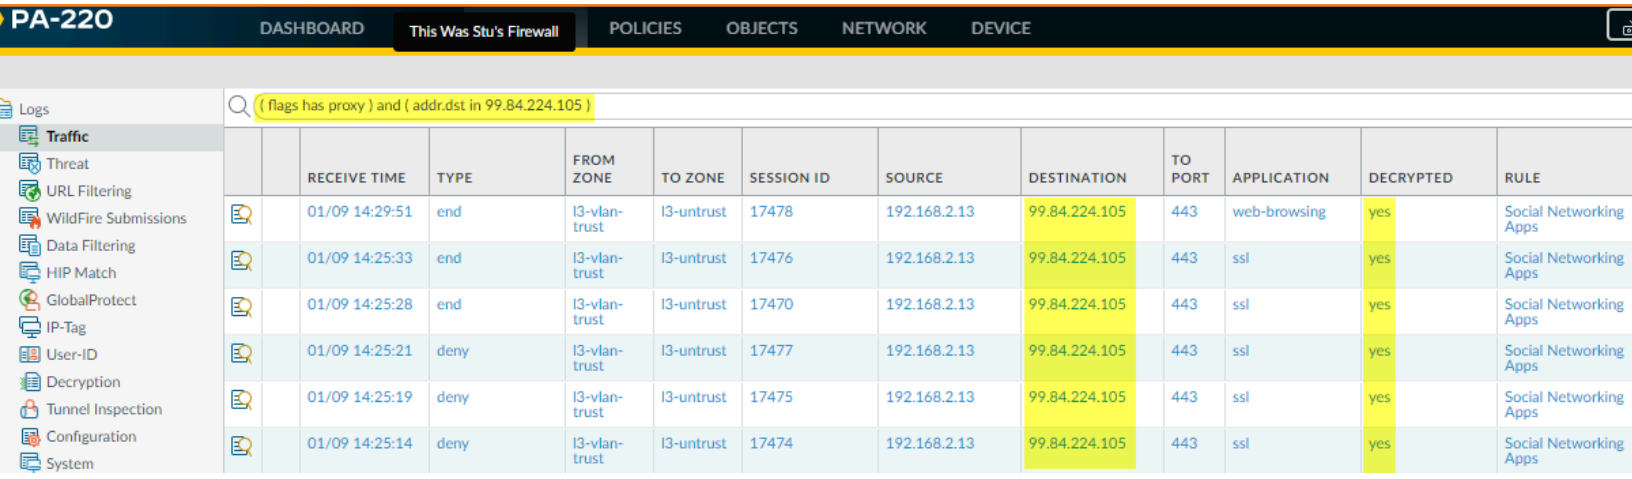
- **View Decrypted Traffic Sessions**—Filter the Traffic Logs (**Monitor** > **Logs** > **Traffic**) using the filter **`( flags has proxy )`**.

  This filter displays only logs in which the SSL proxy flag is on, meaning only decrypted traffic—every log entry has the value yes in the **Decrypted** column.



		RECEIVE TIME	TYPE	FROM ZONE	TO ZONE	SESSION ID	SOURCE	DESTINATION	TO PORT	APPLICATION	DECRYPTED	RULE
		01/09 14:25:38	deny	l3-vlan-trust	l3-untrust	17583	192.168.2.13	92.123.77.73	443	ssl	yes	Social Apps
		01/09 14:25:38	deny	l3-vlan-trust	l3-untrust	17582	192.168.2.13	92.123.77.32	443	ssl	yes	Social Apps
		01/09 14:25:37	deny	l3-vlan-trust	l3-untrust	17581	192.168.2.13	92.123.77.81	443	ssl	yes	Social Apps
		01/09 14:25:37	deny	l3-vlan-trust	l3-untrust	17579	192.168.2.13	92.123.77.73	443	ssl	yes	Social Apps
		01/09 14:25:37	deny	l3-vlan-trust	l3-untrust	17578	192.168.2.13	92.123.77.73	443	ssl	yes	Social Apps
		01/09 14:25:37	deny	l3-vlan-trust	l3-untrust	17580	192.168.2.13	92.123.77.81	443	ssl	yes	Social Apps
		01/09 14:25:37	deny	l3-vlan-trust	l3-untrust	17577	192.168.2.13	92.123.77.72	443	ssl	yes	Social Apps

You can filter the traffic in a more granular fashion by adding more terms to the filter. For example, you can filter for decrypted traffic going only to the destination IP address 99.84.224.105 by adding the filter **`( addr.dst in 99.84.224.105 )`**:

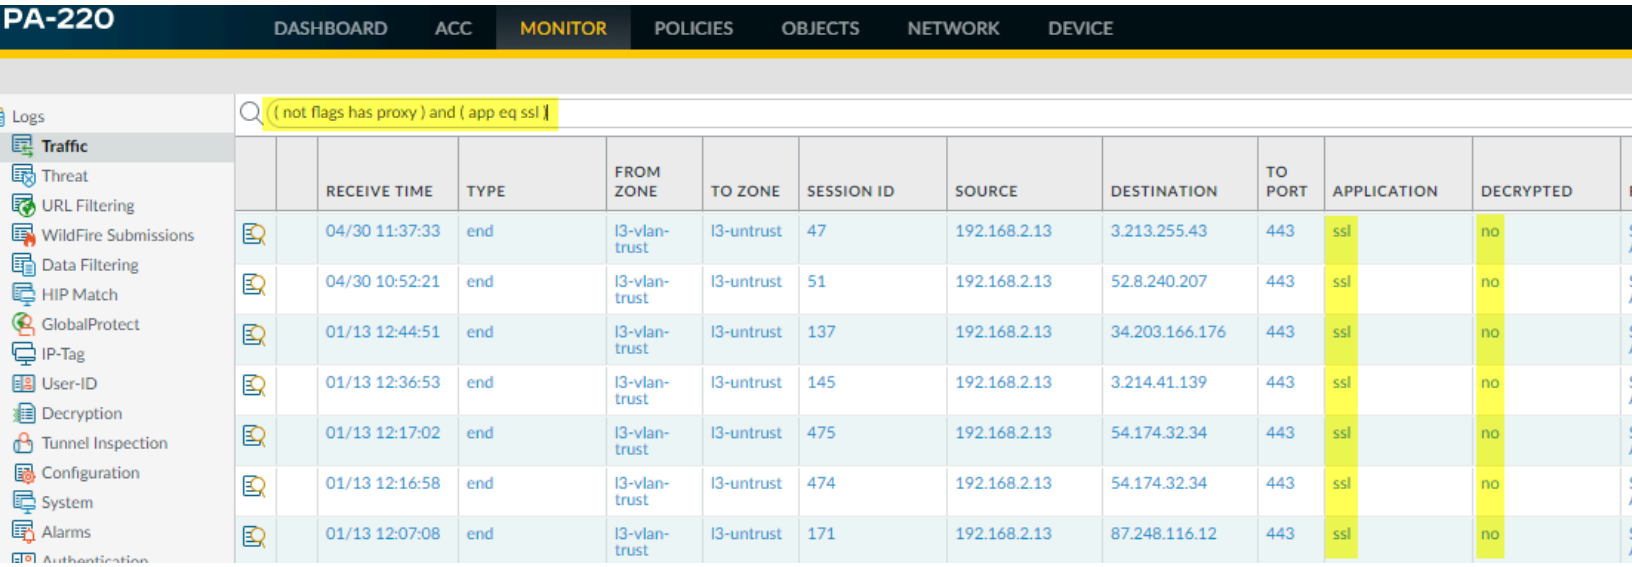


		RECEIVE TIME	TYPE	FROM ZONE	TO ZONE	SESSION ID	SOURCE	DESTINATION	TO PORT	APPLICATION	DECRYPTED	RULE
		01/09 14:29:51	end	l3-vlan-trust	l3-untrust	17478	192.168.2.13	99.84.224.105	443	web-browsing	yes	Social Networking Apps
		01/09 14:25:33	end	l3-vlan-trust	l3-untrust	17476	192.168.2.13	99.84.224.105	443	ssl	yes	Social Networking Apps
		01/09 14:25:28	end	l3-vlan-trust	l3-untrust	17470	192.168.2.13	99.84.224.105	443	ssl	yes	Social Networking Apps
		01/09 14:25:21	deny	l3-vlan-trust	l3-untrust	17477	192.168.2.13	99.84.224.105	443	ssl	yes	Social Networking Apps
		01/09 14:25:19	deny	l3-vlan-trust	l3-untrust	17475	192.168.2.13	99.84.224.105	443	ssl	yes	Social Networking Apps
		01/09 14:25:14	deny	l3-vlan-trust	l3-untrust	17474	192.168.2.13	99.84.224.105	443	ssl	yes	Social Networking Apps

https://t.me/learningnets

- **View SSL Traffic Sessions That Are Not Decrypted**—Filter the Traffic Logs (**Monitor** > **Logs** > **Traffic**) using the filter **`( not flags has proxy ) and ( app eq ssl )`**.

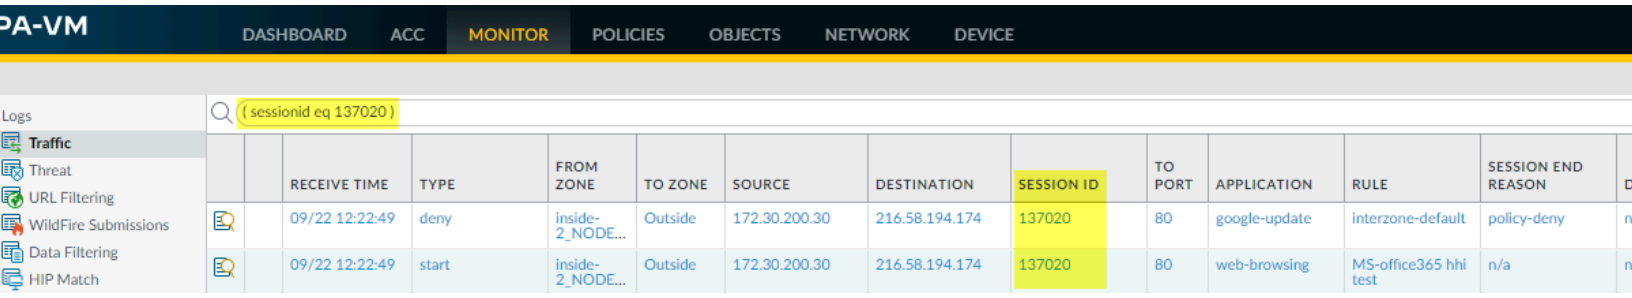
  This filter displays only logs in which the SSL proxy flag is off (meaning only encrypted traffic) and the traffic is SSL traffic; every log entry has the value **no** in the **Decrypted** column and the value **ssl** in the **Application** column.



Similar to the example for viewing decrypted traffic logs, you can add terms to filter the traffic that you don't decrypt in a more granular fashion.

- **View The Log for a Particular Session**—To view the Traffic log for a particular session, filter on the Session ID.

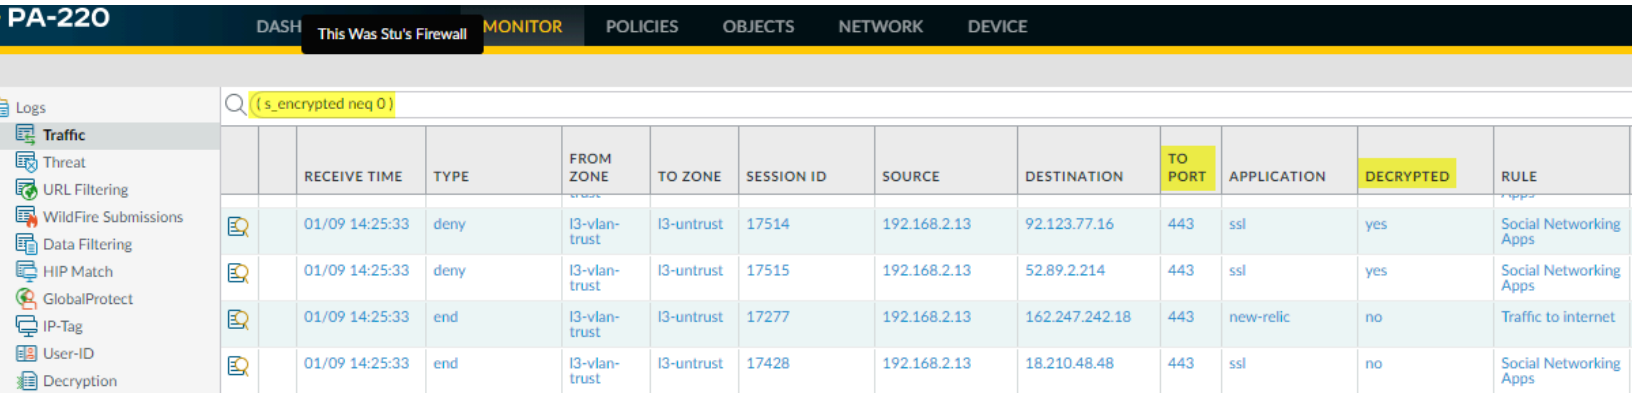
  For example, to see the log for a session with the ID 137020, filter using the term **`( sessionid eq 137020)`**. You can find the ID number in the Session ID column in the log output, as shown in the previous screens. If the Session ID column isn't displayed, add the column to the output.



 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

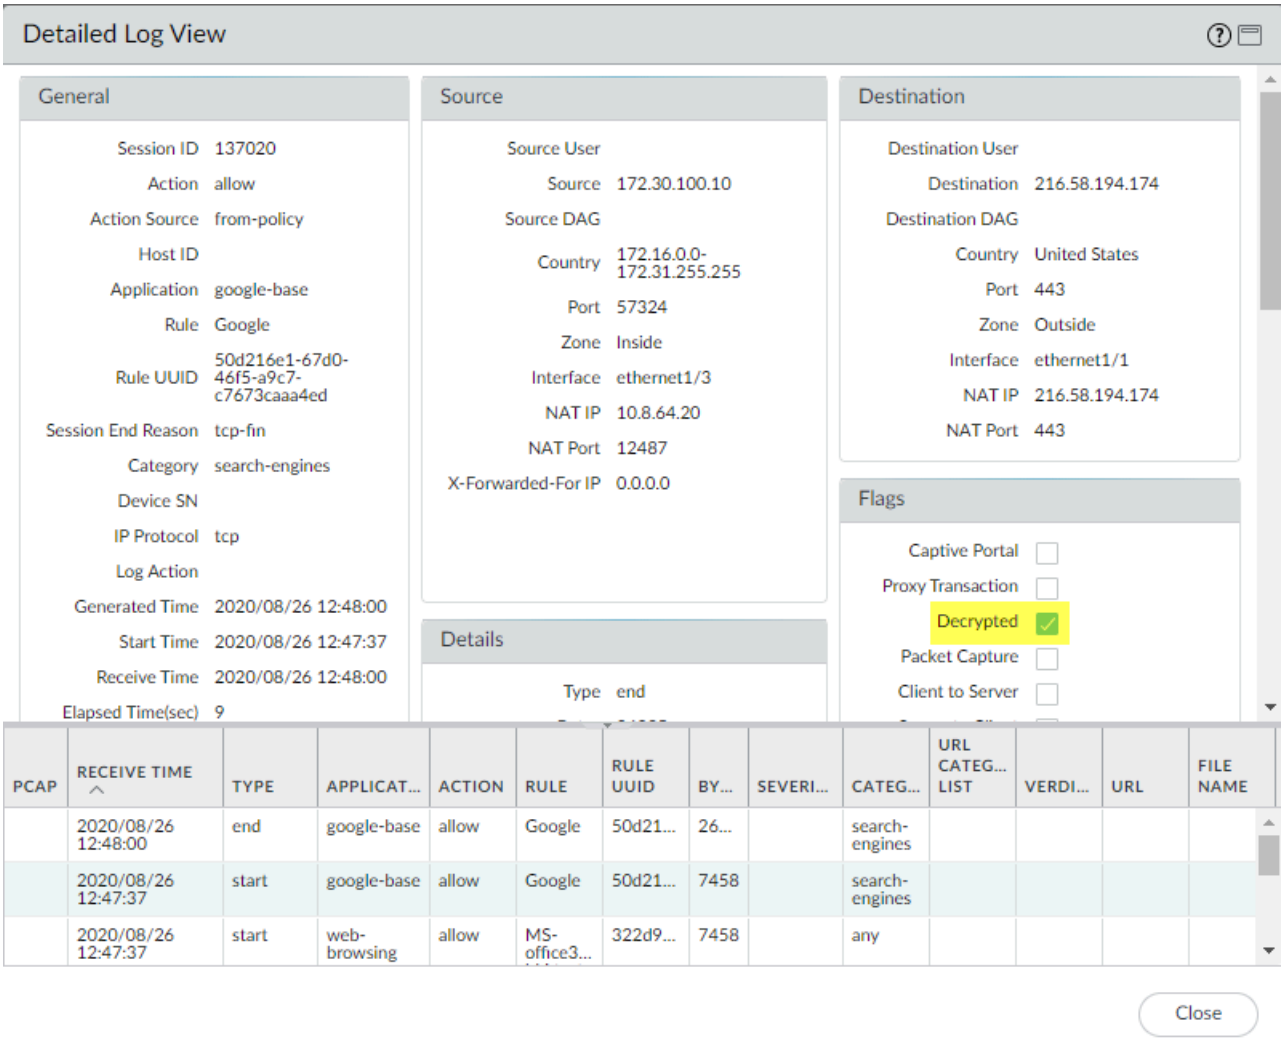
- **View All TLS and SSH Traffic**—Filter the Traffic Logs (**Monitor** > **Logs** > **Traffic**) to view both decrypted and undecrypted TLS and SSH traffic, use the filter **`( s_encrypted neq 0 )`**:



 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

- **Drill Down Into the Details**—To view more information about a particular log entry, click the magnifying glass to see a detailed log view. For example, for Session ID 137020 (shown in the previous bullet), the detailed log looks like this:



The box for the **Decrypted** flag provides a second way to verify if traffic was decrypted.

You can also take upstream and downstream packet captures of decrypted traffic to view how the firewall processes SSL traffic and takes actions on packets, or perform deep packet inspection.

https://t.me/learningnets

Troubleshoot and Monitor Decryption

Troubleshooting tools provide enhanced visibility into TLS traffic so you can monitor your decryption deployment. The tools enable you to diagnose and resolve decryption issues quickly and easily, tighten weaknesses in your decryption deployment, and fix decryption issues to improve your security posture. For example, you can:

- Identify traffic that causes decryption failures by Service Name Identification (SNI) and application.
- Identify traffic that uses weak protocols and algorithms.
- Examine successful and unsuccessful decryption activity in the network.
- View detailed information about individual sessions.
- Profile decryption usage and patterns.
- Monitor detailed decryption statistics and information about adoption, failures, versions, algorithms, etc.

The following tools provide full visibility into the TLS handshake and help you troubleshoot and monitor your decryption deployment:

- **ACC** > **SSL Activity**—The five ACC widgets on this tab (introduced in PAN-OS 10.0) provide details about successful and unsuccessful decryption activity in your network, including decryption failures, TLS versions, key exchanges, and the amount and type of decrypted and undecrypted traffic.
- **Monitor** > **Logs** > **Decryption**—The Decryption Log (introduced in PAN-OS 10.0) provides comprehensive information about individual sessions that match a Decryption policy (use a No Decryption policy for traffic you don't decrypt) and about GlobalProtect sessions when you enable Decryption logging in GlobalProtect Portal or GlobalProtect Gateways configuration. Select which columns to display to view information such as application, SNI, Decryption Policy Name, error index, TLS version, key exchange version, encryption algorithm, certificate key types, and many other characteristics. Filter the information in columns to identify traffic that uses particular TLS versions and algorithms, particular errors, or any other characteristics you want to investigate. By default, Decryption policies log only unsuccessful TLS handshakes. Depending on the available log storage, you can configure Decryption policies to log successful TLS handshakes as well.
- **Local Decryption Exclusion Cache**—There are two constructs for sites that break decryption for technical reasons such as client authentication or pinned certificates and therefore need to be excluded from decryption: the SSL Decryption Exclusion List and the Local Decryption Exclusion Cache. The SSL Decryption Exclusion List contains the sites that Palo Alto Networks has identified that break decryption technically. Content updates keep the list up-to-date and you can add sites to the list manually. The Local Decryption Exclusion Cache automatically adds sites that local users encounter that break decryption for technical reasons and excludes them from decryption, providing that the Decryption profile applied to the traffic allows unsupported modes (if unsupported modes are blocked, then the traffic is blocked instead of added to the local cache).
- **Custom Report Templates for Decryption**—You can create custom reports (**Monitor** > **Manage Custom Reports**) using four predefined templates that summarize decryption activity (introduced in PAN-OS 10.0).

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

The general troubleshooting methodology is to use the new ACC widgets to identify traffic that causes decryption issues and then use the new Decryption Log and custom report templates to drill down into details and gain context about that traffic, which enables you to diagnose issues accurately and much more easily than in the past. Understanding decryption issues and their causes enables you to select the appropriate way to fix each issue, such as:

- Modify Decryption policy rules (a policy rule defines traffic that the rule affects, the action taken on that traffic, log settings, and the Decryption profile applied to the traffic)
- Modify Decryption profiles (acceptable protocols and algorithms for the traffic that a Decryption policy rule defines, plus failure checks, unsupported mode checks for items such as unsupported ciphers and versions, certificate checks, etc.)
- Add sites that break decryption for technical reasons to the SSL Decryption Exclusion List
- Evaluate security decisions about which sites your employees, customers, and partners really need to access and which sites you can block when sites use weak decryption protocols or algorithms

The goals should be to decrypt all the traffic you can decrypt (a decryption best practice) so that you can inspect it and to properly handle traffic that you don't decrypt.

When you upgrade to PAN-OS 10.0, the device takes 1% of the log space and allocates it to Decryption logs. Step 3 in Configure Decryption Logging shows you how to modify the log space allocation to provide more space for Decryption logs.

If you downgrade from PAN-OS 10.1 or later to PAN-OS 9.1 or earlier, the features introduced in PAN-OS 10.1 (Decryption Log, SSL Activity widgets in the ACC, custom report Decryption templates) are removed from the UI. References to Decryption logs are also removed from Log Forwarding profiles. In addition, the Local Decryption Exclusion Cache is only viewable using the CLI in PAN-OS 9.1 and earlier (PAN-OS 10.1 added the local cache to the UI).

If you push configurations from Panorama on PAN-OS 10.1 or later to devices that run PAN-OS 9.1 or earlier, Panorama removes the features introduced in PAN-OS 10.0.

- Decryption Application Command Center Widgets
- Decryption Log
- Custom Report Templates for Decryption
- Decryption Troubleshooting Workflow Examples

Decryption Application Command Center Widgets

The Application Command Center (ACC) widgets for decryption (**ACC** > **SSL Activity**) introduced in PAN-OS 10.1 work with Decryption Log to help you diagnose and resolve decryption issues quickly and easily. Use the **SSL Activity** widget to view and analyze network decryption activity such as the number of decrypted and undecrypted sessions, how much traffic uses different TLS protocol versions, the most common decryption failure reasons, and which applications and Server Name Identifications (SNIs) use weak ciphers and algorithms. Next, use the Decryption logs to drill down into sessions and diagnose the exact issue so you can take appropriate action.

PAN-OS 10.1 introduced five new decryption widgets. Use the information the widgets provide to identify misconfigured Decryption policies and profiles and to make informed decisions about what traffic to allow and what traffic to block:

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

- **Traffic Activity**—Shows SSL/TLS activity compared to non-SSL/TLS activity by total number of sessions or by amount of traffic in bytes.
- **SSL/TLS Traffic**—Shows the amount of decrypted and non-decrypted traffic by number of sessions or amount of traffic in bytes. Reasons for traffic not being decrypted include:
  - No Decryption policy is applied to the traffic.
  - The Decryption policy intentionally exempted the traffic from decryption (for example, a No Decryption policy).
  - The Decryption policy was misconfigured and the traffic was intended to be decrypted but is not.
  - The site is in the SSL Decryption Exclusion List (**Device** > **Certificate Management** > **SSL Decryption Exclusion**), which contains sites Palo Alto Networks has identified that break decryption for technical reasons such as pinned certificates or client authentication. For these sites, the firewall bypasses decryption.
  - The site is in the Local Decryption Exclusion Cache, which contains sites that local users encounter which prevent decryption for technical reasons.

The ACC only populates the next three widgets with data from traffic that a Decryption policy controls. If you don't apply a Decryption policy to traffic, that traffic does not populate these widgets.

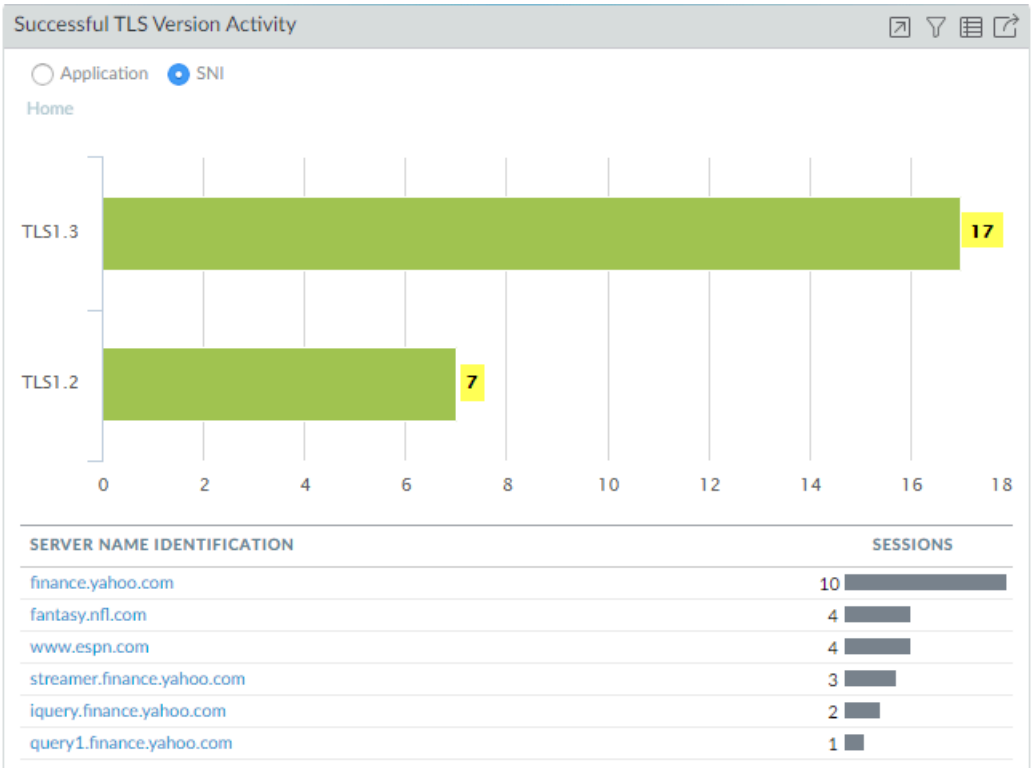
- **Decryption Failure Reasons**—Shows the reasons for decryption failures: protocol, certificate, version, cipher, HSM, resource, resume, or feature issues, by SNI. Use this information to detect problems caused by Decryption policy or profile misconfiguration or by traffic that uses unsupported weak protocols or algorithms. Click a failure reason to drill down and isolate the number of sessions per SNI that experienced the failure or click an SNI to see all of the decryption failures for that SNI.
- **Successful TLS Version Activity**—Shows successful TLS connections by TLS version for applications or SNIs (SNIs are available for Forward Proxy only) so you can evaluate how much risk you are taking on by allowing weaker TLS protocol versions. Identifying applications and SNIs that use weak protocols enables you to evaluate each one and decide whether you need to allow access to it for business reasons. If you don't need the application for business purposes, you may want to block the traffic instead of allowing it to reduce risk. Click a TLS version to drill down and view the SNIs or applications which used that TLS version. Click an application or an SNI to drill down and see how many of those application or SNI sessions used each TLS version.
- **Successful Key Exchange Activity**—Shows successful key exchange activity per algorithm for applications or SNIs (SNIs are available for Forward Proxy only). Click a key exchange algorithm to see the activity for just that algorithm or click an application or SNI to view the key exchange algorithm activity for that application or SNI.

The following example of drilling down into ACC data shows you how to examine successful TLS version activity:

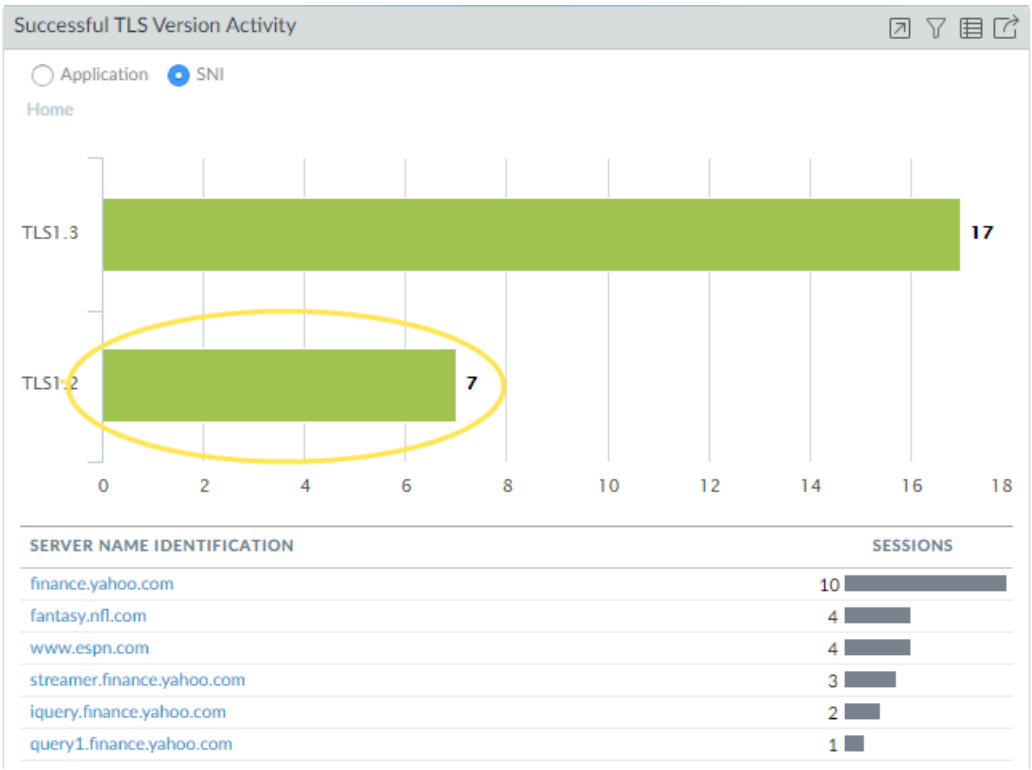
 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

1. The **Successful TLS Version Activity** widget shows that seventeen sessions used TLSv1.3 and seven sessions used TLSv1.2. The SNI list shows the destination SNIs and the number of sessions per SNI.



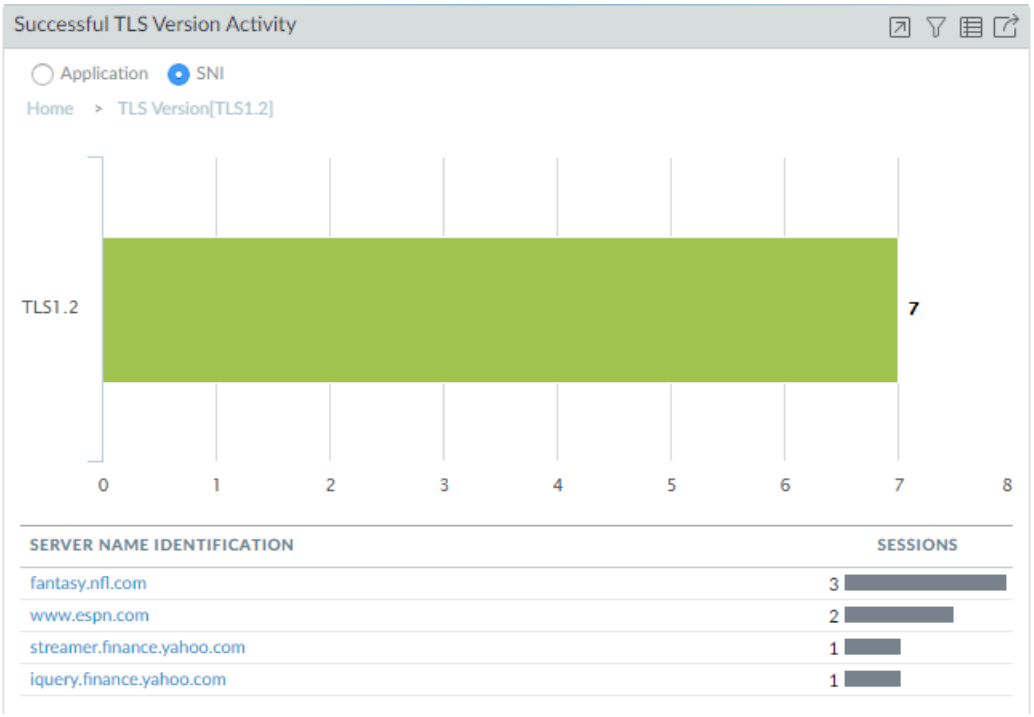
2. To see which SNIs used TLSv1.2, click the green bar labeled TLS1.2.



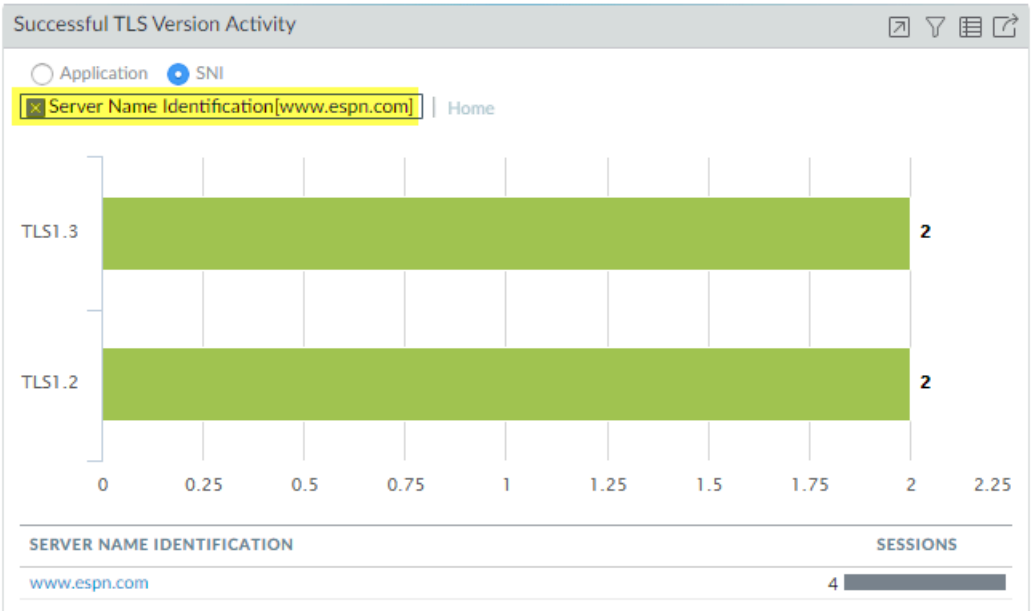
 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

3. Now you can see the seven TLSv1.2 sessions were spread among four servers.



4. Clicking **Home** returns to the home screen. Now, clicking the www.espn.com SNI shows us which TLS versions it used. We can see that two of the four sessions used TLSv1.3 and two used TLSv1.2.



For any Decryption widget, click the Jump to Logs icon to jump directly to the Decryption logs that correspond to the data in the ACC:



 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

In the preceding example, at any point in the investigation you could jump to the Decryption logs for the data to drill down more. For example, you could examine the logs for the individual sessions that used TLSv1.2 to find out why they didn't use TLSv1.3.

Decryption ACC widgets show the name of the decrypted application based on the Palo Alto Networks App-ID. For populating the ACC, the firewall can only identify applications that have a Palo Alto Networks App-ID; the firewall cannot populate the ACC with custom applications or applications that do not have an App-ID. Content updates update App-IDs regularly. Other reasons that the application may be shown as incomplete or unknown are:

- The firewall dropped the session before it could identify the application.
- Decryption logs depend on Traffic logs to populate the Decryption log application field. However, if the Traffic log is not completed in 60 seconds or less, the Traffic log does not populate the application in the Decryption log and the application displays as incomplete or unknown.

Decryption Log

The Decryption Log (**Monitor** > **Logs** > **Decryption**) provides comprehensive information about sessions that match a Decryption policy to help you gain context about that traffic so you can accurately and easily diagnose and resolve decryption issues. The firewall does not log traffic if the traffic does not match a Decryption policy. If you want to log traffic that you don't decrypt, create a policy-based decryption exclusion and for policies that govern TLSv1.2 and earlier traffic, apply a No Decryption profile to the traffic.

PAN-OS supports Decryption logs for the following types of traffic:

- Forward Proxy—Several fields only display information for Forward Proxy traffic, including Root CA (for trusted certificates only) and Server Name Identification (SNI).
- Inbound Inspection.
- No Decrypt (traffic excluded from decryption by Decryption policy).

  Because the session remains encrypted, the firewall displays less information. For undecrypted TLSv1.3 traffic, there is no certificate information because TLSv1.3 encrypts certificate information.
- GlobalProtect—Covers GlobalProtect Gateway, GlobalProtect Portal, and GlobalProtect Clientless VPN (client-to-firewall only).

  GlobalProtect does not support TLSv1.3.
- Decryption Mirror

Not all types of traffic support every parameter. Unsupported Parameters by Proxy Type and TLS Version *provides a complete list of unsupported parameters for each type of decryption traffic.*

The data for Forward Proxy traffic is based on whether the TLS handshake is successful or unsuccessful. For unsuccessful TLS handshakes, the firewall sends error data for the leg of the transaction that caused the error, either client-to-firewall or firewall-to-server. For successful TLS

 ©2022 Palo Alto Networks, Inc.

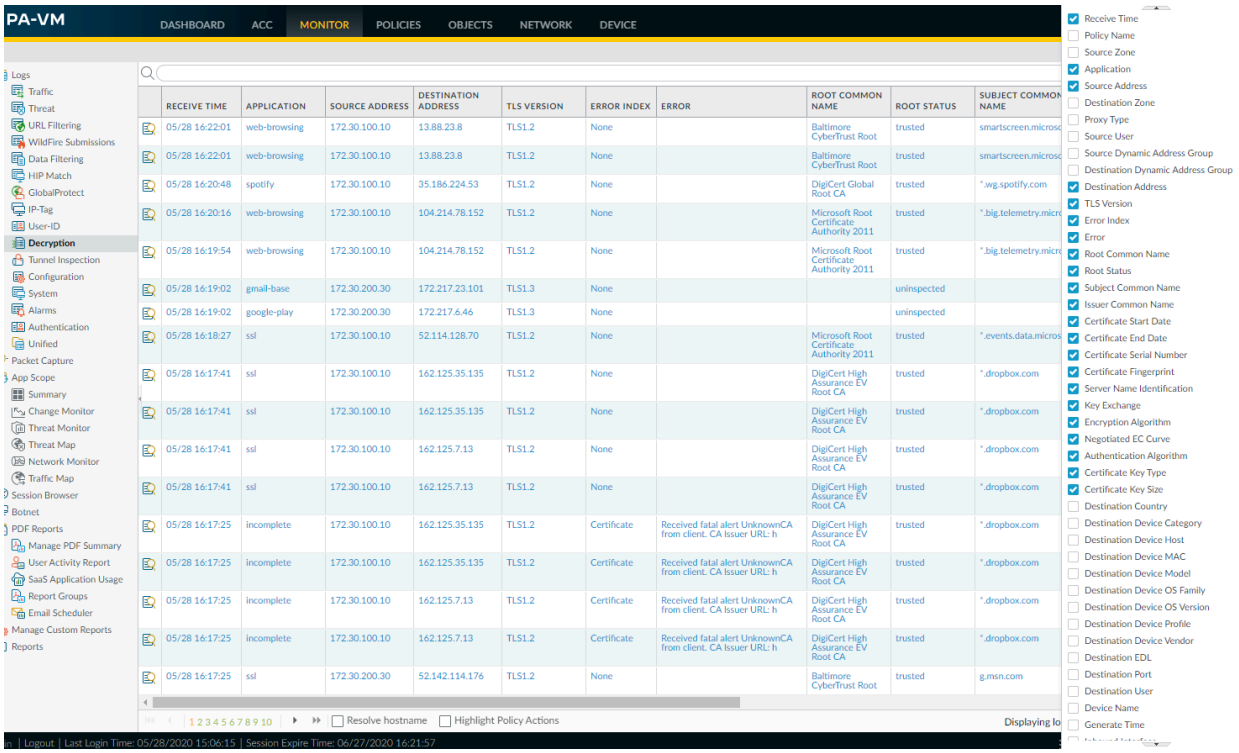
https://t.me/learningnets

handshakes, the data is from the leg that successfully completes first, which is usually client-to-firewall.

The firewall does not generate Decryption log entries for web traffic blocked during SSL/TLS handshake inspection. These sessions do not appear in Decryption logs because the firewall prevents decryption when it resets the SSL/TLS connection, ending the handshake. You can view details of the blocked sessions in the URL Filtering logs.

Decryption logs are not supported for SSH Proxy traffic. In addition, certificate information is not available for session resumption logs.

By default, the firewall logs all unsuccessful TLS handshake traffic. You can also log successful TLS handshake traffic if you choose to do so. You can view up to 62 columns of log information such as application, SNI, Decryption Policy Name, error index, TLS version, key exchange version, encryption algorithm, certificate key types, and many other characteristics:



Click the magnifying glass icon ( ) to see the Detailed Log View of a session.

https://t.me/learningnets

*The Decryption log learns each session's App-ID from the Traffic log, so Traffic logs must be enabled to see the App-ID in the Decryption log. If Traffic logs are disabled, the App-ID shows as **incomplete**. For example, a lot of GlobalProtect traffic is intrazone traffic (Untrust zone to Untrust zone), but the default intra-zone policy does not enable Traffic logs. To see the App-ID for GlobalProtect intrazone traffic, you need to enable the Traffic log for intrazone traffic.*

*Another reason that the App-ID may display as **incomplete** is that for long sessions, the firewall may generate the Decryption log before the Traffic log is complete (the Traffic log is usually generated at session end). In those cases, the App-ID is not available for the Decryption log. In addition, when the TLS handshake fails and generates an error log, the App-ID is not available because the failure terminates the session before the firewall can determine the App-ID. In these cases, the application may display as **ssl** or as **incomplete**.*

To troubleshoot issues, use the Decryption ACC widgets (**ACC** > **SSL Activity**) to identify traffic that causes decryption issues and then use the Decryption log and Custom Report Templates for Decryption to drill down into details.

When you forward Decryption logs for storage, ensure that you properly secure log transport and storage because Decryption logs contain sensitive information.

When the Decryption logs are enabled, the firewall sends HTTP/2 logs as Tunnel Inspection logs (when Decryption logs are disabled, HTTP/2 logs are sent as Traffic logs), so you need to check the Tunnel Inspection logs instead of the Traffic logs for HTTP/2 events. In addition, you must enable Tunnel Content Inspection to obtain the App-ID for HTTP/2 traffic.

- Configure Decryption Logging
- Repair Incomplete Certificate Chains
- Decryption Log Errors, Error Indexes, and Bitmasks

Configure Decryption Logging

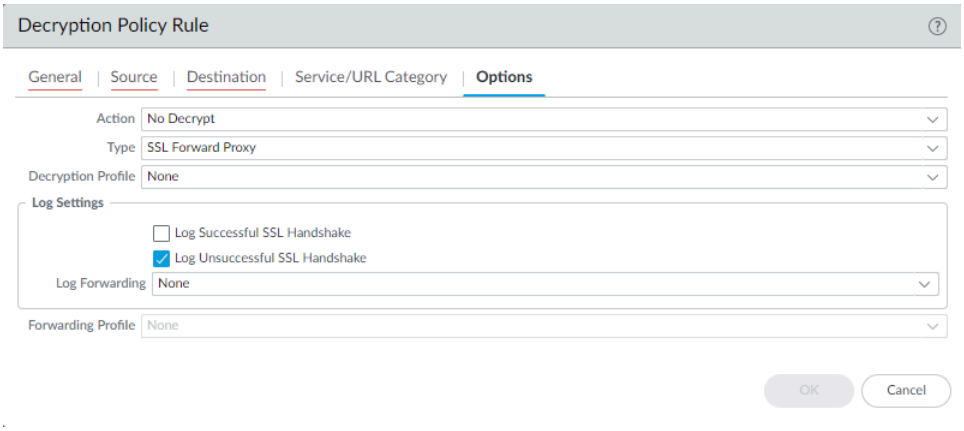
The firewall generates Decryption logs for sessions governed by a Decryption policy, including sessions with a No Decrypt policy. Configure Decryption logging in the Decryption policy that controls the traffic that you want to log.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 1 | Configure the Decryption traffic you want to log in Decryption policy (**Policies** > **Decryption**).

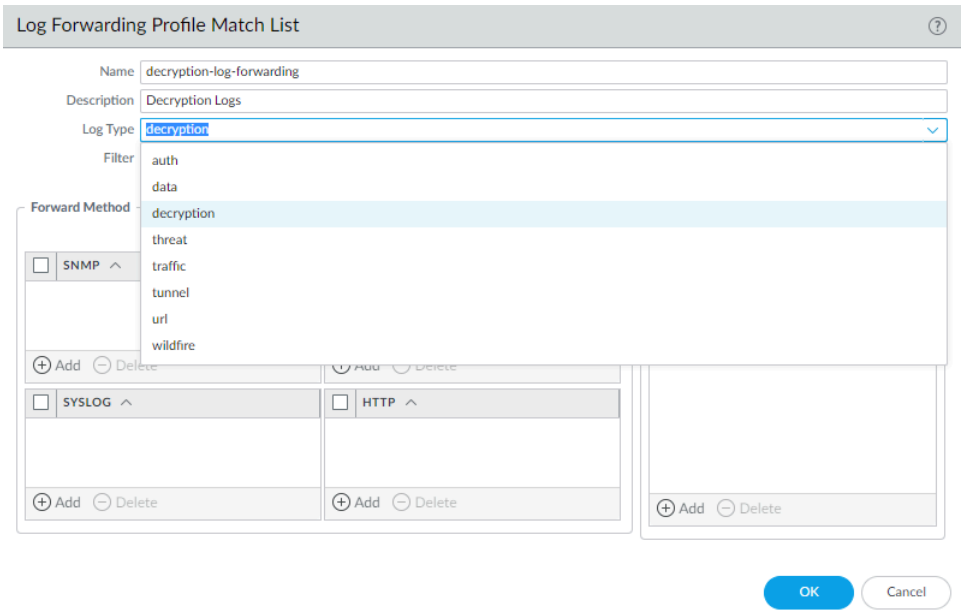
By default, the firewall logs only unsuccessful TLS handshakes:



Log successful handshakes as well as unsuccessful handshakes to gain visibility into as much decrypted traffic as your device's available resources permit (don't decrypt private or sensitive traffic; follow decryption best practices and decrypt as much traffic as you can).

STEP 2 | Create a Log Forwarding profile to forward Decryption logs to Log Collectors, other storage devices, or specific administrators and then specify the profile in the **Log Forwarding** field of the Decryption policy **Options** tab.

To forward Decryption logs, you must configure a Log Forwarding profile (**Objects** > **Log Forwarding**) to specify the Decryption **Log Type** and the method of forwarding the logs.

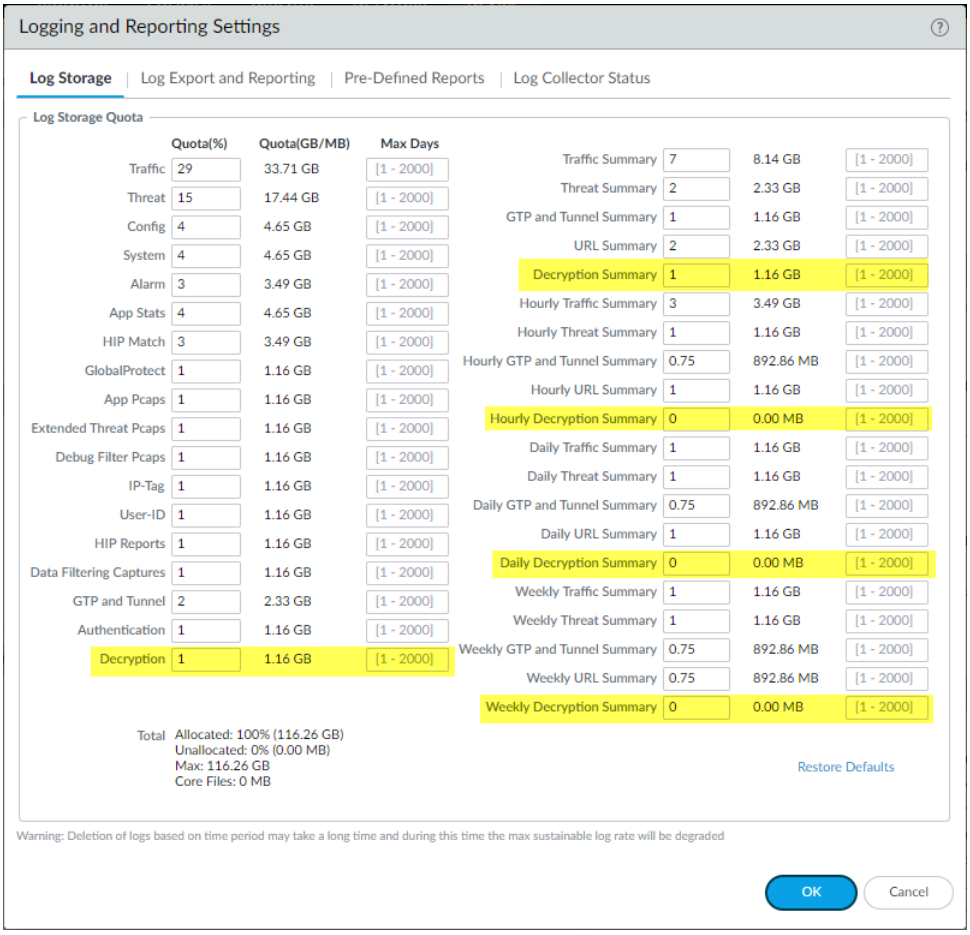


If you forward Decryption logs, be sure that the logs are stored securely because they contain sensitive information.

https://t.me/learningnets

STEP 3 | If you log successful TLS handshakes in addition to unsuccessful TLS handshakes, configure a larger log storage space quota (**Device** > **Setup** > **Management** > **Logging and Reporting Settings** > **Log Storage**) for Decryption logs on the firewall.

The default quota (allocation) is one percent of the device's log storage capacity for Decryption logs and one percent for the general decryption summary. There is no default allocation for hourly, daily, or weekly decryption summaries.



Many factors determine the amount of storage you may need for Decryption logs and they depend on your deployment. For example, take these factors into account:

- The amount of TLS traffic that passes through the firewall.
- The amount of TLS traffic that you decrypt.
- Your usage of other logs (evaluate from which logs you should take capacity to allocate to Decryption logs).
- If you log both successful and unsuccessful TLS handshakes, you probably need significantly more capacity than you need if you only log unsuccessful TLS handshakes. Depending on the amount of traffic you decrypt, Decryption logs could consume as much capacity as

https://t.me/learningnets

Traffic logs or Threat logs and may require a tradeoff among them if the device's capacity is already fully subscribed.

The total combined allocation of log quotas cannot exceed 100% of the available firewall log resources.

You may need to experiment to find the right quota for each log category in your particular deployment. If you only log unsuccessful handshakes, you could start with the default or increase the allocation to two or three percent. If you log both successful and unsuccessful handshakes, you could start by allocating about half of the space to Decryption logs that you allocate to Traffic logs. The logs from which you take the space to allocate to Decryption logs depends on your traffic, your business, and your monitoring requirements.

Decryption Log Errors, Error Indexes, and Bitmasks

The **Error Index** and **Error** columns in the Decryption log provide information about the decryption error category and details, respectively. You can also see error and error index information in the Handshake Details section of the Detailed Log View (click for any log entry). The Decryption log **Error Index** indicates one of eight error categories:

Error Index	Error (possible errors shown for the Error Index)
Certificate	Errors such as invalid certificates, expired certificates, unsupported client certificates, OCSP/CRL check revocations and failures, untrusted issuer CAs (sessions signed by an untrusted root, which includes incomplete certificate chains), and other certificate errors. *When the firewall doesn't have an intermediate certificate because the site did not send the full certificate chain, you can find and install the missing certificate to* Repair Incomplete Certificate Chains.
Cipher	Unsupported cipher errors where: • The client tries to negotiate a cipher that the firewall supports but that the Decryption profile applied to the traffic doesn't support. • The client tries to negotiate a cipher that the firewall doesn't support. • (Rare) Inbound Inspection is enabled and the server's capabilities don't match the Decryption profile settings. The error message includes the supported client cipher bitmask value and the supported Decryption profile cipher bitmask value. Use the bitmask values to identify the cipher the client tried to use and to list the cipher values that the Decryption profile supports as described later in this topic.
Feature	Errors such as oversized TLS handshakes or unknown handshakes, oversized certificate chains (more than five certificates), and other unsupported features.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Error Index	Error (possible errors shown for the Error Index)
HSM	Hardware storage module (HSM) errors such as unknown requests, items not found in the configuration, request timeouts, and other HSM errors and failures.
Protocol	Errors such as TLS handshake failures, private and public key mismatches, Heartbleed errors, TLS key exchange failures, and other TLS protocol errors. Protocol errors show when the server doesn't support the protocols that the client supports, the server uses certificate types that the firewall doesn't support, and general TLS protocol errors.
Resource	Errors such as lack of sufficient memory.
Resume	Session resumption errors concerning resume session IDs and tickets, resume session entries in the firewall cache, and other session resumption errors.
Version	Errors regarding client and Decryption profile version mismatches and client and server version mismatches. The error message includes bitmask values that identify the supported client and Decryption profile versions. Use the bitmask values to identify the cipher the client tried to use and to list the cipher values that the Decryption profile supports as described later in this topic.

*If no suitable error description category exists for an error, the default message is **General TLS protocol error**.*

Version and cipher log error information includes bitmask values that you convert to actual values using operational CLI commands:

- Version error bitmask values identify mismatches between the TLS protocol versions that the client and server use and also identify TLS protocol mismatches between the client and the Decryption profile applied to the traffic. The CLI command to convert version error bitmasks is:

  ```
  admin@vm1>debug dataplane show ssl-decrypt bitmask-version
   <bitmask-value>
  ```

  The command returns the TLS version that matches the bitmask.
- Cipher error bitmask values identify encryption and other mismatches between the client and the Decryption profile applied to the traffic.

  ```
  admin@vm1>debug dataplane show ssl-decrypt bitmask-cipher <bitmask-
  value>
  ```

  The command returns the cipher that matches the bitmask.

Filter the Decryption log to find version and cipher errors, plug the bitmask values for sessions with errors into the appropriate CLI command, obtain the values of the protocol version or cipher

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

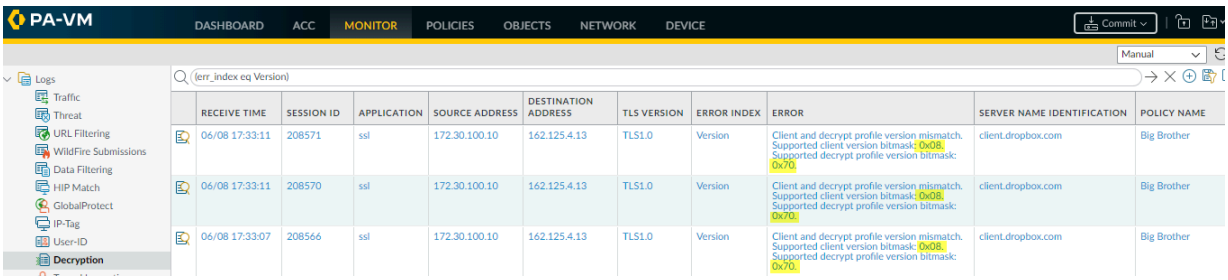
that caused the error, and use the information to update the Decryption policy or profile if you want to allow access to the site in question.

- Version Errors
- Cipher Errors
- Root Status “Uninspected”

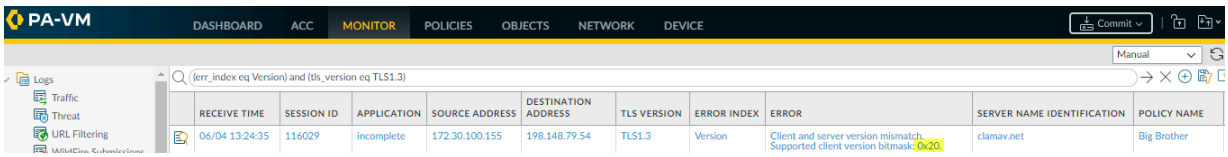
Version Errors

To identify and fix version mismatch errors:

1. Filter the Decryption Log to identify version errors using the filter **`(err_index eq Version)`**. The highlighted values are bitmask values:



   You can filter the Decryption log in many ways. For example, to see only TLSv1.3 version errors, use the filter **`(err_index eq Version) and (tls_version eq TLS1.3)`**:



2. Log in to the CLI and look up the bitmask values. The version errors in the first screen shot (the same errors for all three sessions) show an issue with a client and Decryption profile mismatch

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

—the supported client version bitmask is 0x08 and the supported Decryption profile version bitmask is 0x70:

```
admin@vm1>debug dataplane show ssl-decrypt bitmask-version 0x08
```

```
TLSv1.0
```

This output shows that the client supports only TLSv1.0.

```
admin@vm1>debug dataplane show ssl-decrypt bitmask-version 0x70
```

```
TLSv1.1
```

```
TSLv1.2
```

```
TLSv1.3
```

This output shows that the Decryption profile supports TLSv1.1, TLSv1.2, and TLSv1.3, but not TLSv1.0. Now you know the issue is that the client only supports a very old version of the TLS protocol and the Decryption profile attached to the Decryption policy rule that controls the traffic does not allow TLSv1.0 traffic.

The next thing to do is to decide what action to take. You could update the client so that it accepts a more secure TLS version. If the client requires TLSv1.0 for some reason, you can continue let the firewall continue to block the traffic, or you can update the Decryption profile to allow all TLSv1.0 traffic (not recommended), or you can create a Decryption policy and profile that allow TLSv1.0 and apply it only to the client devices that must use TLSv1.0 and cannot support a more secure protocol (most secure option for allowing the traffic).

The version error in the second screen shot shows a different issue: a client and server version mismatch. The error indicates the supported client bitmask as 0x20:

```
admin@vm1>debug dataplane show ssl-decrypt bitmask-version 0x20
```

```
TLSv1.2
```

The output shows that the client supports only TLSv1.2. Since the server does not support TLSv1.2, it may only support TLSv1.3 or it may support only TLSv1.1 or lower (less secure protocols). You can use Wireshark or another packet analysis tool to find out which version of TLS the server supports. Depending on what the server supports, you can:

- If the server only supports TLSv1.3, you could edit the Decryption profile so that it supports TLSv1.3.
- If the server only supports TLSv1.1 or lower, evaluate whether you need to access that server for business reasons. If not, consider blocking the traffic to increase security. If you need to access the server for business purposes, create or add the server to a Decryption

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

policy that applies only to the servers and sites you need to access for business; don't allow access to all servers that use less secure TLS versions.

3. To find the Decryption policy that controls the session traffic, check the **Policy Name** column in the log (or click the magnifying glass icon next to the Decryption log to see the information in the General section of the Detailed Log View). In the example above, the Decryption policy name is Big Brother. To find the Decryption policy and profile, go to **Policies** > **Decryption**, select the policy named Big Brother, and then select the **Options** tab. **Decryption profile** displays the name of the Decryption profile.

   Go to **Objects** > **Decryption** > **Decryption Profile**, select the appropriate Decryption profile, and edit it to address the version issue.

Cipher Errors

Using the Decryption log to hunt down cipher errors is similar to hunting down version errors —you filter the log to find errors and obtain error bitmasks. Then you go to the CLI, convert the bitmask to the error value, and then take appropriate action to fix the issue. For example:

1. Filter the Decryption Log to identify cipher errors using the filter **`(err_index eq Cipher)`**. For example, let's examine a cipher error with the **Error** message `Unsupported cipher. Supported client cipher bitmask: 0x80000000. Support decrypt profile cipher bitmask 0x60f79980.`
2. Log in to the CLI and look up the bitmask values:

```
admin@vm1>debug dataplane show ssl-decrypt bitmask-cipher
 0x80000000
```

```
CHACHA_PLY1305_SHA256
```

This output shows that client tried to negotiate a cipher that the firewall supports (if the bitmask is all zeros (`0x0000000`, then the client tried to negotiate a cipher that the firewall doesn't support):

```
admin@vm1>debug dataplane show ssl-decrypt bitmask-cipher
 0x80000000
```

```
TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256
TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256
TLS_RSA_WITH_AES_256_CBC_SHA256
TLS_RSA_WITH_AES_128_CBC_SHA256
TLS_DHE_RSA_WITH_AES_256_CBC_SHA256
TLS_DHE_RSA_WITH_AES_128_CBC_SHA256
TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA
TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA
TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA
TLS_RSA_WITH_AES_256_CBC_SHA
TLS_RSA_WITH_AES_128_CBC_SHA
TLS_DHE_RSA_WITH_AES_256_CBC_SHA
TLS13_WITH_AES_256_GCM_SHA384
```


https://t.me/learningnets

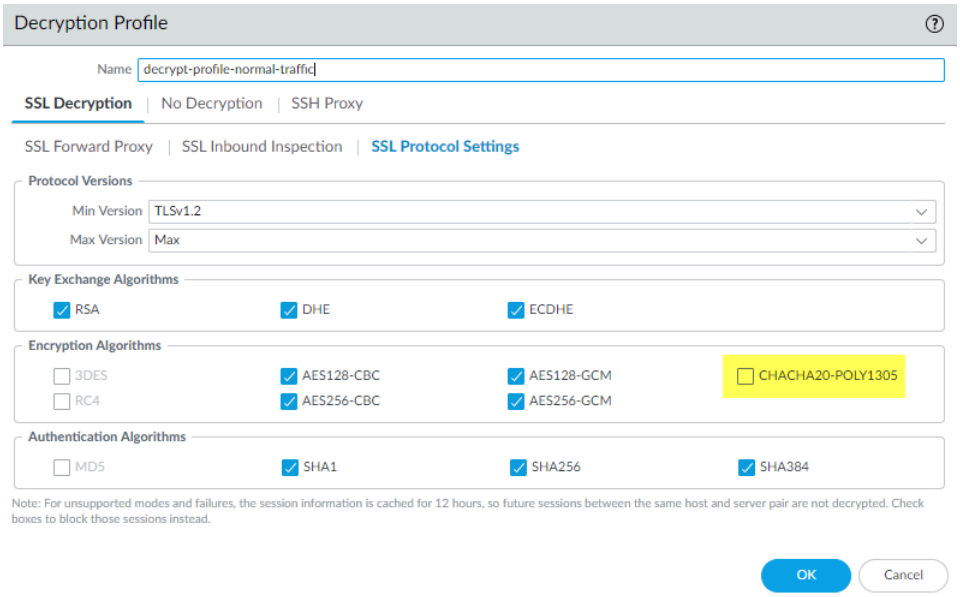
```
TLS13_WITH_AES_128_GCM_SHA256
```

This output shows that the Decryption profile that controls the traffic supports many ciphers, but does not support the cipher the client is trying to use.

To fix this issue so that the firewall allows and decrypts the traffic, you need to add support for the missing cipher to the Decryption profile.

3. Check the Decryption log or the Detailed Log View **Policy Name** to get the name of the Decryption policy that controls the traffic. Go to **Policies** > **Decryption** and select the policy. On the **Options** tab, look up the name of the Decryption profile. Next, Go to **Objects** > **Decryption** > **Decryption Profile**, select the appropriate Decryption profile, and edit it to address the version issue.

   In this example, the Decryption profile does not support the TLS13_WITH_CHACHA_POLY1305_SHA256 cipher, so the client can't connect:



   To fix the issue, select the **CHACHA20-POLY1305** encryption algorithm option (the **Max Version** setting of **Max** means that the profile already supports TLSv1.3 and the Authentication Algorithm setting already includes SHA256, so only the encryption algorithm support was missing) and then **Commit** the configuration. After you commit the configuration, the Decryption profile supports the missing cipher and the decryption sessions for the traffic succeed.

   If the firewall does not support a cipher suite and you need to allow the traffic for business purposes, create a Decryption policy and profile that applies only to that traffic. In the Decryption profile, disable the ***Block sessions with unsupported cipher suites*** *option*

   .

Root Status "Uninspected"

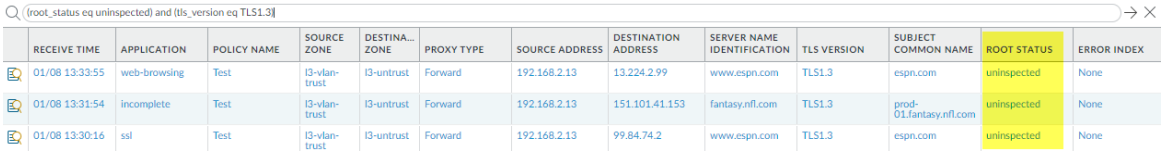
In some cases, the **Root Status** column displays the value **uninspected**. There are a number of reasons why the firewall could not inspect the root status, including:

©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

- Session resumption.
- Traffic was not decrypted because a No Decryption policy controlled the traffic, so the firewall did not decrypt the traffic.
- A decryption failure occurred before the firewall could inspect the server certificate.

Filter the Decryption Log **`(root_status eq uninspected) and (tls_version eq TLS1.3)`** to see Decryption sessions for which the Root Status is uninspected:



RECEIVE TIME	APPLICATION	POLICY NAME	SOURCE ZONE	DESTINA... ZONE	PROXY TYPE	SOURCE ADDRESS	DESTINATION ADDRESS	SERVER NAME IDENTIFICATION	TLS VERSION	SUBJECT COMMON NAME	ROOT STATUS	ERROR INDEX
01/08 13:33:55	web-browsing	Test	l3-vlan-trust	l3-untrust	Forward	192.168.2.13	13.224.2.99	www.espn.com	TLS1.3	espn.com	uninspected	None
01/08 13:31:54	incomplete	Test	l3-vlan-trust	l3-untrust	Forward	192.168.2.13	151.101.41.153	fantasy.nfl.com	TLS1.3	prod-01.fantasy.nfl.com	uninspected	None
01/08 13:30:16	ssl	Test	l3-vlan-trust	l3-untrust	Forward	192.168.2.13	99.84.74.2	www.espn.com	TLS1.3	espn.com	uninspected	None

Repair Incomplete Certificate Chains

Not all websites send their complete certificate chain even though the RFC 5246 TLSv1.2 standard requires authenticated servers to provide a valid certificate chain leading to an acceptable certificate authority. When you enable decryption and apply a Forward Proxy Decryption profile that enables **Block sessions with untrusted issuers** in the Decryption policy, if an intermediate certificate is missing from the certificate list the website's server presents to the firewall, the firewall can't construct the certificate chain to the top (root) certificate. In these cases, the firewall presents its Forward Untrust Certificate to the client because the firewall cannot construct the chain to the root certificate and trust cannot be established without the missing intermediate certificate.

The firewall only has root certificates in its Default Trusted Certificate Authorities store.

If a website you need to communicate with for business purposes has one or more missing intermediate certificates and the Decryption profile blocks sessions with untrusted issuers, then you can find and download the missing intermediate certificate and install it on the firewall as a Trusted Root CA so that the firewall trusts the site's server. (The alternative is to contact the website owner and ask them to configure their server so that it sends the intermediate certificate during the handshake.)

If you allow sessions with untrusted issuers in the Decryption profile, the firewall establishes sessions even if the issuer is untrusted; however, it is a best practice to block sessions with untrusted issuers for better security.

STEP 1 | Find websites that cause incomplete certificate chain errors.

1. Filter the Decryption log to identify Decryption sessions that failed because of an incomplete certificate chain.

   In the filter field, type the query **`(err_index eq Certificate) and (error contains 'http')`**. This query filters the logs for Certificate errors that contain the string "http", which finds all of the error entries that contain the CA Issuer URL (often

https://t.me/learningnets

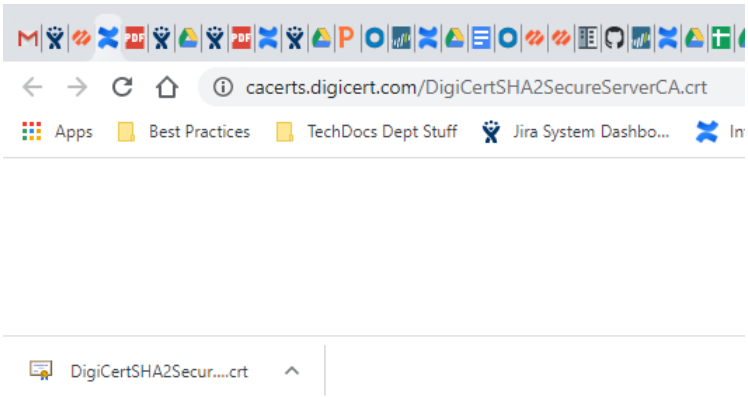
called the URI). The CA Issuer URL is the Authority Information Access (AIA) information for the CA Issuer.

2. Click an **Error** column entry that begins “Received fatal alert UnknownCA from client. CA Issuer URL:” followed by the URI.

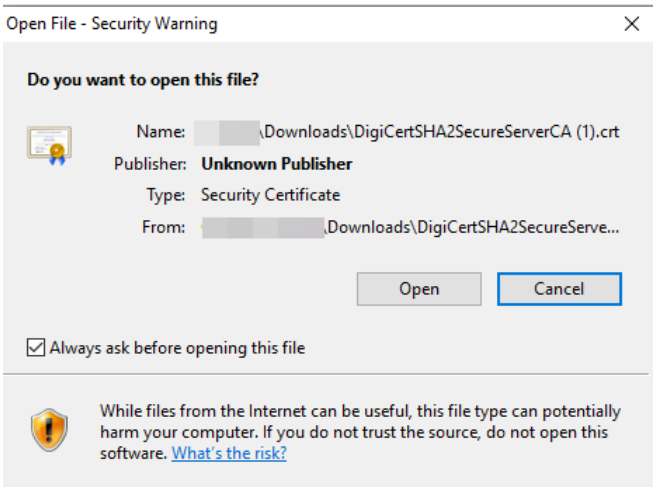


The firewall automatically adds the selected error to the query and shows the full URI path (the full URI path may be truncated in the **Error** column).

STEP 2 | Copy and paste the URI into your browser and then press Enter to download the missing intermediate certificate.

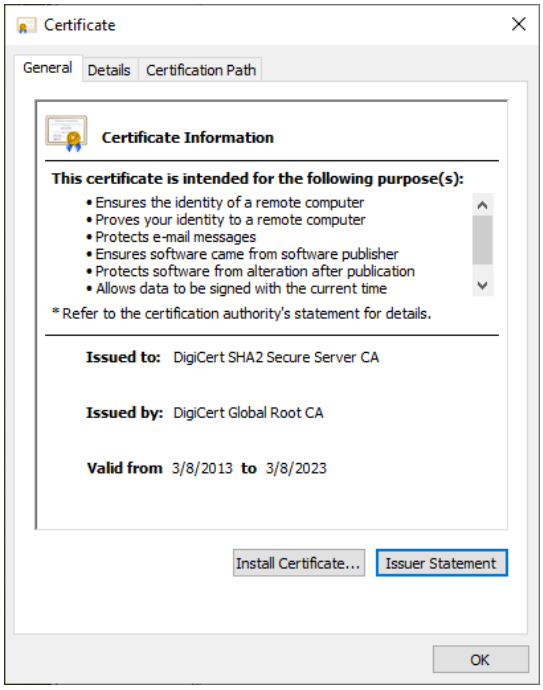


STEP 3 | Click the certificate to open the dialog box.

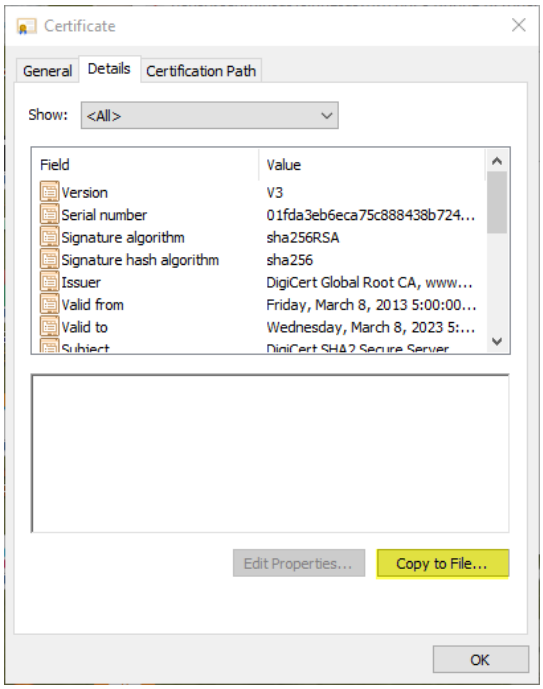


https://t.me/learningnets

STEP 4 | Click **Open** to open the certificate file.



STEP 5 | Select the **Details** tab and then click **Copy to File...**.



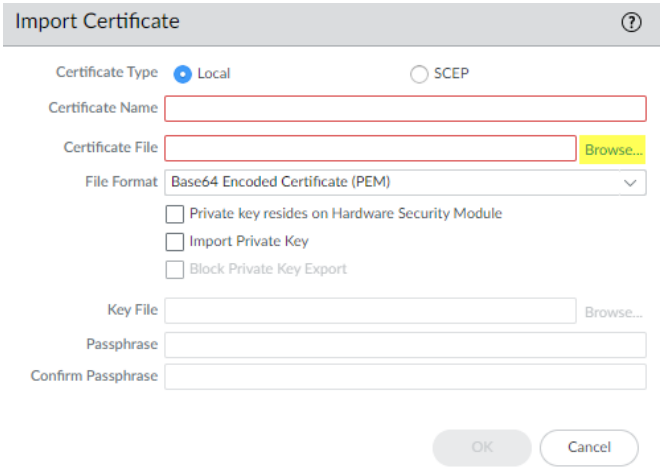
Follow the export directions. The certificate is copied to the folder you designated as you default download folder.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 6 | Import the certificate into the firewall.

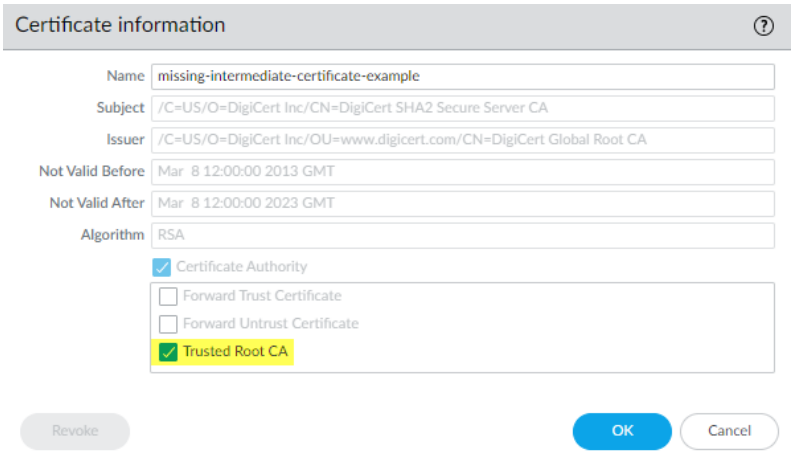
1. Navigate to **Device** > **Certificate Management** > **Certificates** and then select **Import**.
2. **Browse** to the folder where you stored the missing intermediate certificate and select it. Leave the **File Format** as **Base64 Encoded Certificate (PEM)**.



3. Name the certificate and specify any other options you want to use, then click **OK**.

STEP 7 | When the certificate has imported, select the certificate from the **Device Certificates** list to open the Certificate Information dialog.

STEP 8 | Select **Trusted Root CA** to mark the certificate as a Trusted Root CA on the firewall and then click **OK**.



In **Device** > **Certificate Management** > **Certificates** > **Device Certificates**, the imported certificate now appears in the list of certificates. Check the **Usage** column to confirm that the status is **Trusted Root CA Certificate** to verify that the firewall considers the certificate to be a trusted root CA.

STEP 9 | **Commit** the configuration.

STEP 10 | You have now repaired the broken certificate chain.

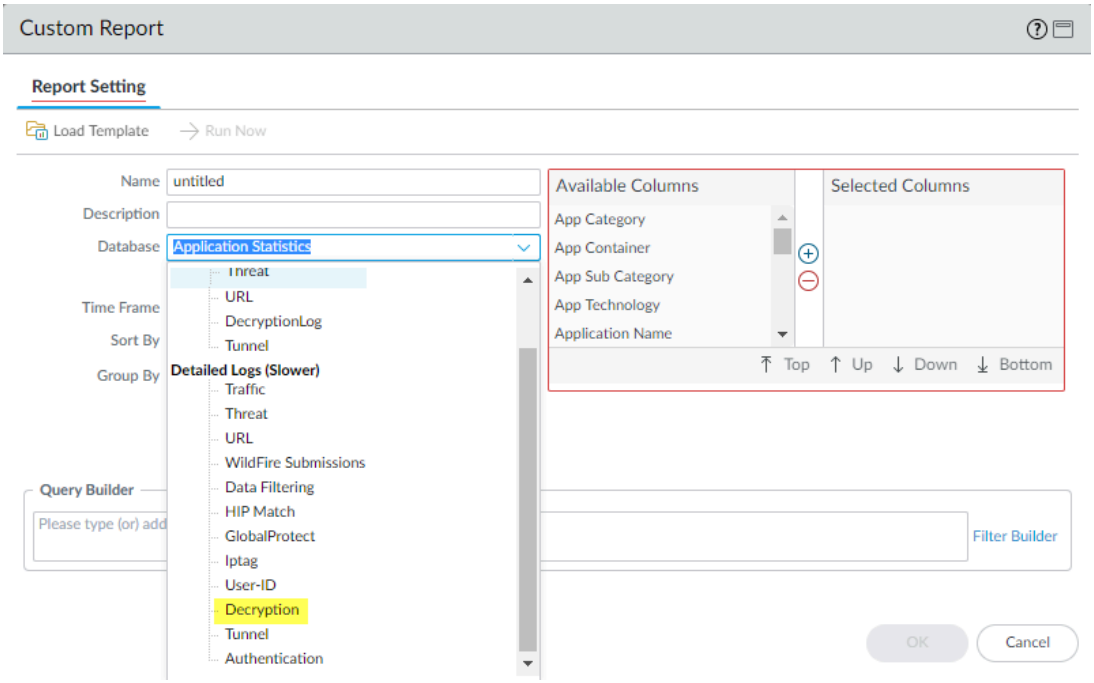
The firewall doesn't block the traffic because the CA issuer is not untrusted anymore. Repeat this process for all missing intermediate certificates to repair their certificate chains.

https://t.me/learningnets

Custom Report Templates for Decryption

You can create Custom Reports and generate them for decryption events based on Decryption log fields and custom templates. Select log fields to include in custom reports and select templates to refine the log query:

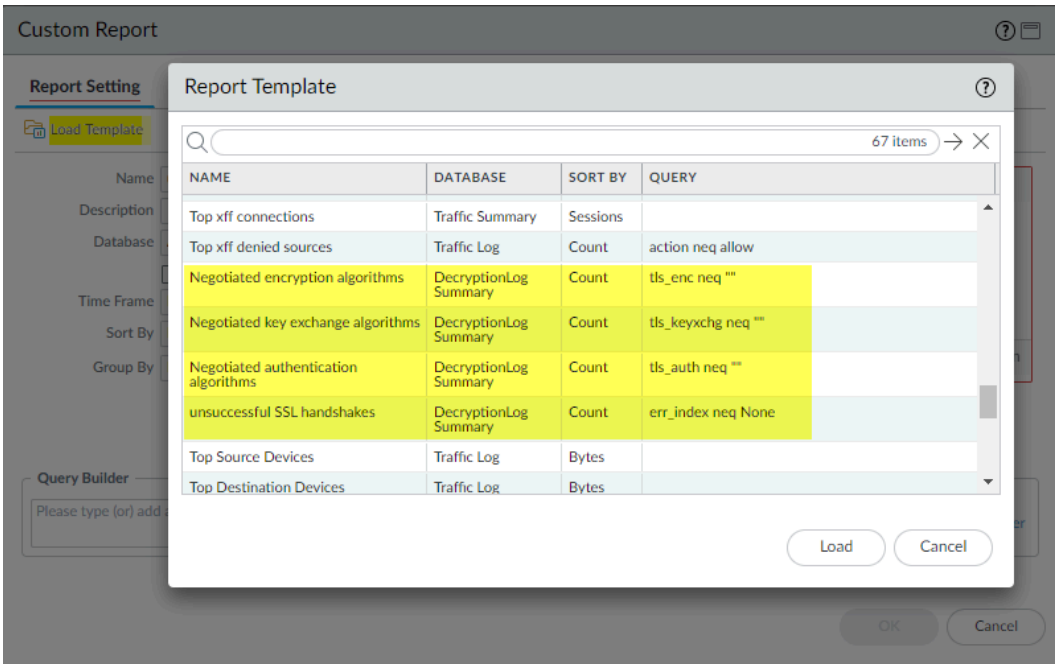
1. **Monitor** > **Manage Custom Reports**.
2. **Add** a custom report.
3. To configure the Decryption log fields to use in the custom report, select **Decryption** as the **Database**.



The **Available Columns** list changes to match the columns available in the Decryption log. Select and add the columns (information) that you want to include in the custom report. If you don't want to refine the custom report any further, click **OK** to generate the report.

https://t.me/learningnets

4. If desired, refine the output of the custom Decryption report using the Query Builder and the four templates introduced in PAN-OS 10.0. To select a template to filter the report output, click **Load Template** and select from the four Decryption templates:



The **Query** column shows the filter query that each template represents. **Load** the desired query and then click **OK** to generate the custom report.

Unsupported Parameters by Proxy Type and TLS Version

Decryption Log fields display decryption session parameters for each decryption proxy type. However, for reasons such as version support, encrypted portions of TLS handshakes, information availability, etc., some parameters are not available for every proxy type or TLS version. The following table shows unsupported Decryption log parameters by proxy type and TLS version.

Proxy Type	Unsupported Parameter	TLS Version
Forward Proxy	Negotiated EC Curve	TLSv1.3
Inbound Inspection	Server Name Identification Root Common Name	All
	Negotiated EC Curve	TLSv1.3
No Decrypt (**No Decrypt** action in the Decryption policy rule)	Negotiated EC Curve Server Name Identification	TLSv1.2
	Negotiated EC Curve Server Name Identification	TLSv1.3

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Proxy Type	Unsupported Parameter	TLS Version
	Certificate Information (all certificate information fields, for example, Certificate Start Date, Certificate End Date, Certificate Key Type, etc.)	
Network Packet Broker	Negotiated EC Curve	TLSv1.3
GlobalProtect Portal	Server Name Identification Root Common Name Decryption policy name App-ID	All
GlobalProtect Gateway	Server Name Identification Decryption policy name App-ID	All
Clientless SSLVPN	Server Name Identification	All
SSH	Decryption Log Not Supported	
Cleartext	Decryption Log Not Supported	

Decryption Troubleshooting Workflow Examples

The Decryption Log and the SSL Activity widgets in the Application Command Center (ACC) provide powerful Decryption troubleshooting tools that work both independently and together. When you gain an understanding of how to use these tools, you can investigate and address a wide range of decryption issues.

The following examples show you how to use the troubleshooting tools to identify, investigate, and address decryption issues. Apply these methods to troubleshoot any issues you encounter in your decryption deployment.

- Investigate Decryption Failure Reasons
- Troubleshoot Unsupported Cipher Suites
- Identify Weak Protocols and Cipher Suites
- Identify Untrusted CA Certificates
- Troubleshoot Expired Certificates
- Troubleshoot Revoked Certificates
- Troubleshoot Pinned Certificates

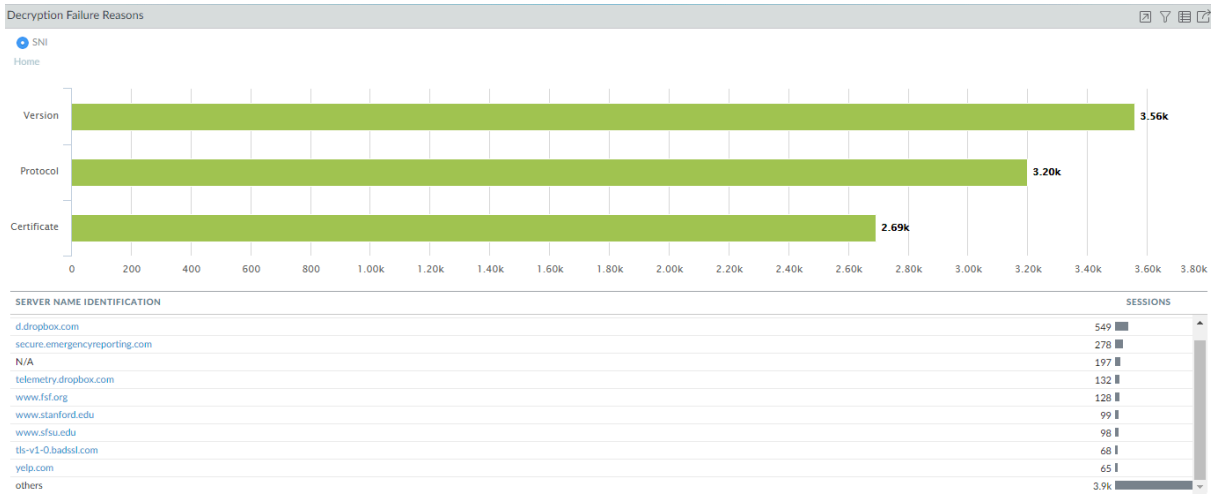
 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Investigate Decryption Failure Reasons

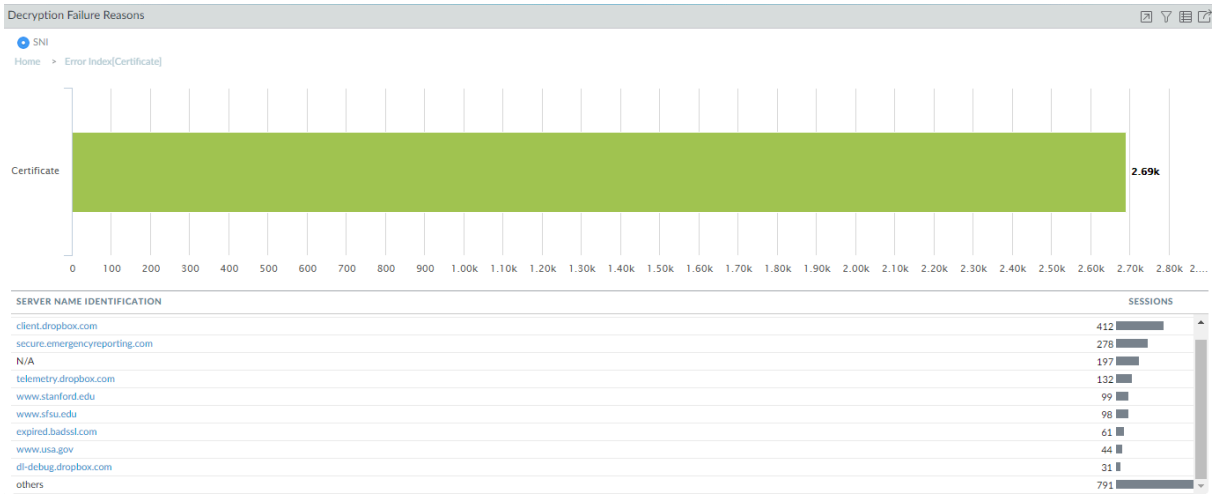
The most common reasons for decryption failures are TLS protocol errors, cipher version errors (client and server version mismatches and also client and Decryption profile version mismatches), and certificate errors. To investigate decryption errors, start with the Application Command Center (ACC) to identify failures and then go to the Decryption logs to drill down into details.

STEP 1 | Begin your investigation at **ACC** > **SSL Activity** and look at the Decryption Failure Reasons widget.



In this example, we investigate certificate errors. You can use the same process to investigate version and protocol errors.

STEP 2 | Click the green bar next to **Certificate** to see which hosts (SNIs) experienced certificate errors and see a list of hosts that experienced the largest number of certificate errors.

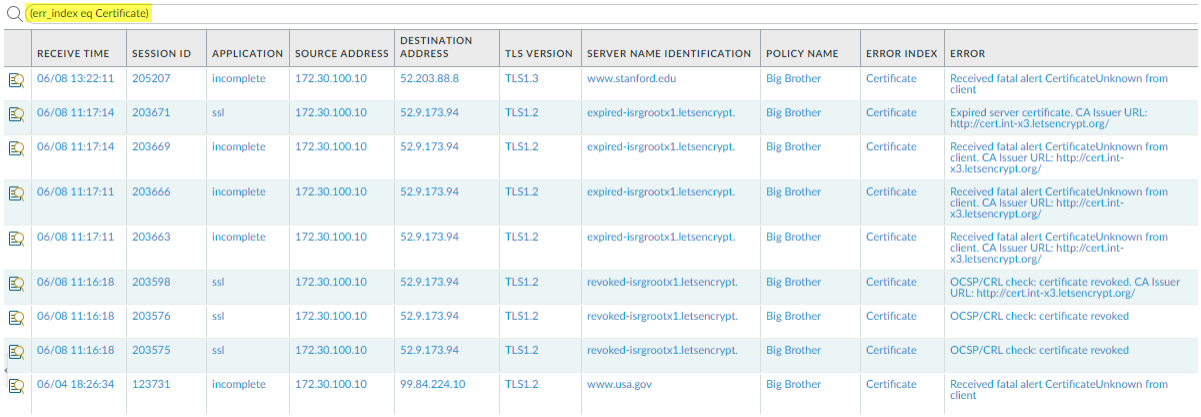


 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 3 | Go to **Monitor** > **Logs** > **Decryption** to drill down into the logs.

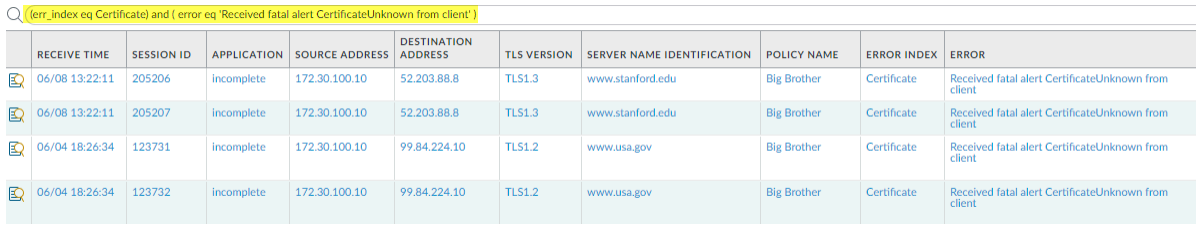
Use the query **`(err_index eq Certificate)`** to filter the Decryption logs to view all Decryption sessions that experienced certificate errors.



(err_index eq Certificate)

RECEIVE TIME	SESSION ID	APPLICATION	SOURCE ADDRESS	DESTINATION ADDRESS	TLS VERSION	SERVER NAME IDENTIFICATION	POLICY NAME	ERROR INDEX	ERROR
06/08 13:22:11	205207	incomplete	172.30.100.10	52.203.88.8	TLS1.3	www.stanford.edu	Big Brother	Certificate	Received fatal alert CertificateUnknown from client
06/08 11:17:14	203671	ssl	172.30.100.10	52.9.173.94	TLS1.2	expired-isrgrootx1.letsencrypt.	Big Brother	Certificate	Expired server certificate. CA Issuer URL: http://cert.int-x3.letsencrypt.org/
06/08 11:17:14	203669	incomplete	172.30.100.10	52.9.173.94	TLS1.2	expired-isrgrootx1.letsencrypt.	Big Brother	Certificate	Received fatal alert CertificateUnknown from client. CA Issuer URL: http://cert.int-x3.letsencrypt.org/
06/08 11:17:11	203666	incomplete	172.30.100.10	52.9.173.94	TLS1.2	expired-isrgrootx1.letsencrypt.	Big Brother	Certificate	Received fatal alert CertificateUnknown from client. CA Issuer URL: http://cert.int-x3.letsencrypt.org/
06/08 11:17:11	203663	incomplete	172.30.100.10	52.9.173.94	TLS1.2	expired-isrgrootx1.letsencrypt.	Big Brother	Certificate	Received fatal alert CertificateUnknown from client. CA Issuer URL: http://cert.int-x3.letsencrypt.org/
06/08 11:16:18	203598	ssl	172.30.100.10	52.9.173.94	TLS1.2	revoked-isrgrootx1.letsencrypt.	Big Brother	Certificate	OCSP/CRL check: certificate revoked. CA Issuer URL: http://cert.int-x3.letsencrypt.org/
06/08 11:16:18	203576	ssl	172.30.100.10	52.9.173.94	TLS1.2	revoked-isrgrootx1.letsencrypt.	Big Brother	Certificate	OCSP/CRL check: certificate revoked
06/08 11:16:18	203575	ssl	172.30.100.10	52.9.173.94	TLS1.2	revoked-isrgrootx1.letsencrypt.	Big Brother	Certificate	OCSP/CRL check: certificate revoked
06/04 18:26:34	123731	incomplete	172.30.100.10	99.84.224.10	TLS1.2	www.usa.gov	Big Brother	Certificate	Received fatal alert CertificateUnknown from client

The **Error** column shows the reason for the certificate error. To filter for all Decryption sessions that had the same error, click the error message to add it to the query and then execute the query. For example, to find all errors based on receiving a fatal alert from the client, clicking the error produces the query **`(err_index eq Certificate) and (error eq 'Received fatal alert CertificateUnknown from client')`**:



(err_index eq Certificate) and (error eq 'Received fatal alert CertificateUnknown from client')

RECEIVE TIME	SESSION ID	APPLICATION	SOURCE ADDRESS	DESTINATION ADDRESS	TLS VERSION	SERVER NAME IDENTIFICATION	POLICY NAME	ERROR INDEX	ERROR
06/08 13:22:11	205206	incomplete	172.30.100.10	52.203.88.8	TLS1.3	www.stanford.edu	Big Brother	Certificate	Received fatal alert CertificateUnknown from client
06/08 13:22:11	205207	incomplete	172.30.100.10	52.203.88.8	TLS1.3	www.stanford.edu	Big Brother	Certificate	Received fatal alert CertificateUnknown from client
06/04 18:26:34	123731	incomplete	172.30.100.10	99.84.224.10	TLS1.2	www.usa.gov	Big Brother	Certificate	Received fatal alert CertificateUnknown from client
06/04 18:26:34	123732	incomplete	172.30.100.10	99.84.224.10	TLS1.2	www.usa.gov	Big Brother	Certificate	Received fatal alert CertificateUnknown from client

To filter for the certificate errors that a specific host received, add that SNI to the query instead of adding error message text. For example, to find all certificate errors for expired.badssl.comm use the query **`(err_index eq Certificate) and (sni eq 'expired.badssl.com')`**:



(err_index eq Certificate) and (sni eq 'expired.badssl.com')

RECEIVE TIME	SESSION ID	APPLICATION	SOURCE ADDRESS	DESTINATION ADDRESS	TLS VERSION	SERVER NAME IDENTIFICATION	POLICY NAME	ERROR INDEX	ERROR
06/02 17:17:20	12959	ssl	172.30.100.10	104.154.89.105	TLS1.2	expired.badssl.com	Big Brother	Certificate	Expired server certificate. CA Issuer URL: htt
06/02 17:17:19	12957	ssl	172.30.100.10	104.154.89.105	TLS1.2	expired.badssl.com	Big Brother	Certificate	Expired server certificate. CA Issuer URL: htt
06/02 17:17:19	12955	ssl	172.30.100.10	104.154.89.105	TLS1.2	expired.badssl.com	Big Brother	Certificate	Expired server certificate. CA Issuer URL: htt
06/02 17:17:19	12958	incomplete	172.30.100.10	104.154.89.105	TLS1.2	expired.badssl.com	Big Brother	Certificate	Received fatal alert CertificateUnknown from client. CA Issuer URL: htt
06/02 17:17:18	12956	incomplete	172.30.100.10	104.154.89.105	TLS1.2	expired.badssl.com	Big Brother	Certificate	Received fatal alert CertificateUnknown from client. CA Issuer URL: htt
06/02 17:17:18	12951	incomplete	172.30.100.10	104.154.89.105	TLS1.2	expired.badssl.com	Big Brother	Certificate	Received fatal alert CertificateUnknown from client. CA Issuer URL: htt
06/02 17:11:48	12802	ssl	172.30.100.10	104.154.89.105	TLS1.2	expired.badssl.com	Big Brother	Certificate	Expired server certificate. CA Issuer URL: htt

The **Error** column shows the specific reason for each certificate error associated with expired.badssl.com.

Once you know the reason for the certificate issue that caused the decryption failure, you can address it. For example, if the certificate chain is incomplete, you can repair the incomplete

https://t.me/learningnets

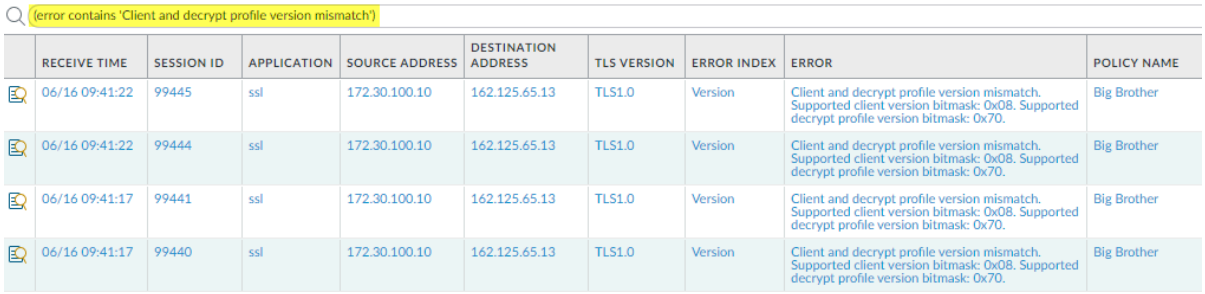
certificate chain. If a certificate is expired, you can notify the site administrator or create a policy-based exception if you need to access the site.

Troubleshoot Unsupported Cipher Suites

Identifying and troubleshooting unsupported cipher suites in the Decryption log is an aspect of version error investigation that is worth examining on its own.

STEP 1 | In the Decryption log (**Monitor** > **Logs** > **Decryption**), use the query **`(error contains 'Client and decrypt profile mismatch'`** to identify all cipher suite version mismatches.

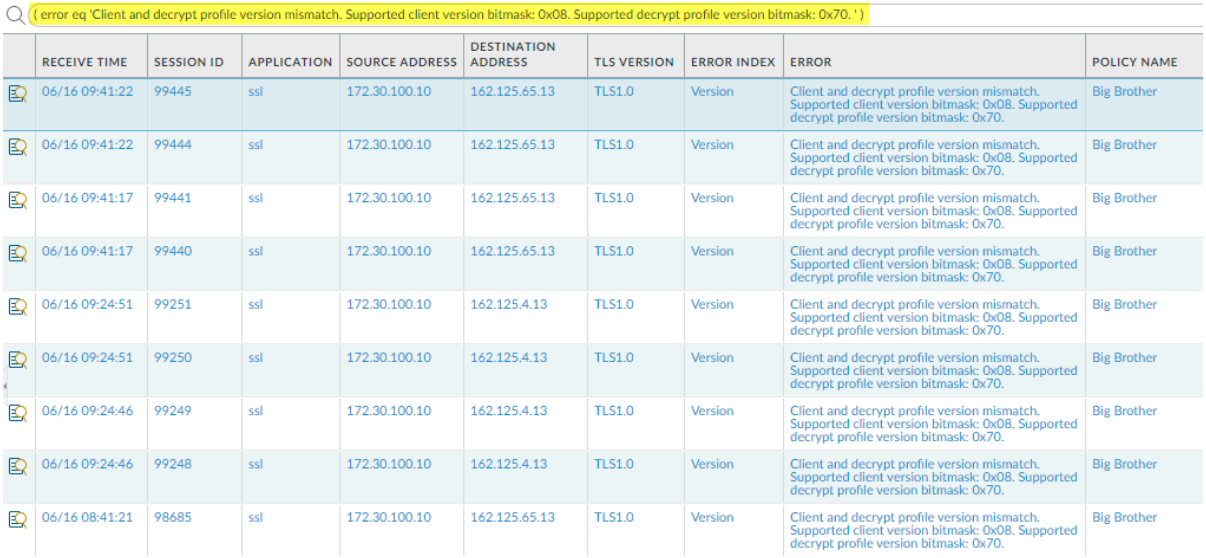
Filtering the logs for these mismatches identifies finds all instances where the client and the Decryption profile cipher suite support don't match.



(error contains 'Client and decrypt profile version mismatch')

RECEIVE TIME	SESSION ID	APPLICATION	SOURCE ADDRESS	DESTINATION ADDRESS	TLS VERSION	ERROR INDEX	ERROR	POLICY NAME
06/16 09:41:22	99445	ssl	172.30.100.10	162.125.65.13	TLS1.0	Version	Client and decrypt profile version mismatch. Supported client version bitmask: 0x08. Supported decrypt profile version bitmask: 0x70.	Big Brother
06/16 09:41:22	99444	ssl	172.30.100.10	162.125.65.13	TLS1.0	Version	Client and decrypt profile version mismatch. Supported client version bitmask: 0x08. Supported decrypt profile version bitmask: 0x70.	Big Brother
06/16 09:41:17	99441	ssl	172.30.100.10	162.125.65.13	TLS1.0	Version	Client and decrypt profile version mismatch. Supported client version bitmask: 0x08. Supported decrypt profile version bitmask: 0x70.	Big Brother
06/16 09:41:17	99440	ssl	172.30.100.10	162.125.65.13	TLS1.0	Version	Client and decrypt profile version mismatch. Supported client version bitmask: 0x08. Supported decrypt profile version bitmask: 0x70.	Big Brother

To find all Decryption sessions that experienced the same error, click the error message to add it to the query and remove the original query, for example:



(error eq 'Client and decrypt profile version mismatch. Supported client version bitmask: 0x08. Supported decrypt profile version bitmask: 0x70. ')

RECEIVE TIME	SESSION ID	APPLICATION	SOURCE ADDRESS	DESTINATION ADDRESS	TLS VERSION	ERROR INDEX	ERROR	POLICY NAME
06/16 09:41:22	99445	ssl	172.30.100.10	162.125.65.13	TLS1.0	Version	Client and decrypt profile version mismatch. Supported client version bitmask: 0x08. Supported decrypt profile version bitmask: 0x70.	Big Brother
06/16 09:41:22	99444	ssl	172.30.100.10	162.125.65.13	TLS1.0	Version	Client and decrypt profile version mismatch. Supported client version bitmask: 0x08. Supported decrypt profile version bitmask: 0x70.	Big Brother
06/16 09:41:17	99441	ssl	172.30.100.10	162.125.65.13	TLS1.0	Version	Client and decrypt profile version mismatch. Supported client version bitmask: 0x08. Supported decrypt profile version bitmask: 0x70.	Big Brother
06/16 09:41:17	99440	ssl	172.30.100.10	162.125.65.13	TLS1.0	Version	Client and decrypt profile version mismatch. Supported client version bitmask: 0x08. Supported decrypt profile version bitmask: 0x70.	Big Brother
06/16 09:24:51	99251	ssl	172.30.100.10	162.125.4.13	TLS1.0	Version	Client and decrypt profile version mismatch. Supported client version bitmask: 0x08. Supported decrypt profile version bitmask: 0x70.	Big Brother
06/16 09:24:51	99250	ssl	172.30.100.10	162.125.4.13	TLS1.0	Version	Client and decrypt profile version mismatch. Supported client version bitmask: 0x08. Supported decrypt profile version bitmask: 0x70.	Big Brother
06/16 09:24:46	99249	ssl	172.30.100.10	162.125.4.13	TLS1.0	Version	Client and decrypt profile version mismatch. Supported client version bitmask: 0x08. Supported decrypt profile version bitmask: 0x70.	Big Brother
06/16 09:24:46	99248	ssl	172.30.100.10	162.125.4.13	TLS1.0	Version	Client and decrypt profile version mismatch. Supported client version bitmask: 0x08. Supported decrypt profile version bitmask: 0x70.	Big Brother
06/16 08:41:21	98685	ssl	172.30.100.10	162.125.65.13	TLS1.0	Version	Client and decrypt profile version mismatch. Supported client version bitmask: 0x08. Supported decrypt profile version bitmask: 0x70.	Big Brother

The hexadecimal codes identify the exact version that the client supports and the exact version that the Decryption profile supports.

https://t.me/learningnets

STEP 2 | Log in to the CLI and look up the bitmask values.

The errors show a client and Decryption profile mismatch. The supported client bitmask is 0x08 and the supported Decryption profile bitmask is 0x70:

```
admin@vm1>debug dataplane show ssl-decrypt bitmask-version 0x08
```

```
TLSv1.0
```

This output shows that the client supports only TLSv1.0.

```
admin@vm1>debug dataplane show ssl-decrypt bitmask-version 0x70
```

```
TLSv1.1
```

```
TSLv1.2
```

```
TLSv1.3
```

This output shows that the Decryption profile supports TLSv1.1, TLSv1.2, and TLSv1.3, but not TLSv1.0. Now you know that the client only supports an old version of the TLS protocol and the Decryption profile attached to the Decryption policy rule that controls the traffic does not allow that version.

STEP 3 | Decide what action to take.

You could update the client so that it accepts a more secure TLS version. If the client requires TLSv1.0 for some reason, you can continue let the firewall continue to block the traffic, or you can update the Decryption profile to allow all TLSv1.0 traffic (not recommended), or you can create a Decryption policy and profile that allow TLSv1.0 and apply it only to the client devices that must use TLSv1.0 and cannot support a more secure protocol (most secure option for allowing the traffic).

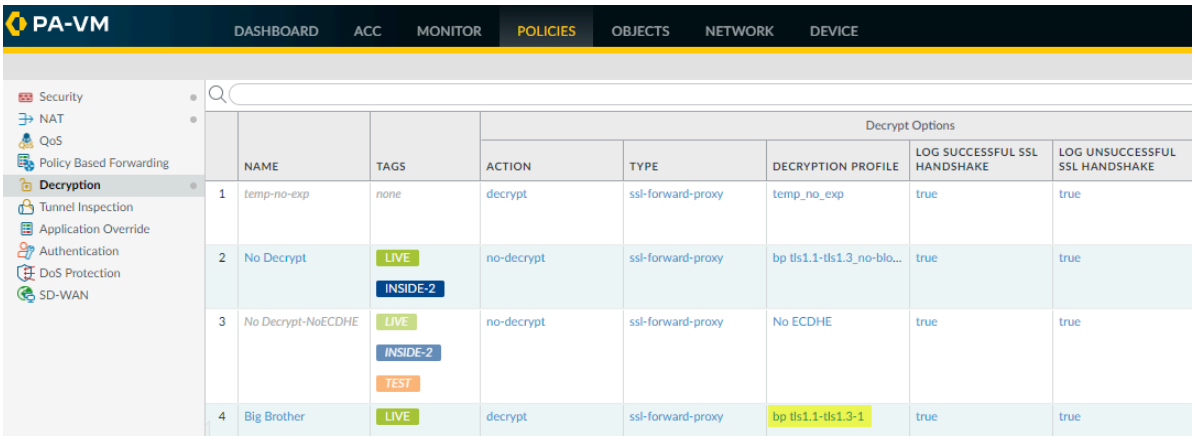
STEP 4 | If you choose to edit the Decryption profile, to find the Decryption policy that controls the session traffic, check the **Policy Name** column in the log (or click the magnifying glass icon

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

next to the Decryption log to see the information in the General section of the Detailed Log View).

1. In this example, the Decryption policy name is Big Brother; to find the Decryption profile, go to **Policies** > **Decryption** and check the **Decryption Profile** column.

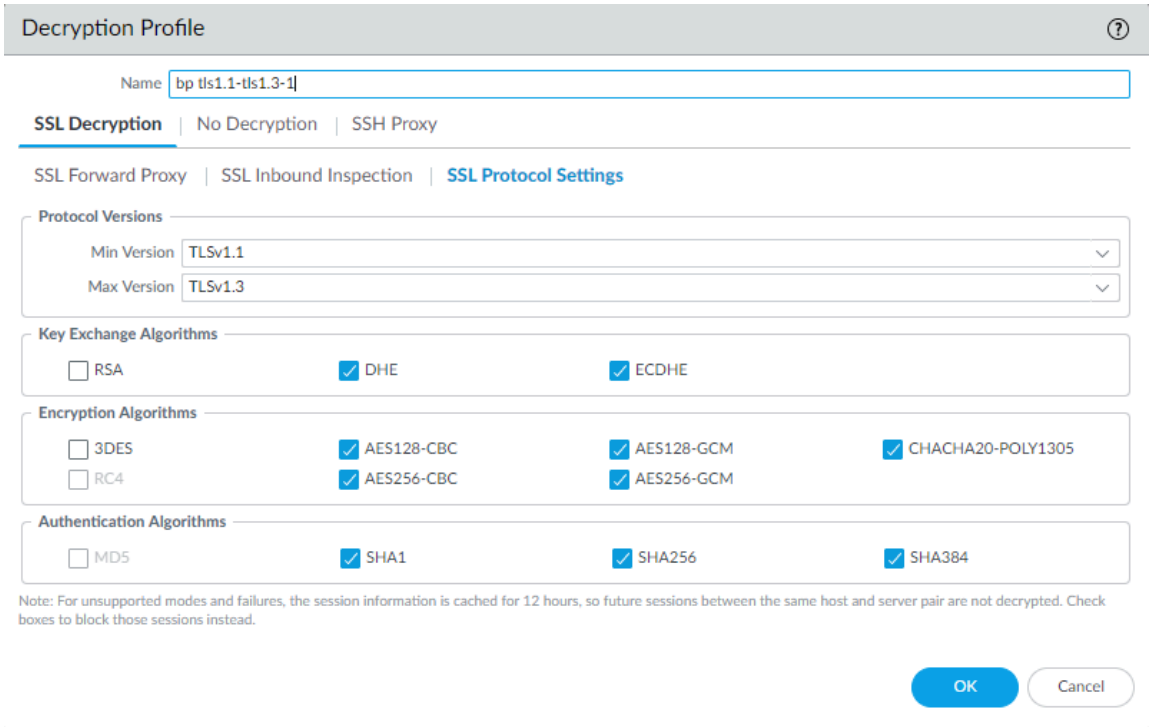


The name of the Decryption profile is **bp tls1.1-tls1.3-1**. You can also select the Big Brother policy and then select the **Options** tab to see the name of the Decryption profile.

Go to **Objects** > **Decryption** > **Decryption Profile**, select the appropriate Decryption profile, and edit it to address the version issue.

2. Go to **Objects** > **Decryption** > **Decryption Profile**.

Select the **bp tls1.1-tls1.3-1** Decryption profile and click the **SSL Protocol Settings** tab.



The minimum TLS protocol version (**Min Version**) that the profile supports is TLSv1.1. To allow the traffic that the version mismatch blocks, you could change the **Min Version** to TLSv1.0. However, a more secure option is to update the client to use a recent TLS

https://t.me/learningnets

protocol version. If you can't update the client, you can create a Decryption policy and profile that apply only to that user, device, or source address (and to any similar users, devices, or source addresses so that one policy and profile control all of this traffic) instead of applying a general Decryption policy that allows TLSv1.0 traffic.

Identify Weak Protocols and Cipher Suites

Weak TLS protocols and weak cipher suites (encryption algorithms, authentication algorithms, key exchange algorithms, and negotiated EC curves) weaken your security posture and are easier for bad actors to exploit than strong TLS protocols and strong cipher suites.

Five fields in the Decryption log entries show the protocol and cipher suites for a decryption session:

TLS VERSION	ENCRYPTION ALGORITHM	KEY EXCHANGE	AUTHENTICATI... ALGORITHM	NEGOTIATED EC CURVE
TLS1.2	AES_128_GCM	ECDHE	SHA256	secp256r1
TLS1.2	AES_256_GCM	ECDHE	SHA384	secp256r1

Track down old, vulnerable TLS versions and cipher suites so that you can make informed decisions about whether to allow connections with servers and applications that may compromise your security posture.

The examples in this topic show how to:

- Identify traffic that uses less secure TLS protocol versions.
- Identify traffic that uses a particular key exchange algorithm.
- Identify traffic that uses a particular authentication algorithm.
- Identify traffic that uses a particular encryption algorithm.

These examples show you how to use the decryption troubleshooting tools in various ways so that you can learn to use them to troubleshoot any decryption issues you may encounter.

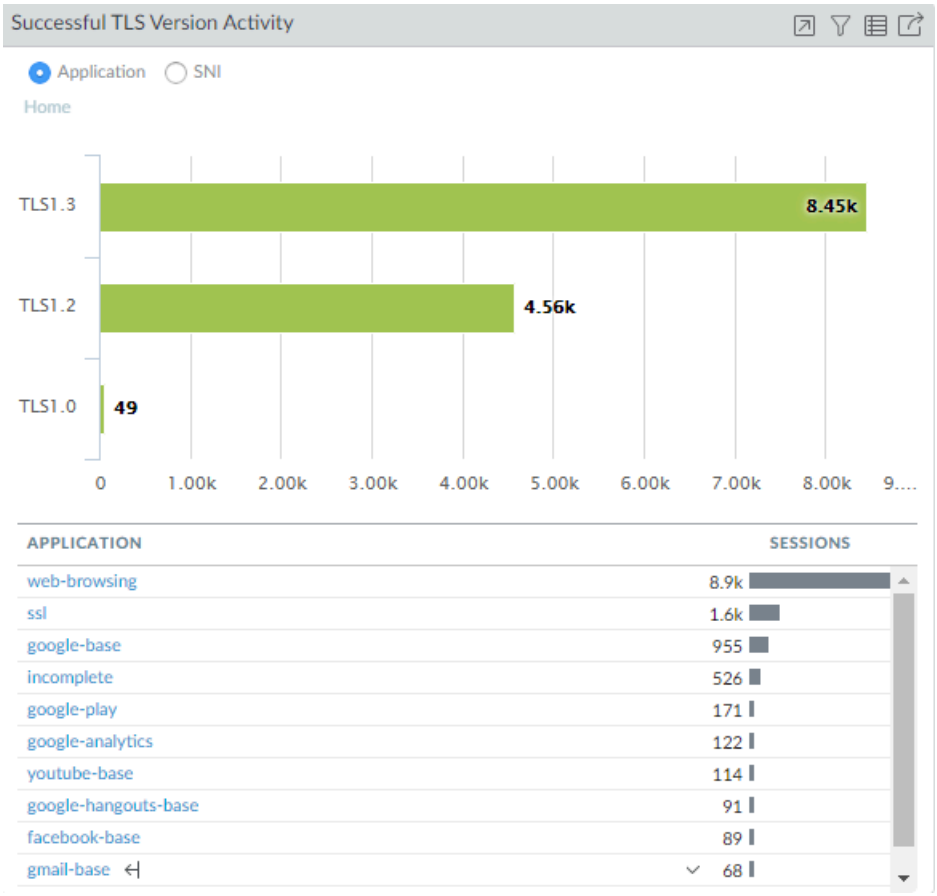


You can use Wireshark or other packet analyzers to double-check whether the client or the server caused an issue, TLS client and server versions, and other cipher suite information. This can help analyze version mismatches and other issues.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

- **TLS Protocols**—Identify traffic that uses older, less secure versions of the TLS protocol so that you can evaluate whether to allow access to servers and applications that use weak protocols.
  1. Start by checking the Application Command Center (ACC) to see if the firewall allows weak protocols (**ACC** > **SSL Activity** > **Successful TLS Version Activity**) and to get an overall view of activity.

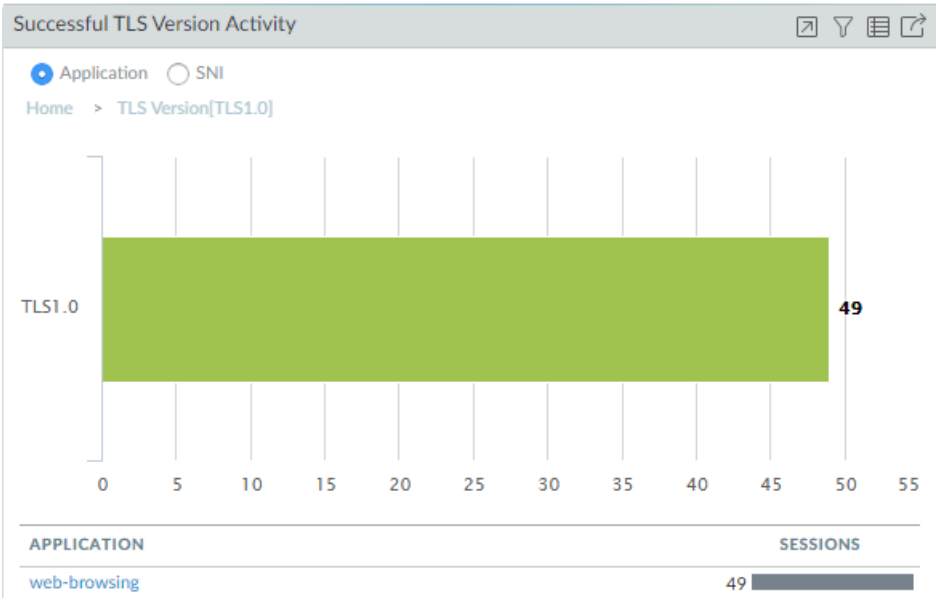


The majority of successful TLS activity in this example is TLSv1.2 and TLSv1.3 activity. However, there are a few instances of allowed TLSv1.0 traffic. Let's click the number **49**

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

to drill down into the TLSv1.0 activity and see which applications are making successful TLSv1.0 connections:

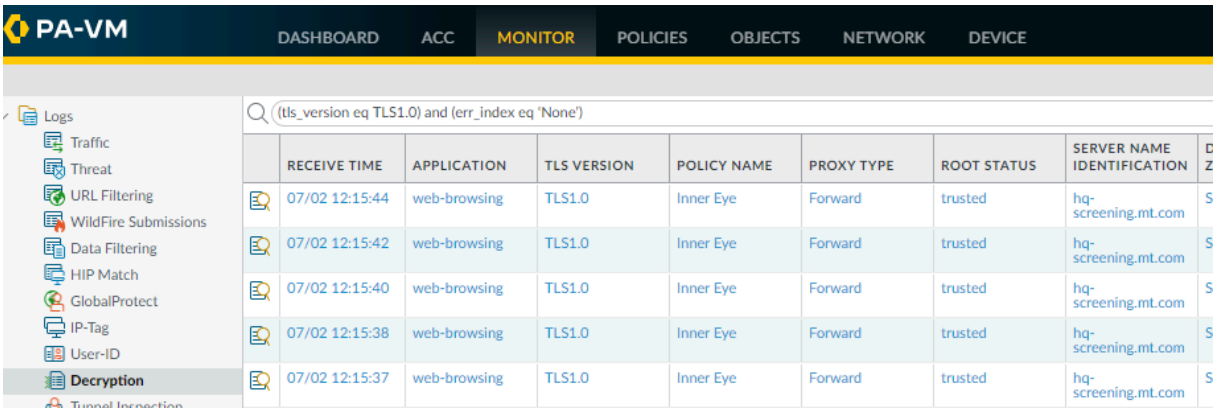


We see that the firewall is allowing traffic identified as web-browsing traffic. To gain insight into what that TLSv1.0 web-browsing traffic is and why it's allowed, we go next to the Decryption logs.

2. Filter the Decryption log to check TLSv1.0 activity details.

   Use the query **`(tls_version eq TLS1.0) and (err_index eq 'None')`** to show successful TLSv1.0 Decryption sessions.

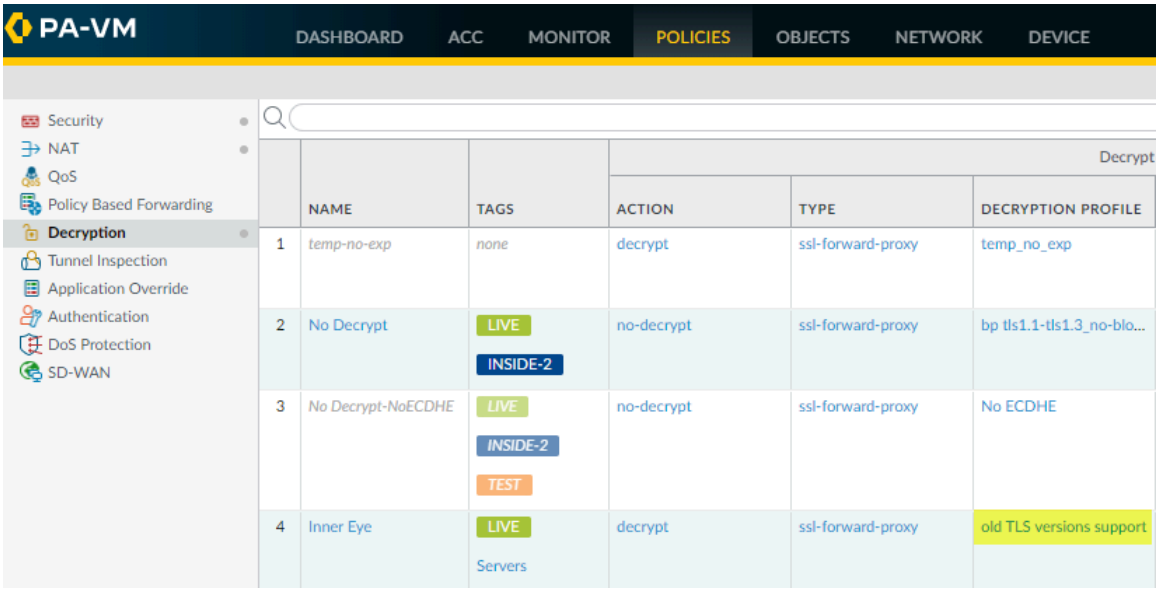
   Decryption logs show successful TLS activity only if you enable logging successful TLS handshakes in Decryption policy when you Configure Decryption Logging. If logging successful TLS handshakes is disabled, you can't check this information.



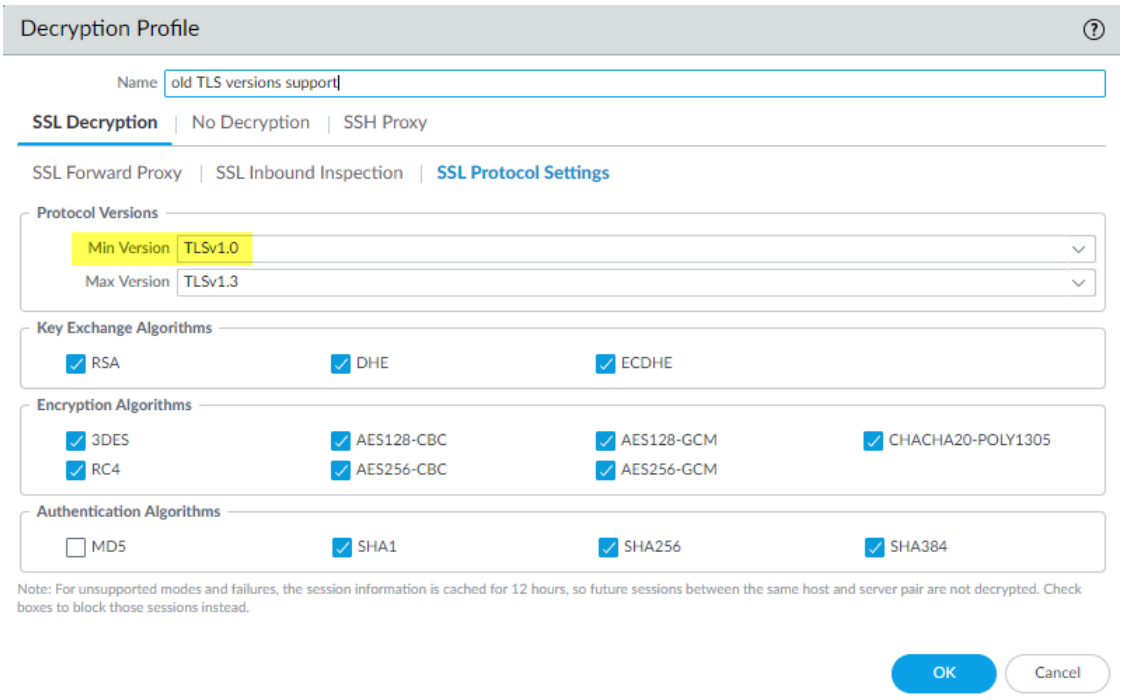
The Decryption log shows us that the name of the Decryption policy that controls the traffic is **Inner Eye** and that the name of the host is **hq-screening.mt.com**. Now we know the site that uses TLSv1.0 and we can check the Decryption policy (**Policies** >

https://t.me/learningnets

Decryption) to find the Decryption profile that controls the traffic and learn why the traffic is allowed:



We see that the Decryption profile associated with the policy is **old TLS versions support**. We check the profile (**Objects** > **Decryption** > **Decryption Profile**) and look at the SSL Protocol Settings to find out exactly what traffic the profile allows:



The profile allows TLSv1.0 traffic. The next thing to do is to decide if you want to allow access to the site (do you need access for business purposes?) or if you want to block it.

Another common scenario that results in the firewall allowing traffic that uses less secure protocols is when that traffic is not decrypted. When you filter the Decryption log for TLSv1.0 traffic, if the **Proxy Type** column contains the value **No Decrypt**, then a No

https://t.me/learningnets

Decryption policy controls the traffic, so the firewall does not decrypt or inspect it. If you don't want to allow the weak protocol, modify the Decryption profile so that it blocks TLSv1.0 traffic.

There are many ways you can filter the Decryption log to find applications and sites that use weak protocols, for example:

- Instead of filtering only for successful TLSv1.0 handshakes, filter for both successful and unsuccessful TLSv1.0 handshakes using the query **`(tls_version eq TLS1.0).`**
- Filter only for unsuccessful TLSv1.0 handshakes using the query **`(tls_version eq TLS1.0) and (err_index neq 'None')`**.
- Filter for all less secure protocols (TLSv1.1 and earlier) using the query **`(tls_version leq tls1.1)`**.

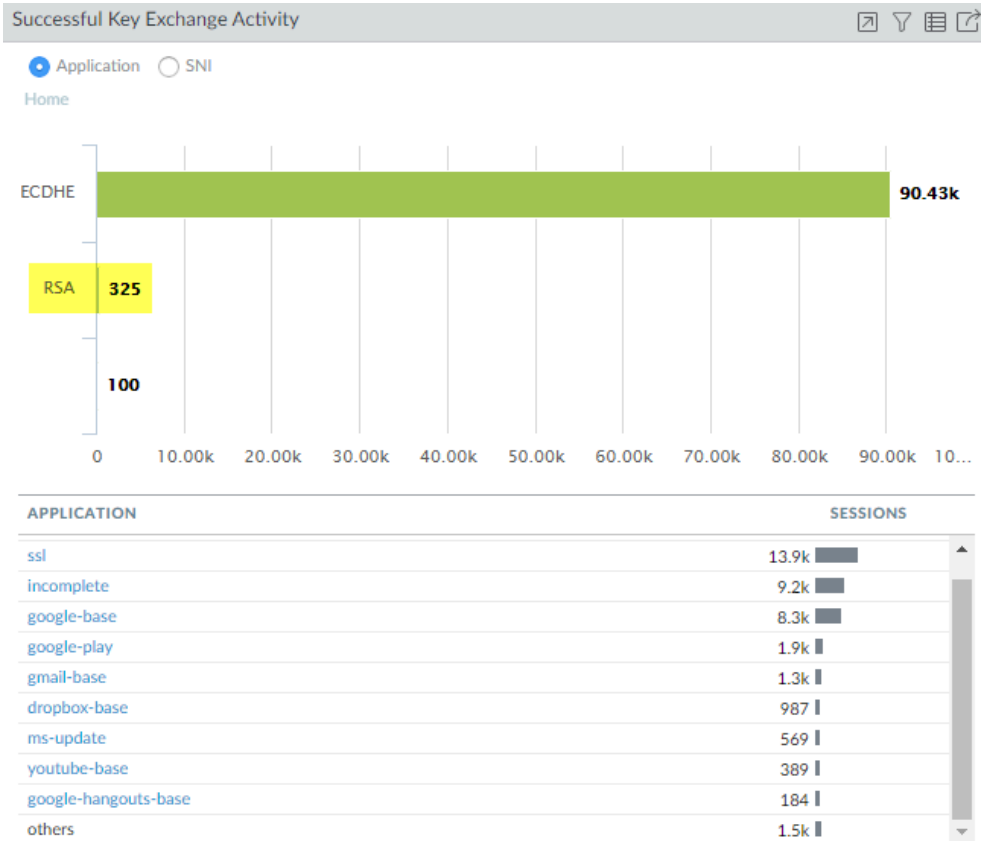
If you want to filter the logs for other TLS versions, simply replace **`TLS1.0`** or **`TLS1.1`** with another TLS version.

3. Decide what action to take for sites that use weak TLS protocols.
   - If you don't need to access the site for business purposes, the safest action is to block access to the site by editing the Decryption policy and Decryption profile that control the traffic. The Decryption log **Policy Name** column provides the policy name and the Decryption policy shows the attached Decryption profile (**Options** tab).
   - If you need to access the site for business purposes, consider creating a Decryption policy and Decryption profile that apply only to that site (or to that site and other similar sites) and block all other traffic that uses less secure protocols.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

- **Key Exchange**—Identify traffic that uses less secure key exchange algorithms.
  1. Start by checking the Application Command Center (ACC) to see which key exchange algorithms the firewall allows (**ACC** > **SSL Activity** > **Successful Key Exchange Activity**) and to get an overall view of activity.

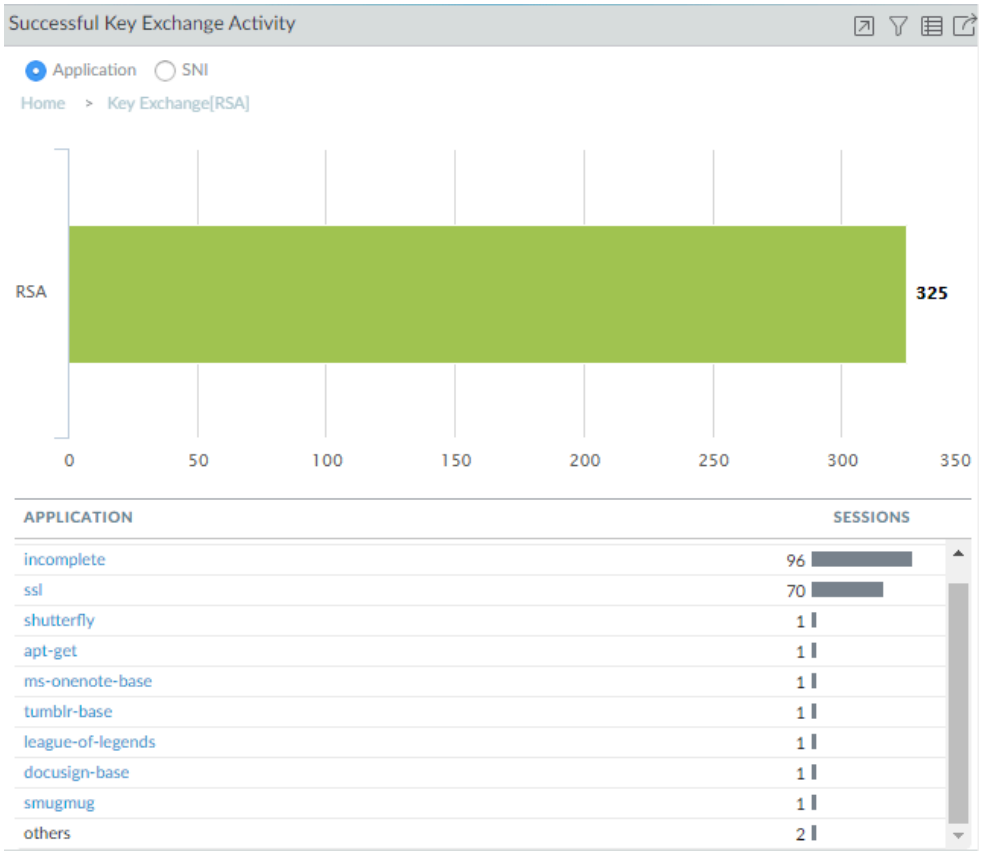


The majority of the key exchanges use the secure ECDHE key exchange algorithm. However, some key exchange sessions use the less secure RSA algorithm and a few use

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

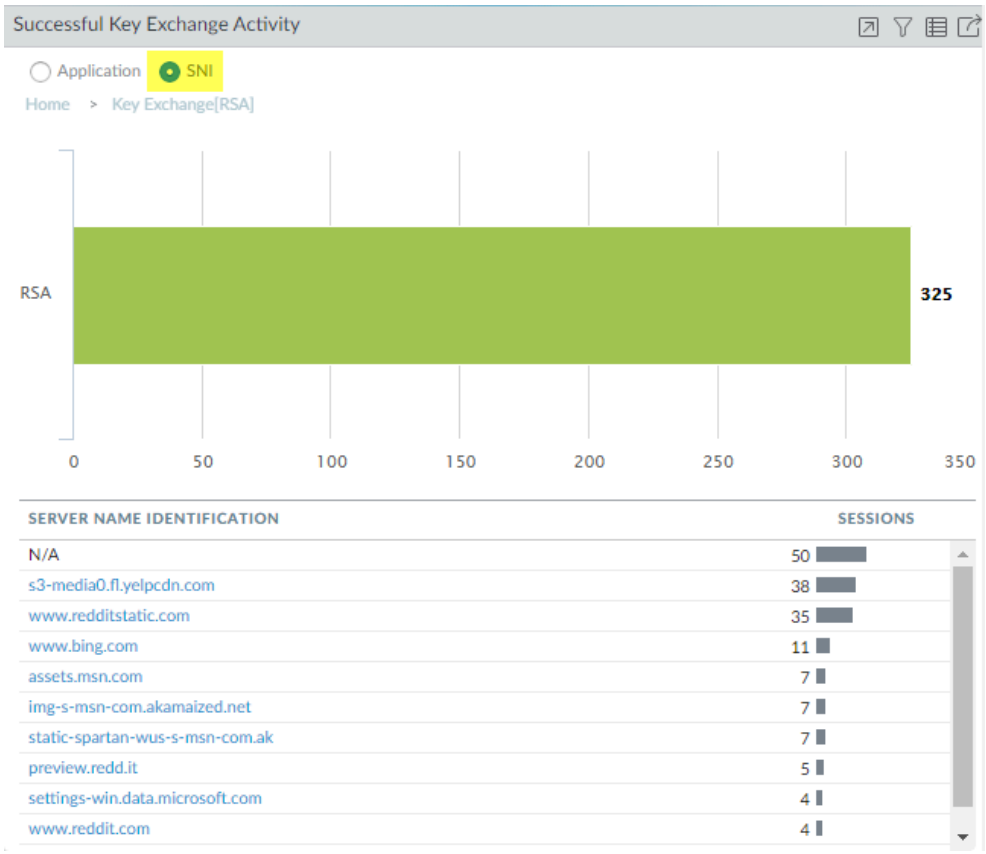
another key algorithm. To begin investigating traffic that uses RSA key exchanges, for example, click the number **325** to drill down into the data.



The drill-down shows the applications that use RSA key exchanges. We can also click the **SNI** radio button to view the RSA key exchanges by SNI:

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets



Armed with this information, we can go to the logs to gain more context about RSA key exchange usage.

2. Go to the Decryption log (**Monitor** > **Logs** > **Decryption**)) and filter them for decryption sessions that use the RSA key exchange using the query **`(tls_keyxchg eq RSA)`**:

(tls_keyxchg eq RSA)

RECEIVE TIME	SESSION ID	APPLICATION	SOURCE ADDRESS	DESTINATION ADDRESS	TLS VERSION	ERROR INDEX	ERROR	POLICY NAME
06/04 09:29:50	92884	ssl	172.30.200.30	185.31.128.129	TLS1.2	None		No Decrypt
06/04 09:29:50	92887	ssl	172.30.200.30	185.31.128.129	TLS1.2	None		No Decrypt
06/04 09:29:44	92998	ssl	172.30.200.30	74.120.19.22	TLS1.2	None		No Decrypt
06/04 09:29:24	92882	ssl	172.30.200.30	192.132.33.46	TLS1.2	Certificate	Expired server certificate	No Decrypt
06/04 09:29:24	92880	ssl	172.30.200.30	192.132.33.46	TLS1.2	Certificate	Expired server certificate	No Decrypt
06/04 09:29:23	92874	ssl	172.30.200.30	192.132.33.46	TLS1.2	Certificate	Expired server certificate	No Decrypt
06/04 09:29:23	92873	ssl	172.30.200.30	192.132.33.46	TLS1.2	Certificate	Expired server certificate	No Decrypt
06/03 22:30:11	36522	vudu	172.30.100.155	208.79.221.210	TLS1.2	None		Big Brother
06/03 20:08:57	16896	ssl	172.30.200.30	66.117.28.86	TLS1.2	None		No Decrypt
06/03 20:08:22	16947	ssl	172.30.200.30	185.31.128.129	TLS1.2	None		No Decrypt

From the **Policy Name** column in the log, we see that the **No Decrypt** Decryption policy controls most of the traffic that uses RSA key exchanges and can infer that the firewall does not decrypt the traffic and allows it without inspection. Because the traffic isn't

https://t.me/learningnets

decrypted, the firewall can't identify the application and lists it as **ssl**. If you don't want to allow traffic that uses RSA key exchanges, modify the Decryption profile attached to the Decryption policy that controls the traffic.

You can add to the query to further filter the results for a particular SNI or application that you saw in the ACC or in the first Decryption log query.

3. Decide what action to take for traffic that uses less secure key exchange algorithms.

   Block access to sites that use less secure key exchange protocols unless you need to access them for business purposes. For those sites, consider creating a Decryption policy and Decryption profile that apply only to that site (or to that site and other similar sites) and block all other traffic that uses less secure key exchange algorithms.

- Use the Decryption logs to identify sessions that uses older, less secure authentication algorithms.

  Filter the Decryption log to identify older, less secure authentication algorithms.

  For example, to identify all sessions that use the SHA1 algorithm, use the query **`(tls_auth eq SHA)`**:


(tls_auth eq SHA)

	RECEIVE TIME	SESSION ID	APPLICATION	TLS VERSION	ERROR INDEX	ERROR	POLICY NAME	SERVER NAME IDENTIFICATION	AUTHENTICATION ALGORITHM
	06/08 23:12:02	213635	ssl	TLS1.2	None		No Decrypt		SHA
	06/08 11:16:02	203438	incomplete	TLS1.2	None		Big Brother	p.sfx.ms	SHA
	06/08 11:16:02	203439	web-browsing	TLS1.2	None		Big Brother	p.sfx.ms	SHA
	06/08 11:15:01	203437	web-browsing	TLS1.2	None		Big Brother	p.sfx.ms	SHA
	06/08 02:45:32	196795	incomplete	TLS1.2	None		Big Brother	p.sfx.ms	SHA
	06/08 02:44:30	196794	web-browsing	TLS1.2	None		Big Brother	p.sfx.ms	SHA
	06/08 02:44:30	196793	web-browsing	TLS1.2	None		Big Brother	p.sfx.ms	SHA
	06/04 13:38:36	117329	web-browsing	TLS1.2	None		Big Brother	inegi.org.mx	SHA
	06/04 13:35:01	116980	web-browsing	TLS1.2	None		Big Brother	rupress.org	SHA

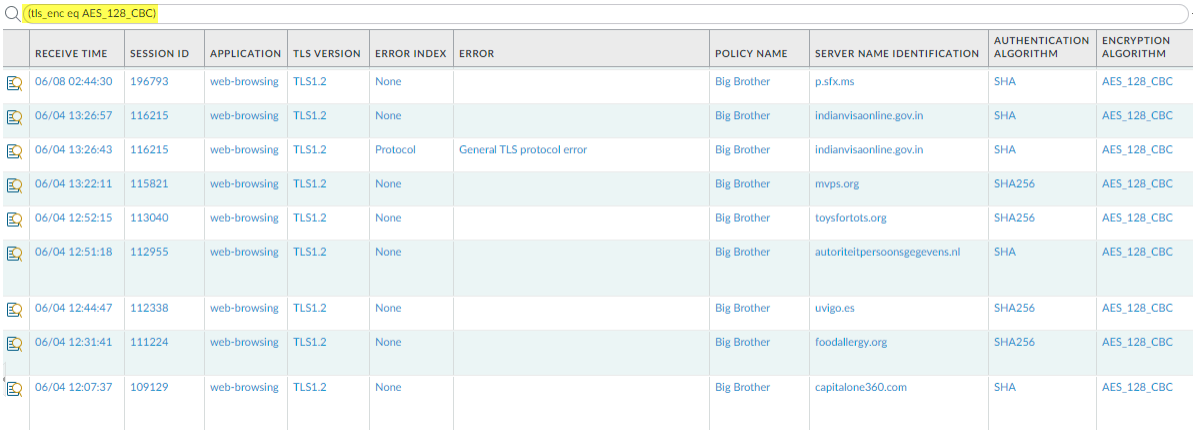
  You can add to the query to further drill down into the results. For example, you can add a particular SNI, a key exchange version (such as filtering for SHA1 sessions that also use RSA key exchanges), a TLS version, or any other metric found in a Decryption log column.

©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

- Use the Decryption logs to identify sessions that use a particular encryption algorithm.

  For example, to identify all sessions that use the AES-128-CBC encryption algorithm, use the query **`(tls_enc eq AES_128_CBC)`**:


(tls_enc eq AES_128_CBC)

	RECEIVE TIME	SESSION ID	APPLICATION	TLS VERSION	ERROR INDEX	ERROR	POLICY NAME	SERVER NAME IDENTIFICATION	AUTHENTICATION ALGORITHM	ENCRYPTION ALGORITHM
	06/08 02:44:30	196793	web-browsing	TLS1.2	None		Big Brother	p.sfx.ms	SHA	AES_128_CBC
	06/04 13:26:57	116215	web-browsing	TLS1.2	None		Big Brother	indianvisaonline.gov.in	SHA	AES_128_CBC
	06/04 13:26:43	116215	web-browsing	TLS1.2	Protocol	General TLS protocol error	Big Brother	indianvisaonline.gov.in	SHA	AES_128_CBC
	06/04 13:22:11	115821	web-browsing	TLS1.2	None		Big Brother	mvps.org	SHA256	AES_128_CBC
	06/04 12:52:15	113040	web-browsing	TLS1.2	None		Big Brother	toysfortots.org	SHA256	AES_128_CBC
	06/04 12:51:18	112955	web-browsing	TLS1.2	None		Big Brother	autoriteitpersoonsgegevens.nl	SHA	AES_128_CBC
	06/04 12:44:47	112338	web-browsing	TLS1.2	None		Big Brother	uvigo.es	SHA256	AES_128_CBC
	06/04 12:31:41	111224	web-browsing	TLS1.2	None		Big Brother	foodallergy.org	SHA256	AES_128_CBC
	06/04 12:07:37	109129	web-browsing	TLS1.2	None		Big Brother	capitalone360.com	SHA	AES_128_CBC

  You can add to the query to further drill down into the results.

  Examples of queries to find other older encryption algorithms include: **`(tls_enc eq DES_CBC)`**, **`(tls_enc eq 3DES_EDE_CBC)`**, and **`(tls_enc eq DES40_CBC)`**.

- Use this methodology and the log filter builder to create queries to investigate negotiated ECC curves and any other information you find in the Decryption log.

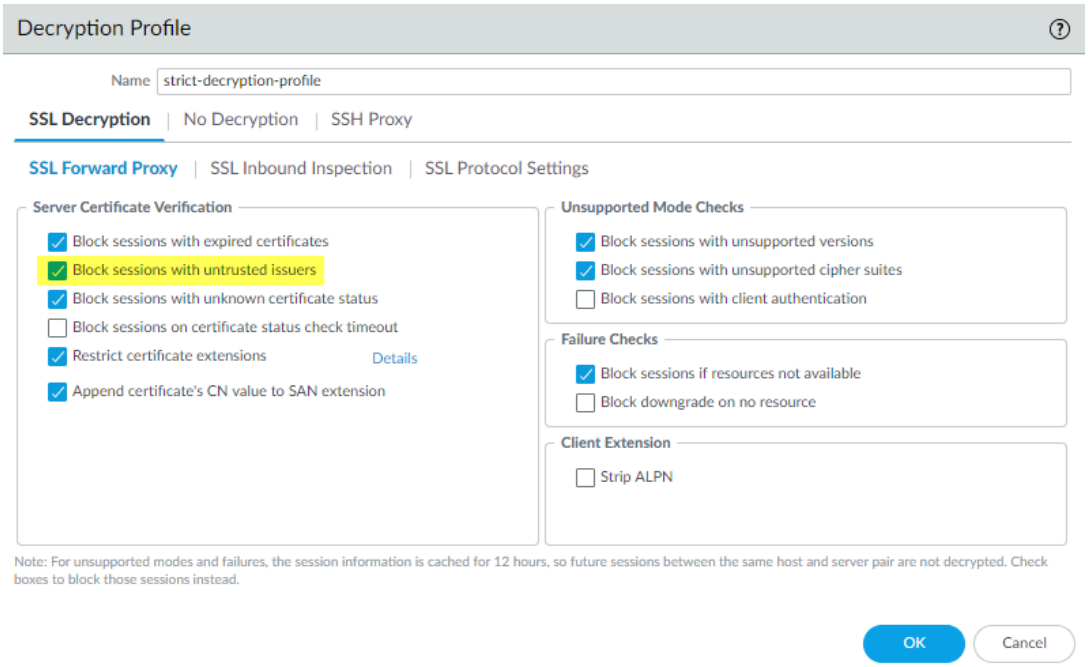
Identify Untrusted CA Certificates

Blocking access to sites with untrusted CA certificates and certificates self-signed by an untrusted root CA is a best practice because sites with untrusted CAs may indicate a man-in-the-middle attack, a replay attack, or other malicious activity.

 ©2022 Palo Alto Networks, Inc.

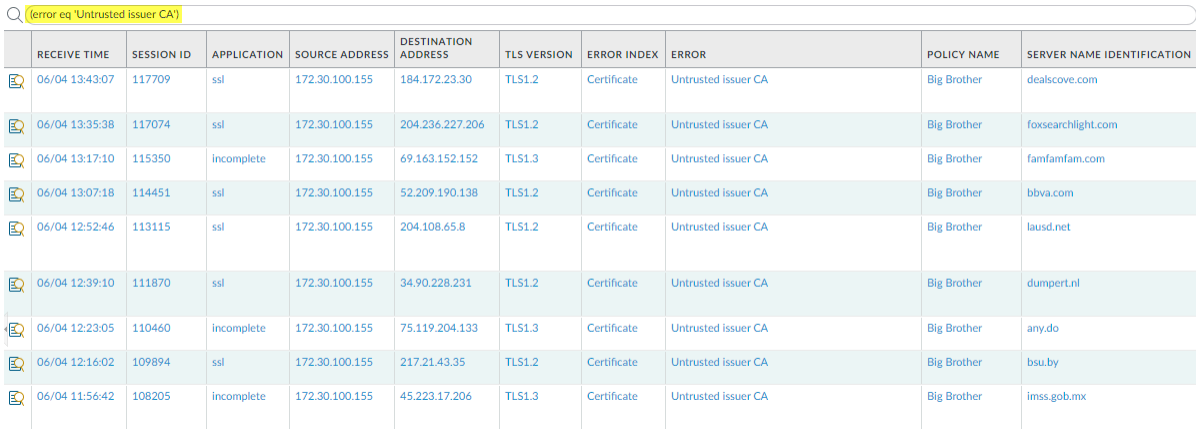
https://t.me/learningnets

STEP 1 | Ensure that you **Block sessions with untrusted issuers** in the Forward Proxy Decryption profile (**Objects** > **Decryption** > **Decryption Profiles**) to block sites with untrusted CAs.



When you block sessions with untrusted issuers in the Decryption profile, the Decryption log (**Monitor** > **Logs** > **Decryption**) logs the error.

STEP 2 | Filter the log to identify sessions that failed due to revoked certificates using the query **`(error eq 'Untrusted issuer CA')`**.



RECEIVE TIME	SESSION ID	APPLICATION	SOURCE ADDRESS	DESTINATION ADDRESS	TLS VERSION	ERROR INDEX	ERROR	POLICY NAME	SERVER NAME IDENTIFICATION
06/04 13:43:07	117709	ssl	172.30.100.155	184.172.23.30	TLS1.2	Certificate	Untrusted issuer CA	Big Brother	dealscove.com
06/04 13:35:38	117074	ssl	172.30.100.155	204.236.227.206	TLS1.2	Certificate	Untrusted issuer CA	Big Brother	foxsearchlight.com
06/04 13:17:10	115350	incomplete	172.30.100.155	69.163.152.152	TLS1.3	Certificate	Untrusted issuer CA	Big Brother	famfamfam.com
06/04 13:07:18	114451	ssl	172.30.100.155	52.209.190.138	TLS1.2	Certificate	Untrusted issuer CA	Big Brother	bbva.com
06/04 12:52:46	113115	ssl	172.30.100.155	204.108.65.8	TLS1.2	Certificate	Untrusted issuer CA	Big Brother	lausd.net
06/04 12:39:10	111870	ssl	172.30.100.155	34.90.228.231	TLS1.2	Certificate	Untrusted issuer CA	Big Brother	dumpert.nl
06/04 12:23:05	110460	incomplete	172.30.100.155	75.119.204.133	TLS1.3	Certificate	Untrusted issuer CA	Big Brother	any.do
06/04 12:16:02	109894	ssl	172.30.100.155	217.21.43.35	TLS1.2	Certificate	Untrusted issuer CA	Big Brother	bsu.by
06/04 11:56:42	108205	incomplete	172.30.100.155	45.223.17.206	TLS1.3	Certificate	Untrusted issuer CA	Big Brother	imss.gob.mx

STEP 3 | (Optional) Double-check the certificate expiration date at the Qualys SSL Labs site.

Enter the hostname of the server (**Server Name Identification** column of the Decryption log) in the **Hostname** field and **Submit** it to view certificate information for the host.

Troubleshoot Expired Certificates

If you follow Decryption best practices and **Block sessions with expired certificates** in the Forward Proxy Decryption profile or in the No Decryption profile, then if a server presents an expired certificate, the firewall blocks the session. However, if site that you need to access for

https://t.me/learningnets

business reasons allows its certificate to expire, connections to that site may be blocked and you may not know why.

You can use the Decryption log to check for expired certificates and to check for certificates that will expire soon so you can be aware of the situation and take appropriate action.

STEP 1 | Filter the Decryption log for expired certificates using the query **`(error eq 'Expired server certificate')`**.



(error eq 'Expired server certificate')

RECEIVE TIME	SESSION ID	APPLICATION	SOURCE ADDRESS	DESTINATION ADDRESS	TLS VERSION	ERROR INDEX	ERROR	SERVER NAME IDENTIFICATION	POLICY NAME
06/04 16:19:49	121352	incomplete	172.30.100.10	34.225.62.221	TLS1.3	Certificate	Expired server certificate	www.stanford.edu	Big Brother
06/04 13:43:26	117747	incomplete	172.30.100.155	104.197.149.89	TLS1.3	Certificate	Expired server certificate	phone.com	Big Brother
06/04 13:41:03	117572	incomplete	172.30.100.155	208.117.9.16	TLS1.3	Certificate	Expired server certificate	netcarshow.com	Big Brother
06/04 13:38:51	117379	ssl	172.30.100.155	69.172.200.184	TLS1.2	Certificate	Expired server certificate	royal.gov.uk	Big Brother
06/04 13:36:27	117150	ssl	172.30.100.155	107.21.104.61	TLS1.2	Certificate	Expired server certificate	www.uthscsa.edu	Big Brother
06/04 13:34:53	117004	incomplete	172.30.100.155	66.115.56.251	TLS1.3	Certificate	Expired server certificate	gunsamerica.com	Big Brother
06/04 13:33:17	116853	incomplete	172.30.100.155	34.107.140.234	TLS1.3	Certificate	Expired server certificate	skiplagged.com	Big Brother
06/04 13:32:45	116798	ssl	172.30.100.155	104.236.4.58	TLS1.2	Certificate	Expired server certificate	uploading.com	Big Brother
06/04 13:31:28	116655	incomplete	172.30.100.155	35.186.201.59	TLS1.3	Certificate	Expired server certificate	shared.com	Big Brother
06/04 13:29:32	116507	ssl	172.30.100.155	147.139.136.53	TLS1.2	Certificate	Expired server certificate	beautynesia.id	Big Brother
06/04 13:28:56	116426	incomplete	172.30.100.155	45.55.105.190	TLS1.3	Certificate	Expired server certificate	designbundles.net	Big Brother

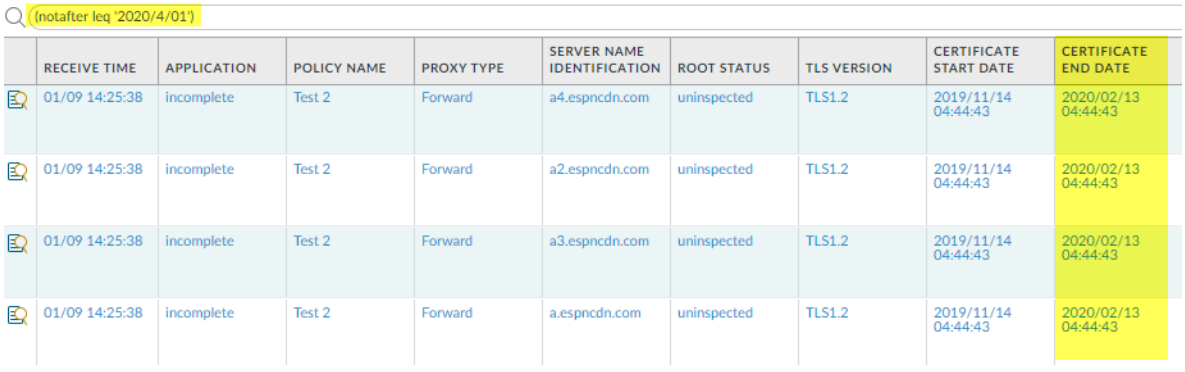
This query identifies servers that generate `Expired server certificate` errors. The firewall blocks access to these servers because of the expired certificate.

STEP 2 | (Optional) Double-check the certificate expiration date at the Qualys SSL Labs site.

Enter the hostname of the server (**Server Name Identification** column of the Decryption log) in the **Hostname** field and **Submit** it to view certificate information for the host.

STEP 3 | Filter the Decryption log (**Monitor** > **Logs** > **Decryption**) for certificates that will expire soon using a query that identifies upcoming certificate end dates.

For example, if today's date is February 1, 2020 and you want to give yourself two months to evaluate and prepare in case sites don't update their certificates, query the Decryption log for certificates that expire April 1 2020 or earlier (**`notafter leq '2020/4/01'`**):



(notafter leq '2020/4/01')

RECEIVE TIME	APPLICATION	POLICY NAME	PROXY TYPE	SERVER NAME IDENTIFICATION	ROOT STATUS	TLS VERSION	CERTIFICATE START DATE	CERTIFICATE END DATE
01/09 14:25:38	incomplete	Test 2	Forward	a4.espncdn.com	uninspected	TLS1.2	2019/11/14 04:44:43	2020/02/13 04:44:43
01/09 14:25:38	incomplete	Test 2	Forward	a2.espncdn.com	uninspected	TLS1.2	2019/11/14 04:44:43	2020/02/13 04:44:43
01/09 14:25:38	incomplete	Test 2	Forward	a3.espncdn.com	uninspected	TLS1.2	2019/11/14 04:44:43	2020/02/13 04:44:43
01/09 14:25:38	incomplete	Test 2	Forward	a.espncdn.com	uninspected	TLS1.2	2019/11/14 04:44:43	2020/02/13 04:44:43

The **Certificate End Date** column shows the eact date on which the certificate expires.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 4 | Determine the action to take for sites with expired certificates.

- If you don't need to access the site for business purposes, the safest action is to continue to block access to the site.
- If you need to access the site for business purposes, take one of the following actions:
  - Contact the administrator of the site with the expired certificate and notify them that they need to update or renew their certificate.
  - Create a Decryption policy that applies only to the sites with expired certificates that you need for business purposes and a Decryption profile that allows sites with expired certificates. Do not apply the policy to any sites that you don't need for business purposes. When a site updates its certificate, remove it from the policy.

Troubleshoot Revoked Certificates

A revoked certificate is no longer valid. It may indicate that there are security issues with a site and that the certificate is not trustworthy, although there are also benign reasons why a certificate may be revoked.



Don't trust revoked certificates; enable certificate revocation checking to deny access to sites with revoked certificates.

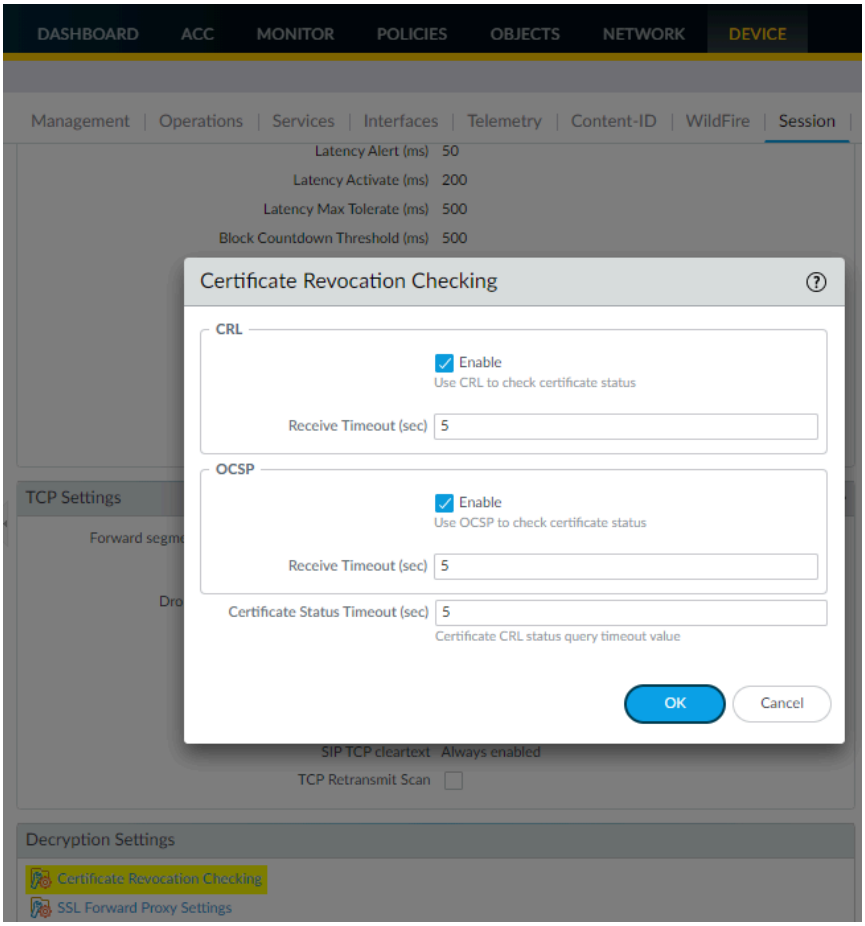
In order to drop sessions with revoked certificates and troubleshoot revoked certificates, you need to enable certificate revocation checking. If you don't enable certificate revocation checking, the firewall doesn't check for revoked certificates and you won't know if a site has a revoked certificate.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

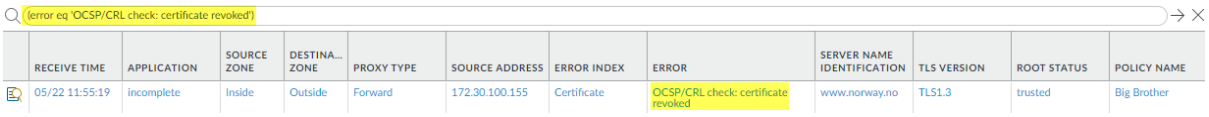
STEP 1 | Enable certificate revocation checking if you haven't already enabled it.

1. Go to **Device** > **Setup** > **Session** > **Decryption Settings**.
2. Enable both OCSP and CRL certificate checking.



If you **Block sessions on certificate status check timeout** in the Forward Proxy Decryption profile and are concerned that 5 seconds is not enough time and may result in too many sessions blocked by timeouts, set the **Receive Timeout (sec)** to a longer amount of time.

STEP 2 | Filter the Decryption log (**Monitor** > **Logs** > **Decryption**) to find certificate revocation errors using the query **`(error eq 'OCSP/CRL check: certificate revoked')`**.



STEP 3 | (Optional) Double-check the certificate expiration date at the Qualys SSL Labs site.

Enter the hostname of the server (**Server Name Identification** column of the Decryption log) in the **Hostname** field and **Submit** it to view certificate information for the host.

Troubleshoot Pinned Certificates

Certificate pinning forces the client application to validate the server's certificate against a known copy to ensure that certificate really comes from the server. The intent of pinned certificates is

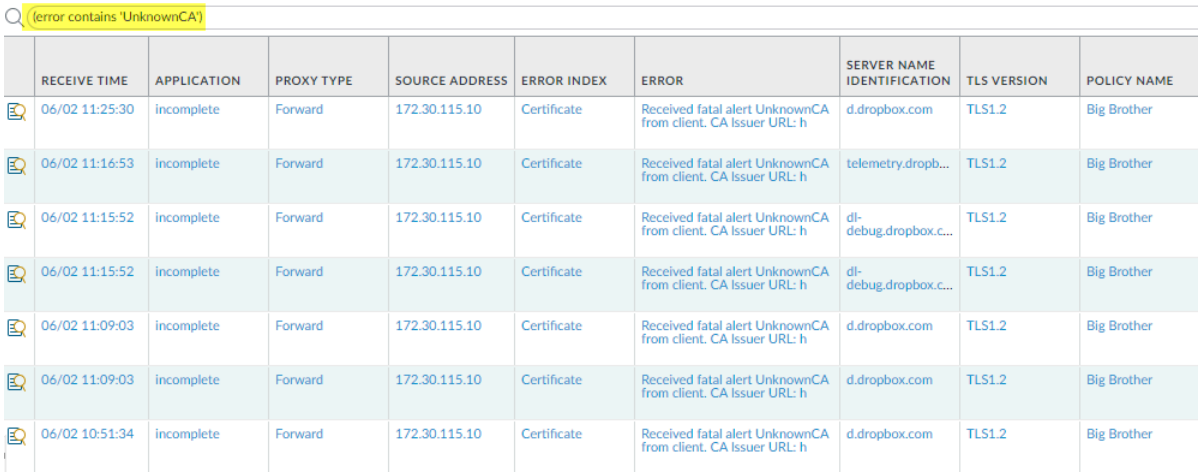
https://t.me/learningnets

to protect against man-in-the-middle (MITM) attacks where a device between the client and the server replaces the server certificate with another certificate.

Although this prevents malicious actors from intercepting and manipulating connections, it also prevents forward proxy decryption because the firewall creates an impersonation certificate instead of the server certificate to present to the client. Instead of one session that connects the client and server directly, forward proxy creates two sessions, one between the client and the firewall and another between the firewall and the server. This establishes trust with the client so that the firewall can decrypt and inspect the traffic.

However, when a certificate is pinned, the firewall cannot decrypt the traffic because the client does not accept the firewall's impersonation certificate—the client only accepts the certificate that is pinned to the application.

STEP 1 | Filter the Decryption log (**Monitor** > **Logs** > **Decryption**) to find pinned certificates using the query **`(error contains 'UnknownCA')`**.



The application generates a TLS error code (Alert) when it fails to verify the server's certificate. Different applications may use different error codes to indicate a pinned certificate. The most common error indicators for pinned certificates are UnknownCA and BadCertificate. After running the **`(error contains 'UnknownCA')`** query, run the query **`(error contains 'BadCertificate')`** to catch more pinned certificate errors.

You can use Wireshark or other packet analyzers to double-check the error. Look for the client breaking the connection immediately after the TLS handshake to confirm that it is a pinned certificate issue.

STEP 2 | Decide what to do about pinned certificates.

If you don't need access for business purposes, you can let the firewall continue to block access. If you need access, then you can Exclude a Server from Decryption for Technical Reasons by adding it to the SSL Decryption Exclusion List (**Device** > **Certificate Management** > **SSL Decryption Exclusion**.

The firewall bypasses decryption for sites on the SSL Decryption Exclusion List. The firewall cannot inspect the traffic, but the traffic is allowed.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Activate Free Licenses for Decryption Features

Decrypting SSH traffic and SSL traffic (SSL internet traffic or SSL traffic to an internal server) does not require a license. However, you must activate a free license in order to enable Decryption Mirroring. The free license requirement ensures that this feature can only be used after the approved personnel purposefully activates the associated license.

In PAN-OS 10.1, the Decryption Broker feature and free license were replaced with Network Packet Broker (see the Networking Administrator's Guide), which expands the broker's capabilities to non-decrypted TLS traffic and non-TLS traffic in addition to decrypted TLS traffic. Network Packet Broker licenses are also free to download and install from the Customer Support Portal.

Follow these steps on the Palo Alto Networks Customer Support Portal to activate a decryption mirroring feature license.

STEP 1 | Log in to the Customer Support Portal.

STEP 2 | Select **Assets** > **Devices** on the left-hand navigation pane.

STEP 3 | Find the device on which you want to enable decryption port mirroring and select **Actions** (the pencil icon).

STEP 4 | Under Activate Licenses, select **Activate Feature License**.

STEP 5 | Select the feature for which you want to activate a free license: **Decryption Port Mirror**.

STEP 6 | **Agree and Submit**.

STEP 7 | Install the decryption mirroring license on the firewall.

1. Select **Device** > **Licenses**.
2. Click **Retrieve license keys from the license server**.
3. Verify that the **Decryption Port Mirror** license is now active on the firewall.
4. Restart the firewall (**Device** > **Setup** > **Operations**). Decryption port mirroring is not available for configuration until the firewall reloads.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

URL Filtering

Palo Alto Networks URL filtering solution allows you to monitor and control the sites users can access, to prevent phishing attacks by controlling the sites to which users can submit valid corporate credentials, and to enforce safe search for search engines like Google and Bing.

- About Palo Alto Networks URL Filtering Solution
- How Advanced URL Filtering Works
- URL Filtering Inline ML
- URL Filtering Use Cases
- URL Categories
- Plan Your URL Filtering Deployment
- URL Filtering Best Practices
- Activate The Advanced URL Filtering Subscription
- Test URL Filtering Configuration
- Configure URL Filtering
- Configure URL Filtering Inline ML
- Monitor Web Activity
- Log Only the Page a User Visits
- Create a Custom URL Category
- URL Category Exceptions
- Use an External Dynamic List in a URL Filtering Profile
- Allow Password Access to Certain Sites
- Prevent Credential Phishing
- Safe Search Enforcement
- URL Filtering Response Pages
- Customize the URL Filtering Response Pages
- HTTP Header Logging
- Request to Change the Category for a URL
- Troubleshoot URL Filtering
- PAN-DB Private Cloud
- Enable SSL/TLS Handshake Inspection

https://t.me/learningnets

About Palo Alto Networks URL Filtering Solution

Palo Alto Networks URL filtering solution, Advanced URL Filtering, is a subscription-based feature that protects against web-based threats by giving you a way to safely enable web access while controlling how your users interact with online content. The Advanced URL Filtering subscription provides all of the functionality offered by the legacy URL Filtering subscription, while also providing the added benefit of real-time URL analysis.



Legacy URL Filtering subscription holders are able to continue using their URL filtering deployment until the end of the license term.

You can create policy rules to limit access to sites based on URL categories, users, and groups. (See URL Filtering Use Cases for different ways you can leverage your Advanced URL Filtering subscription to meet your organization's web security needs.)

With Advanced URL Filtering enabled, URL requests are:

- Compared against the PAN-DB URL database, which contains millions of websites that have been categorized. You can use these URL categories in URL Filtering profiles or as match criteria to enforce Security policy. You can also use URL filtering to enforce safe search settings for your users and to prevent credential theft based on URL category.
- Analyzed in real-time using the cloud-based Advanced URL Filtering detection modules to provide protection against new and unknown threats that do not currently exist in the URL filtering database.
- Inspected for phishing and malicious JavaScript using inline machine learning (ML), a firewall-based analysis solution, which can block unknown malicious web pages in real-time.

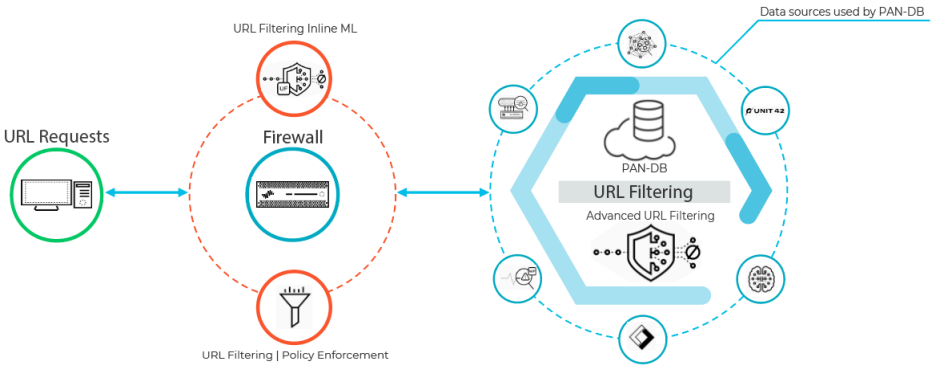
If the network security requirements in your enterprise prohibit the firewalls from directly accessing the Internet, Palo Alto Networks provides an offline URL filtering solution with the *PAN-DB Private Cloud*. This allows you to deploy a PAN-DB private cloud on one or more M-600 appliances that function as PAN-DB servers within your network.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

How Advanced URL Filtering Works

Advanced URL Filtering classifies websites based on site content, features, and safety. A URL can have up to four URL categories, including risk categories (high, medium, and low) that indicate the likelihood that the site will expose you to threats. As PAN-DB, the Advanced URL Filtering URL database, categorizes sites, firewalls with Advanced URL Filtering enabled can leverage that knowledge to enforce your organization's security policies. In addition to the protection offered by the PAN-DB database, Advanced URL Filtering provides real-time analysis using machine learning to defend against new and unknown threats. This provides protection against malicious URLs that are updated or introduced before URL filtering databases have an opportunity to analyze and add the content, giving attackers an open period from which they can launch precision attack campaigns. Advanced URL Filtering compensates for the coverage gaps inherent in database solutions by providing real time URL analysis on a per request basis. The ML-based models used by advanced URL filtering have been trained, and are continuously updated, to detect various malicious URLs, phishing web pages, and C2.



When a user requests a web page, the firewall queries user-added exceptions and PAN-DB for the site's risk category. PAN-DB uses URL information from Unit 42, WildFire, passive DNS, Palo Alto Networks telemetry data, data from the Cyber Threat Alliance, and applies various analyzers to determine the category. If the URL displays risky or malicious characteristics, it is also submitted to Advanced URL Filtering in the cloud for real-time analysis and generates additional analysis data. The resulting risk category is then retrieved by the firewall and is used to enforce the web-access rules based on your policy configuration. Additionally, the firewall caches site categorization information for new entries to enable fast retrieval for subsequent requests, while it removes URLs that users have not accessed recently so that it accurately reflects the traffic in your network. Additionally, checks built into PAN-DB cloud queries ensure that the firewall receives the latest URL categorization information. If you do not have Internet connectivity or an active URL filtering license, no queries are made to PAN-DB.



 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

The firewall determines a website's URL category by comparing it to entries in 1) custom URL categories, 2) external dynamic lists (EDLs), and 3) predefined URL categories, in order of precedence.

Firewalls configured to analyze URLs in real-time using machine learning on the dataplane provides an additional layer of security against phishing websites and JavaScript exploits. The inline ML models used to identify these URL-based threats extend to currently unknown as well as future variants of threats that match characteristics that Palo Alto Networks has identified as malicious. To keep up with the latest changes in the threat landscape, inline ML models are added or updated via content releases.

When the firewall checks PAN-DB for a URL, it also looks for critical updates, such as URLs that previously qualified as benign but are now malicious.

If you believe PAN-DB has incorrectly categorized a site, you can submit a URL category change request in your browser through Test A Site or directly from the firewall logs.



Did you know?

Technically, the firewall caches URLs on both the management plane and the dataplane:

- *PAN-OS 9.0 and later releases do not download PAN-DB seed databases. Instead, upon activation of the URL filtering license, the firewall populates the cache as URL queries are made.*
- *The management plane holds more URLs and communicates directly with PAN-DB. When the firewall cannot find a URL's category in the cache and performs a lookup in PAN-DB, it caches the retrieved category information in the management plane. The management plane passes that information along to the dataplane, which also caches it and uses it to enforce policy.*
- *The dataplane holds fewer URLs and receives information from the management plane. After the firewall checks URL category exception lists (custom URL categories and external dynamic lists) for a URL, the next place it looks is the dataplane. Only if the firewall cannot find the URL categorized in the dataplane does it check the management plane and, if the category information is not there, PAN-DB.*

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

URL Filtering Inline ML

URL Filtering inline ML enables the firewall dataplane to apply machine learning on webpages to alert users when phishing variants are detected while preventing malicious variants of JavaScript exploits from entering your network. Inline ML dynamically analyzes and detects malicious content by evaluating various web page details using a series of ML models. Each inline ML model detects malicious content by evaluating file details, including decoder fields and patterns, to formulate a high probability classification and verdict, which is then used as part of your larger web security policy. URLs classified as malicious are forwarded to PAN-DB for additional analysis and validation. You can specify URL exceptions to exclude any false-positives that might be encountered. This allows you to create more granular rules for your profiles to support your specific security needs. To keep up with the latest changes in the threat landscape, inline ML models are updated regularly and added via content releases. An active Advanced URL Filtering or legacy URL Filtering license is required to configure URL Filtering inline ML.

Inline ML-based protection can also be enabled to detect malicious PE (portable executables), ELF and MS Office files, and PowerShell and shell scripts in real-time as part of your Antivirus profile configuration. For more information, refer to: WildFire Inline ML

URL Filtering inline ML is not supported on the VM-50 or VM50L virtual appliance.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

URL Filtering Use Cases

There are many ways to enforce web page access beyond only blocking and allowing certain sites. For example, you can use multiple categories per URL to allow users to access a site, but block particular functions like submitting corporate credentials or downloading files. You can also use URL categories to enforce different types of policy, such as Authentication, Decryption, QoS, and Security.

Read on for more about the different ways that you can deploy URL filtering.

Control web access based on URL category

You can create a URL Filtering profile that specifies an action for a URL category and attach the profile to a policy rule. The firewall enforces policy against traffic based on the settings in the profile. For example, to block all gaming websites you could set the block action for the URL category *games* in the URL Filtering profile and attach it to the Security policy rule(s) that allow web access.

Multi-Category URL Filtering

Every URL can have up to four categories, including a risk category that indicates the likelihood a site will expose you to threats. More granular URL categorizations means that you can move beyond a basic "block-or-allow" approach to web access. Instead, you can control how your users interact with online content that, while necessary for business, is more likely to be used as part of a cyberattack.

For instance, you might consider certain URL categories risky to your organization, but are hesitant to block them outright as they also provide valuable resources or services (like cloud storage services or blogs). Now, you can allow users to visit sites that fall into these types of URL categories while you protect your network by decrypting and inspecting traffic and enforcing read-only access to the content.

For a URL category that you want to tightly control, set the URL Filtering profile action to alert as part of the steps to configure URL Filtering. Then continue to follow the URL Filtering best practices: decrypt the URL category, block dangerous file downloads, and turn on credential phishing prevention.

You can also define a custom URL category by selecting **Category Match** and specifying two or more PAN-DB categories of which the new category will consist. Creating a custom category from multiple categories allows you to target enforcement for a website or page that matches all of the categories specified in the custom URL category object.

Block or allow corporate credential submissions based on URL category

Prevent credential phishing by enabling the firewall to detect corporate credential submissions to sites, and then control those submissions based on URL category. Block users from submitting credentials to malicious and untrusted sites, warn users against entering corporate credentials on unknown sites or reusing corporate credentials on non-corporate sites, and explicitly allow users to submit credentials to corporate and sanctioned sites.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Enforce Safe Search Settings

Many search engines have a safe search setting that filters out adult images and videos from search results. You can enable the firewall to block search results if the end user is not using the strictest safe search settings, and you can transparently enable safe search for your users. The firewall supports safe search enforcement for the following search providers: Google, Yahoo, Bing, Yandex, and YouTube. See how to get started with Safe Search Enforcement.

Enforce Password Access to Certain Sites

You can block access to a site for most users while allowing certain users to access the site. See how to allow password access to certain sites.

Block high-risk file downloads from certain URL categories

You can block high-risk file downloads from specific URL categories by creating a Security policy with a File Blocking profile attached.

Enforce Security, Decryption, Authentication, and QoS policies based on URL category

You can enforce different types of firewall policies based on URL categories. For example, suppose you have enabled Decryption, but you want to exclude certain personal information from being decrypted. In this case you could create a Decryption policy rule that excludes websites that match the URL categories *financial-services* and *health-and-medicine* from decryption. Another example would be to use the URL category *streaming-media* in a QoS policy to apply bandwidth controls to websites that fall in to this category.

The following table describes the policies that accept URL categories as match criteria:

Policy Type	Description
Decryption	You can also use URL categories to phase-in decryption, and to exclude URL categories that might contain sensitive or personal information from decryption (like financial-services and health-and-medicine). Plan to decrypt the riskiest traffic first (URL categories most likely to harbor malicious traffic, such as gaming or high-risk) and then decrypt more as you gain experience. Alternatively, decrypt the URL categories that don't affect your business first (if something goes wrong, it won't affect business), for example, news feeds. In both cases, decrypt a few URL categories, listen to user feedback, run reports to ensure that decryption is working as expected, and then gradually decrypt a few more URL categories, and so on. Plan to make decryption exclusions to exclude sites from decryption if you can't decrypt them for technical reasons or because you choose not to decrypt them.  *Decrypting traffic based on URL categories is a best practice for both URL Filtering and Decryption.*

https://t.me/learningnets

Policy Type	Description
Authentication	To ensure that users authenticate before being allowed access to a specific category, you can attach a URL category as a match criterion for Authentication policy rules.
QoS	Use URL categories to allocate throughput levels for specific website categories. For example, you may want to allow the *streaming-media* category, but limit throughput by adding the URL category to a QoS policy rule.
Security	In Security policy rules, you can use URL categories in two ways: • Enforce policy based on URL categories by selecting them as match criteria. • Attach a URL Filtering profile that specifies the policy action for each category. If for example, the IT-security group in your company needs access to the *hacking* category, but all other users are denied access to the category, you must create the following rules: • A Security policy rule that allows the IT-Security group to access content categorized as *hacking*. The Security policy rule references the *hacking* category in the **Services/URL Category** tab and IT-Security group in the **Users** tab. • Another Security policy rule that allows general web access for all users. To this rule you attach a URL Filtering profile that blocks the *hacking* category. You must list the policy that allows access to *hacking* before the policy that blocks *hacking*. This is because the firewall evaluates Security policy rules from the top down, so when a user who is part of the security group attempts to access a *hacking* site, the firewall evaluates the policy rule that allows access first and grants the user access. The firewall evaluates users from all other groups against the general web access rule that blocks access to the *hacking* sites.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

URL Categories

PAN-DB classifies websites based on site content, features, and safety. A URL can have up to four categories, including risk categories (high, medium, and low), which indicate how likely it is that the site will expose you to threats. For a complete list of predefined URL categories, see PAN-DB URL Filtering Categories.

Visit Test A Site to see how PAN-DB categorizes a URL, and to learn about all available URL categories. You can also use Test A Site to submit a URL category change request, or you can submit the request directly in the firewall: select **Monitor** > **Logs** and open the details for a log entry. Under the URL category, you'll see the option to submit a change request.

Read on to learn more about URL categories:

- URL Filtering Use Cases
- Security-Focused URL Categories
- Malicious URL Categories
- Verified URL Categories
- Policy Actions You Can Take Based on a URL Category

Security-Focused URL Categories

Security-focused URL categories can help you to reduce your attack surface by providing targeted decryption and enforcement for sites that pose varying levels of risk, but are not confirmed malicious. Websites are classified with a security-related category only so long as they meet the criteria for that category; as site content changes, policy enforcement dynamically adapts. You cannot submit a change request for security-focused URL categories.

Security-Focused URL Categories	
High-Risk	High-risk sites include: • Sites previously confirmed to be malware, phishing, or C2 sites. These sites will remain in this category for at least 30 days. • Unknown domains are classified as high-risk until PAN-DB completes site analysis and categorization. • Sites that are associated with confirmed malicious activity. For example, a page might be high-risk if there are malicious hosts on the same domain, even if the page itself does not contain malicious content. • Bulletproof ISP-hosted sites. • Domains classified as DDNS due to the presence of an active dynamic DNS configuration. • Sites hosted on IPs from ASNs that are known to allow malicious content.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Security-Focused URL Categories	
	Default and Recommended Policy Action: Alert
Medium-Risk	Medium-risk sites include: • All cloud storage sites (with the URL category **online-storage-and-backup**). • Sites previously confirmed to be malware, phishing, or C2 sites that have displayed only benign activity for at least 30 days. These sites will remain in this category for an additional 60 days. • Unknown IP addresses are categorized as medium-risk until PAN-DB completes site analysis and categorization. **Default and Recommended Policy Action: Alert**
Low-Risk	Sites that are not medium or high risk are considered low risk. These sites have displayed benign activity for a minimum of 90 days. **Default and Recommended Policy Action: Allow**
Newly-Registered Domains	Identifies sites that have been registered within the last 32 days. New domains are frequently used as tools in malicious campaigns. **Default Policy Action: Alert** **Recommended Policy Action: Block** *Newly-registered domains are often generated purposefully or by domain generation algorithms and used for malicious activity. It is a best practice to block this URL category.*

Malicious URL Categories

We strongly recommend that you block the URL categories that identify malicious or exploitive content. To get started, you can clone the default URL Filtering profile which blocks malware, phishing, and command-and-control URL categories by default. The default URL Filtering profile also blocks the abused-drugs, adult, gambling, hacking, questionable, and weapons URL categories. Whether to block these URL categories depends on your business requirements. For example, a university probably won't want to restrict student access to most of these sites because availability is important, but a business that values security first may block some or all of them.

- **command-and-control**—Command-and-control URLs and domains used by malware and/or compromised systems to surreptitiously communicate with an attacker's remote server to receive malicious commands or exfiltrate data.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

- **malware**—Sites known to host malware or used for command and control (C2) traffic. May also exhibit Exploit Kits.
- **phishing**—Known to host credential phishing pages or phishing for personal identification. This includes web content that covertly attempts to fool the user in order to harvest information, including login credentials, credit card information – voluntarily or involuntarily, account numbers, PINs, and any information considered to be personally identifiable information (PII) from victims via social engineering techniques. Technical support scams and scareware are also included as phishing.
- **grayware**—Websites and services that do not meet the definition of a virus or pose a direct security threat but displays obtrusive behavior and influences users to grant remote access or perform other unauthorized actions. Grayware includes scams, illegal activities, criminal activities, get rich quick sites, adware, and other unwanted or unsolicited applications, such as embedded crypto miners or hijackers that change the elements of the browser. Typosquatting domains that do not exhibit maliciousness and is not owned by the targeted domain will be categorized as grayware. Prior to Content release version 8206, the firewall placed grayware in either the malware or questionable URL category. If you are unsure about whether to block grayware, start by alerting on grayware, investigate the alerts, and then decide whether to block grayware or continue to alert on grayware.
- **dynamic-dns**—Hosts and domain names for systems with dynamically assigned IP addresses and which are oftentimes used to deliver malware payloads or C2 traffic. Also, dynamic DNS domains do not go through the same vetting process as domains that are registered by a reputable domain registration company, and are therefore less trustworthy.
- **unknown**—Sites that have not yet been identified by PAN-DB. If availability is critical to your business and you must allow the traffic, alert on unknown sites, apply the best practice Security profiles to the traffic, and investigate the alerts.

  PAN-DB Real-Time Updates learns unknown sites after the first attempt to access an unknown site, so unknown URLs are identified quickly and become known URLs that the firewall can then handle based on the actual URL category.
- **newly-registered-domain**—Newly registered domains are often generated purposely or by domain generation algorithms and used for malicious activity.
- **copyright-infringement**—Domains with illegal content, such as content that allows illegal download of software or other intellectual property, which poses a potential liability risk. This category was introduced to enable adherence to child protection laws required in the education industry as well as laws in countries that require internet providers to prevent users from sharing copyrighted material through their service.
- **extremism**—Websites promoting terrorism, racism, fascism, or other extremist views discriminating against people or groups of different ethnic backgrounds, religions or other beliefs. This category was introduced to enable adherence to child protection laws required in the education industry. In some regions, laws and regulations may prohibit allowing access to extremist sites, and allowing access may pose a liability risk.
- **proxy-avoidance-and-anonymizers**—URLs and services often used to bypass content filtering products.
- **questionable**— Websites containing tasteless humor, offensive content targeting specific demographics of individuals, or groups of people.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

- **parked**—Domains registered by individuals, oftentimes later found to be used for credential phishing. These domains may be similar to legitimate domains, for example, pal0alto0netw0rks.com, with the intent of phishing for credentials or personal identify information. Or, they may be domains that an individual purchases rights to in hopes that it may be valuable someday, such as panw.net.

For categories that you decide to alert on, instead of block, you can very strictly control how users interact with site content. For example, give users access to the resources they need (like developer blogs for research purposes or cloud storage services), but take the following precautions to reduce exposure to web-based threats:

- Follow the Anti-Spyware, Vulnerability Protection, and File Blocking best practices. A protective measure would be to block downloads of dangerous file types and blocking obfuscated JavaScript for sites that you are alerting on.
- Target decryption based on URL category. A good start would be to decrypt high-risk and medium-risk sites.
- Display a response page to users when they visit high-risk and medium-risk sites. Alert them that the site they are attempting to access is potentially malicious, and advise them on how to take precautions if they decide to continue to the site.
- Stop credential theft by blocking users from submitting their corporate credentials to sites including those that are high-risk and medium-risk.

Verified URL Categories

Palo Alto Networks verifies the URL categories to which a requested site belongs. If PAN-DB determines that a URL belongs to malicious URL categories, it does not assign the site a security-focused URL category (high-risk, medium-risk, low-risk). The firewall automatically blocks sites in these categories because they pose an unacceptable level of risk for most environments. Risk levels are reserved only for URLs that have *not* been classified as malicious or are *no longer* classified as malicious due to a lack of evidence.

The following table lists the categories that PAN-DB considers malicious and blocks by default, with the exception of `Private IP Addresses`. Private IP addresses (and hosts) are unique to the host environment and are invisible to PAN-DB. As a result, Palo Alto Networks does not assign a risk rating to sites in this category.

Category	Default Action
Malware	Block
Phishing	
Command and Control	
Grayware	
Private IP Addresses	Allowed (no default action)

For more information about current URL categories, see Complete List of PAN-DB URL Filtering Categories.

©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Policy Actions You Can Take Based on URL Categories

On the firewall, you can use a URL Filtering profile to specify how you would like to enforce URL categories. By default, site access for all URL categories is set to allow when you create a new URL Filtering profile. This means that the users will be able to browse to all sites freely and the traffic is not logged. Customize the URL Filtering profile by deciding what type of **Site Access** you want to enforce for each category. To prevent credential phishing, you can also allow or disallow **User Credential Submissions** based on URL category (for example, you can block user credential submissions to medium and high-risk sites). Users can still access these sites, but cannot enter submit their corporate credentials to them.

To start enforcing the actions you've defined in a URL Filtering, you'll need to attach the profile to a Security policy rule. The firewall enforces the profile actions on traffic that matches the Security policy rule (for details, see Configure URL Filtering).



Learn more about configuring a best practice URL Filtering profile to ensure protection against URLs that have been observed hosting malware or exploitative content.

Action	Description
Site Access	
alert	The website is allowed and a log entry is generated in the URL filtering log.  *Set **alert** as the Action for categories of traffic you don't block to log and provide visibility into the traffic.*
allow	The website is allowed and no log entry is generated.  *Don't set **allow** as the Action for categories of traffic you don't block because you lose visibility into traffic you don't log. Instead, set **alert** as the Action for categories of traffic you don't block to log and provide visibility into the traffic.*
block	The website is blocked and the user will see a response page and will not be able to continue to the website. A log entry is generated in the URL filtering log. Blocking site access for a URL category also sets User Credential Submissions for that URL category to block.
continue	The user will be prompted with a response page indicating that the site has been blocked due to company policy, but the user is prompted with the option to continue to the website. The **continue** action is typically used for categories that are considered benign and is used to improve the user experience by giving them the option to continue if they feel the site is incorrectly categorized. The response page message can be

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Action	Description
	customized to contain details specific to your company. A log entry is generated in the URL filtering log. *The Continue page doesn't display properly on client systems configured to use a proxy server.*
override	The user will see a response page indicating that a password is required to allow access to websites in the given category. With this option, the security admin or helpdesk person would provide a password granting temporary access to all websites in the given category. A log entry is generated in the URL filtering log. See Allow Password Access to Certain Sites. In earlier release versions, URL Filtering category overrides had priority enforcement ahead of custom URL categories. As part of the upgrade to PAN-OS 9.0, URL category overrides are converted to custom URL categories, and no longer receive priority enforcement over other custom URL categories. Instead of the action you defined for the category override in previous release versions, the new custom URL category is enforced by the security policy rule with the strictest URL Filtering profile action. From most strict to least strict, possible URL Filtering profile actions are: block, override, continue, alert, and allow. This means that, if you had URL category overrides with the action allow, there's a possibility the overrides might be blocked after they are converted to custom URL category in PAN-OS 9.0. *The Override page doesn't display properly on client systems configured to use a proxy server.*
none	The **none** action only applies to custom URL categories. Select **none** to ensure that if multiple URL Filtering profiles exist, the custom category will not have any impact on other profiles. For example, if you have two URL Filtering profiles and the custom URL category is set to **block** in one profile, if you do not want the block action to apply to the other profile, you must set the action to **none**. Also, in order to delete a custom URL category, it must be set to **none** in any profile where it is used.
User Credential Permissions *These settings require you to first* set up credential phishing prevention.	
alert	Allow users to submit corporate credentials to sites in this URL category, but generate a URL Filtering alert log each time this occurs.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Action	Description
allow (default)	Allow users to submit corporate credentials to websites in this URL category.
block	Block users from submitting corporate credentials to websites in this category. A default anti-phishing response page is displayed to users when they access sites to which corporate credential submissions are blocked. You can choose to create a custom block page to display.
continue	Display a response page to users that prompts them to select Continue to access to access the site. By default, the Anti Phishing Continue Page is shown to user when they access sites to which credential submissions are discouraged. You can also choose to create a custom response page to display—for example, if you want to warn users against phishing attempts or reusing their credentials on other websites.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Plan Your URL Filtering Deployment

To deploy URL filtering in your network, we recommend that you start with a basic setup that'll give you visibility into web activity patterns while blocking confirmed malicious content:

- Start with a (mostly) passive URL Filtering profile that alerts on most categories. This gives you visibility into the sites your users are accessing, so you can decide what you want allow, limit, and block.
- Block URL categories that we know are bad: malware, C2, and phishing.

Because alerting on all web activity might create a large amount of log files, you might decide you only want to do this as you're initially deploying URL Filtering.



*At that time, you can also reduce URL filtering logs by enabling the **Log container page only** option in the URL Filtering profile so only the main page that matches the category will be logged, not subsequent pages/categories that may be loaded within the container page.*

STEP 1 | At any time, you can use Test A Site to see how PAN-DB—the URL Filtering cloud database—categorizes a specific URL, and to learn about all possible URL categories.

You can also use Test A Site to submit a change request, if you disagree with how a specific URL is categorized.

STEP 2 | Create a passive URL Filtering profile, that alerts on all categories so you have visibility into web traffic.

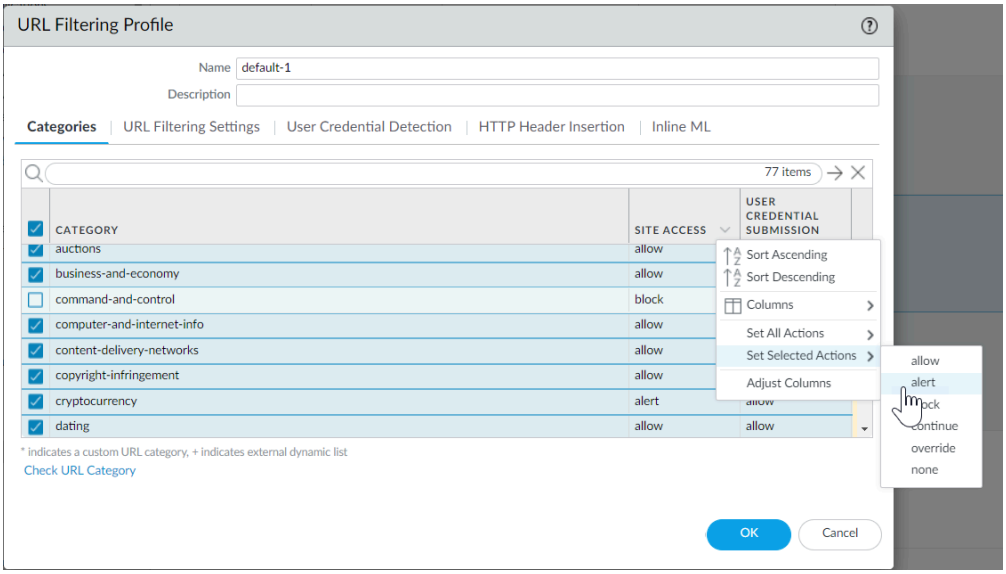
1. Select **Objects** > **Security Profiles> URL Filtering**.
2. Select the default profile and then click **Clone**. The new profile will be named **default-1**.
3. Select the **default-1** profile and rename it. For example, rename it to URL-Monitoring.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 3 | Configure the action for all categories to **alert**, except for malware, command-and-control, and phishing, which should remain blocked.

1. In the section that lists all URL categories, select all categories and then de-select malware, command-and-control, and phishing.
2. To the right of the *Action* column heading, mouse over and select the down arrow and then select **Set Selected Actions** and choose **alert**.



3. **Block** access to known dangerous URL categories.

   Block access to malware, phishing, dynamic-dns, unknown, command-and-control, extremism, copyright-infringement, proxy-avoidance-and-anonymizers, newly-registered-domain, grayware, and parked URL categories.

4. Click **OK** to save the profile.

STEP 4 | Apply the URL Filtering profile to the Security policy rule(s) that allows web traffic for users.

1. Select **Policies** > **Security** and select the appropriate Security policy to modify it.
2. Select the **Actions** tab and in the **Profile Setting** section, click the drop-down for **URL Filtering** and select the new profile.
3. Click **OK** to save.

STEP 5 | Save the configuration.

Click **Commit**.

STEP 6 | View the URL filtering logs to see all of the website categories that your users are accessing. The categories you've set to block are also logged.

For information on viewing the logs and generating reports, see Monitor Web Activity.

Select **Monitor** > **Logs** > **URL Filtering**. A log entry will be created for any website that exists in the URL filtering database that is in a category set to any action other than **allow**. URL Filtering reports give you a view of web activity in a 24-hour period. (**Monitor** > **Reports**).

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 7 | Next Steps:

- PAN-DB categorizes every URL with up to four categories, and every URL has a risk category (high, medium, and low). While high and medium-risk sites are not confirmed malicious, they are closely associated with malicious sites. For example, they might be on the same domain as malicious sites or maybe they hosted malicious content until only very recently. For everything that you do not allow or block, you can use risk categories to write simple policy based on website safety.

  You can take precautionary measures to limit your users' interaction high-risk sites especially, as there might be some cases where you want to give your users access to sites that might also present safety concerns (for example, you might want to allow your developers to use developer blogs for research, yet blogs are a category known to commonly host malware).
- Pair URL Filtering with User-ID to control web access based on organization or department and to block corporate credential submissions to unsanctioned sites:
  - URL Filtering prevents credential theft by detecting corporate credential submissions to sites based on the site category. Block users from submitting credentials to malicious and untrusted sites, warn users against entering corporate credentials on unknown sites or reusing corporate credentials on non-corporate sites, and explicitly allow users to submit credentials to corporate sites.
  - Add or update a Security policy rule with the passive URL Filtering profile so that it applies to a department user group, for example, Marketing or Engineering (**Policies** > **Security** > **User**). Monitor the department activity, and get feedback from department members to understand the web resources that are essential to the work they do.
- Consider all the ways you can use URL Filtering to reduce your attack surface and to control web usage. For example, if you're a school, you can use URL Filtering to enforce strict safe search settings, where search engines filter out adult images and videos from search results. Or, if you have a security operations center, you might give threat analysts password access to compromised or dangerous sites for research, that you might not want to otherwise open up to entire organizations or teams.
- Follow the URL Filtering best practices.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

URL Filtering Best Practices

Palo Alto Networks URL filtering solution protects you from web-based threats, and gives you a simple way to monitor and control web activity. To get the most out of your URL filtering deployment, you should start by creating allow rules for the applications you rely on to do business. Then, review the URL categories that classify malicious and exploitive content—we recommend that you block these outright. Then, for everything else, these best practices can guide you how to reduce your exposure to web-based threats, without limiting your users' access to web content that they need.

- Before you get started, identify the applications you want to allow and create application allow rules as part of building a best practice internet gateway security policy.

  Allowed applications include not only the applications you provision and administer for business and infrastructure purposes, but also the applications that your users need to get their jobs done and applications you might want to allow for personal use.

  After you've identified these sanctioned applications, you can use URL filtering to control and secure all the web activity that is not on the allow list.
- Get visibility in to your users web activity so you can plan the most effective URL filtering policy for your organization, and roll it out smoothly. This includes:
  - Using Test A Site to see how PAN-DB—the Palo Alto Networks URL filtering cloud database—categorizes a specific URL, and to learn about all possible URL categories.
  - Starting with a (mostly) passive URL Filtering profile that alerts on URL categories. This gives you visibility into the sites your users are accessing, so you can decide what you want to allow, limit, and block.
  - Monitoring web activity to assess the sites your users are accessing and see how they align with your business needs.
- Block URL categories that classify malicious and exploitive web content. While we know that these categories are dangerous, always keep in mind that the URL categories that you decide to block might depend on your business needs.
- Use URL categories to phase-in decryption, and to exclude sensitive or personal information (like financial-services and health-and-medicine) from decryption.

  Plan to decrypt the riskiest traffic first (URL categories most likely to harbor malicious traffic, such as gaming or high-risk) and then decrypt more as you gain experience. Alternatively, decrypt the URL categories that don't affect your business first (if something goes wrong, it won't affect business), for example, news feeds. In both cases, decrypt a few URL categories, listen to user feedback, run reports to ensure that decryption is working as expected, and then gradually decrypt a few more URL categories, and so on. Plan to make decryption exclusions to exclude sites from decryption if you can't decrypt them for technical reasons or because you choose not to decrypt them.

  Targeting decryption based on URL categories is also a Decryption *best practice.*
- Prevent credential theft by enabling the firewall to detect corporate credential submissions to sites, and then control those submissions based on URL category. Block users from submitting credentials to malicious and untrusted sites, warn users against entering corporate credentials

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

on unknown sites or reusing corporate credentials on non-corporate sites, and explicitly allow users to submit credentials to corporate and sanctioned sites.

- Block malicious variants of JavaScript exploits and phishing attacks in real-time. Enabling URL Filtering Inline ML allows you to dynamically analyze web pages using machine learning on the firewall.
- Decrypt, inspect, and strictly limit how users interact with high-risk and medium-risk content (if you decided not to block any of the malicious URL categories for business reasons, you should also strictly limit how users interact with those categories).

  The web content that you sanction and the malicious URL categories that you block outright are just one portion of your overall web traffic. The rest of the content your users are accessing is a combination of benign (low-risk) and risky content (high-risk and medium-risk). High-risk and medium-risk content is not confirmed malicious but is closely associated with malicious sites. For example, a high-risk URL might be on the same domain as a malicious site, or maybe it hosted malicious content in the past.

  However, many sites that pose a risk to your organization also provide valuable resources and services to your users (cloud storage services are a good example). While these resources and services are necessary for business, they are also more likely to be used as part of a cyberattack. Here's how to control how users interact with this potentially-dangerous content, while still providing them a good user experience:

  - In a URL Filtering profile, set the high-risk and medium-risk categories to **continue** to display a response page that warns users they're visiting a potentially-dangerous site. Advise them how to take precautions if they decide to continue to the site. If you don't want to prompt users with a response page, alert on the high-risk and medium-risk categories instead.
  - Decrypt decrypt high-risk and medium-risk sites.
  - Follow the Anti-Spyware, Vulnerability Protection, and File Blocking best practices for high-risk and medium-risk sites. A protective measure would be to block downloads of dangerous file types and blocking obfuscated JavaScript.
  - Stop credential theft by blocking users from submitting their corporate credentials to high-risk and medium-risk sites.
- Schools or educational institutions should use safe search enforcement to make sure that search engines filter out adult images and videos from search results. You can even transparently enable safe search for users.
- Enable the firewall to hold an initial web request as it looks up a website's URL category with PAN-DB.

  When a user visits a website, a firewall with Advanced URL Filtering enabled checks its local cache of URL categories to categorize the site. If the firewall doesn't find the URL's category in the cache, it performs a lookup in PAN-DB, the Palo Alto Networks URL database. By default, the firewall allows the user's web request during this cloud lookup and enforces policy when the server responds.

  But when you choose to hold web requests, the firewall blocks the request until it either finds the URL category or times out. If the lookup times out, the firewall considers the URL category not-resolved.

  1. In **Device** > **Setup** > **Content-ID**, check the box for

     Hold client request for category lookup.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

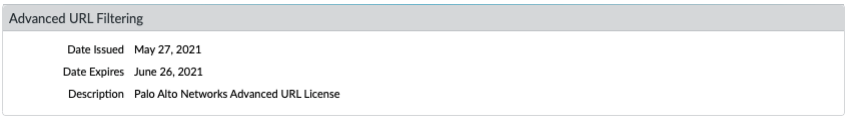
Activate The Advanced URL Filtering Subscription

Palo Alto Networks URL filtering solution, the Advanced URL Filtering subscription, provides real time URL analysis and malware prevention. In addition to PAN-DB access, the Palo Alto Networks-developed URL filtering database for high-performance URL lookups, it also offers coverage against malicious URLs and IP addresses. This multi-layered protection solution is configured through your URL filtering profile.

STEP 1 | Obtain and install an Advanced URL Filtering license and confirm that it is installed.

The Advanced URL Filtering license includes access to PAN-DB; if the license expires, the firewall ceases to perform all URL filtering functions, URL category enforcement, and URL cloud lookups. Additionally, all other cloud based updates will not function until you install a valid license.

1. Select **Device** > **Licenses** and, in the License Management section, select the license installation method:
   - **Retrieve license keys from license server**
   - **Activate feature using authorization code**
2. After installing the license, confirm that the Advanced URL Filtering section, **Date Expires** field, displays a valid date.



When you activate the Advanced URL Filtering license, your license entitlements for PAN-DB and Advanced URL Filtering might not display correctly on the firewall — this is a display anomaly, not a licensing issue, and does not affect access to the services. You can update the licenses on the firewall to rectify the display issue by using the following CLI command: ***`request license fetch`***.

STEP 2 | Download and install the latest PAN-OS content release. PAN-OS Applications and Threats content release 8390-6607 and later allows firewalls operating PAN-OS 9.x and later to identify URLs that have been categorized using the new real-time-detection category, identifying URLs classified by advanced URL filtering. For more information about the update, refer to the Applications and Threat Content Release Notes. You can also review Content Release Notes for apps and threats on the Palo Alto Networks Support Portal or directly in the firewall web interface: select **Device** > **Dynamic Updates** and open the Release Note for a specific content release version.

Follow the Best Practices for Applications and Threats Content Updates when updating to the latest content release version.

https://t.me/learningnets

STEP 3 | Schedule the firewall to download dynamic updates for Applications and Threats.

A Threat Prevention license is required to receive content updates, which covers Antivirus and Applications and Threats.

1. Select **Device** > **Dynamic Updates**.
2. In the Schedule field in the Applications and Threats section, click the **None** link to schedule periodic updates.

   You can only schedule dynamic updates if the firewall has direct Internet access. If updates are already scheduled in a section, the link text displays the schedule settings.

   The Applications and Threats updates sometimes contain updates for URL filtering related to Safe Search Enforcement.

Next Steps:

1. Configure a URL filtering profile to define your organization's web usage policies.
2. Verify Advanced URL Filtering

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Configure URL Filtering

After you determine URL filtering policy requirements, you should have a basic understanding of the types of websites and website categories your users are accessing. Use this information to create custom URL Filtering profiles and attach them to the Security policy rules that allow web access. In addition to managing web access with a URL Filtering profile, if you configure User-ID™, you can manage the sites to which users can submit corporate credentials.

STEP 1 | Create a URL Filtering profile.



If you didn't already, configure a best practice URL Filtering profile to ensure protection against URLs hosting malware or exploitive content.

Select **Objects** > **Security Profiles** > **URL Filtering** and **Add** or modify a URL Filtering profile.

STEP 2 | Define site access for each URL category.

Select **Categories** and set the Site Access for each URL category:

- **allow** traffic destined for that URL category; allowed traffic is not logged.
- Select **alert** to have visibility into sites that users are accessing. Traffic matching that category is allowed but a URL filtering log is generated to record when a user accesses a site in that category.
- Select **block** to deny access to traffic that matches that category and to enable logging of the blocked traffic.
- Select **continue** to display a page to users with a warning and require them to click **Continue** to proceed to a site in that category.
- To only allow access if users provide a configured password, select **override**. For more details, see Allow Password Access to Certain Sites.

STEP 3 | Configure the URL Filtering profile to detect corporate credential submissions to websites that are in allowed URL categories.



To ensure the best performance and a low false positive rate, the firewall automatically skips checking the credential submissions for any App-ID™ associated with sites that have never been observed hosting malware or phishing content—even if you enable checks in the corresponding category. The list of sites for which the firewall skips credential checking is automatically updated through Applications and Threats content updates.

1. Select **User Credential Detection**.
2. Select one of the methods to check for corporate credential submissions to web pages from the **User Credential Detection** drop-down:
   - **Use IP User Mapping**—Checks for valid corporate username submissions and verifies that the username matches the user logged in to the source IP address of the session. The firewall matches the submitted username against its IP address-to-username

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

mapping table. You can use any of the user mapping methods described in Map IP Addresses to Users.

- **Use Domain Credential Filter**—Checks for valid corporate usernames and password submissions and verifies that the username maps to the IP address of the logged-in user. See Configure User Mapping Using the Windows User-ID Agent for instructions on how to set up User-ID to enable this method.
- **Use Group Mapping**—Checks for valid username submissions based on the user-to-group mapping table populated when you configure the firewall to map users to groups.

  With group mapping, you can apply credential detection to **any** part of the directory or to a specific group, such as groups like IT that have access to your most sensitive applications.

  This method is prone to false positives in environments that do not have uniquely structured usernames, so you should only use this method to protect your high-value user accounts.

3. Set the **Valid Username Detected Log Severity** that the firewall uses to log detection of corporate credential submissions (default is medium).

STEP 4 | Configure the URL Filtering profile to detect phishing and malicious JavaScript in real-time using URL Filtering Inline ML.

STEP 5 | Allow or block users from submitting corporate credentials to sites based on URL category to prevent credential phishing.

To ensure the best performance and a low false positive rate, the firewall automatically skips checking the credential submissions for any App-ID associated with sites that have never been observed hosting malware or phishing content—even if you enable checks in the corresponding category. The list of sites for which the firewall skips credential checking is automatically updated through Applications and Threats content updates.

1. For each URL category to which you allow **Site Access**, select how you want to treat **User Credential Submissions**:
   - **alert**—Allow users to submit credentials to the website but generate a URL filtering alert log each time a user submits credentials to sites in this URL category.
   - **allow** (default)—Allow users to submit credentials to the website.
   - **block**—Displays the Anti Phishing Block Page to block users from submitting credentials to the website.
   - **continue**—Present the Anti Phishing Continue Page to require users to click **Continue** to access the site.
2. Configure the URL Filtering profile to detect corporate credential submissions to websites that are in allowed URL categories.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 6 | Define URL category exception lists to specify websites that should always be blocked or allowed, regardless of URL category.

For example, to reduce URL filtering logs, you may want to add your corporate websites to the allow list so that no logs are generated for those sites or, if there is a website that is being overly used and is not work-related, you can add that site to the block list.

The policy actions configured for custom URL categories have priority enforcement over matching URLs in external dynamic lists.

Traffic to websites in the block list is always blocked regardless of the action for the associated category and traffic to URLs in the allow list is always allowed.

For more information on the proper format and wildcard usage, review the URL category exception list guidelines.

STEP 7 | Enable Safe Search Enforcement.

STEP 8 | Log only Container Pages for URL filtering events.

1. Select **URL Filtering Settings**. Enable **Log container page only** (default) so that the firewall logs only the main page that matches the category, not subsequent pages or categories that loaded within the container page.
2. To enable logging for all pages and categories, disable the **Log container page only** option.

STEP 9 | Enable HTTP Header Logging for one or more of the supported HTTP header fields.

Select **URL Filtering Settings** and select one or more of the following fields to log:

- **User-Agent**
- **Referer**
- **X-Forwarded-For**

STEP 10 | Save the URL Filtering profile and commit your changes.

1. Click **OK**.
2. Click **Commit**.

STEP 11 | Test your URL filtering policy configuration.

1. Access a website in the desired URL category and observe the firewall's behavior.

   Use Palo Alto Networks URL Filtering Test Pages (urlfiltering.paloaltonetworks.com/test-<*url-category*>) if you want to avoid directly accessing a site. Palo Alto Networks has test URLs for benign and malicious URL categories. For example, to test your block policy for malware, visit https://urlfiltering.paloaltonetworks.com/test-malware.
2. Review the Traffic and URL filtering logs (**Monitor** > **Logs**) to confirm that the correct policy rule is logged.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 12 | Enable **Hold client request for category lookup** to block client requests while the firewall performs URL category lookups.

1. Select **Device** > **Setup** > **Content-ID**.
2. Select **Hold client request for category lookup**.
3. Commit your changes.

Enable this feature as a URL Filtering best practice.

STEP 13 | Set the amount of time, in seconds, before a URL category lookup times out.

1. Select **Device** > **Setup** > **Content-ID** > **gear icon**.
2. Enter a number in **Category lookup timeout (sec)**.
3. Click **OK**.
4. Commit your changes.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Test URL Filtering Configuration

To test your URL Filtering and Advanced URL Filtering policy configurations, use Palo Alto Networks URL Filtering Test Pages. Test pages have been created for the safe testing of all predefined URL categories, including real-time-detection categories applicable only to firewalls running Advanced URL Filtering.

You must enable SSL decryption for test pages to work over an HTTPS connection.

Advanced URL filtering test pages contain "real-time-detection" in the URL and confirm that firewalls correctly categorize and analyze malicious URLs in real-time. They do not verify firewall behavior for any other categories.

You can check the classification of a specific website using Palo Alto Networks URL category lookup tool, Test A Site.

Follow the procedure corresponding to your URL Filtering subscription:

- Verify URL Filtering
- Verify Advanced URL Filtering

Verify URL Filtering

If you have the legacy URL Filtering subscription, test and verify that the firewall correctly categorizes, enforces, and logs URLs in the categories that you access.

STEP 1 | Access a website in the URL category of interest.

Consider testing sites in blocked URL categories. You can use a test page (urlfiltering.paloaltonetworks.com/test-*<url-category>*) to avoid directly accessing a site. For example, to test your block policy for malware, visit https://urlfiltering.paloaltonetworks.com/test-malware.

STEP 2 | Review the Traffic and URL Filtering logs (**Monitor** > **Logs**) to verify that your firewall processes the site correctly.

For example, if you configured a block page to display when someone accesses a site that violates your organization's policy, check that one appears when you visit the test site.

Verify Advanced URL Filtering

If you have an Advanced URL Filtering subscription, test and verify that real-time URL analysis is happening.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Palo Alto Networks recommends setting the real-time-detection action setting to alert for your active URL filtering profiles. This provides visibility into URLs analyzed in real-time and will block (or allow, depending on your policy settings) based on the category settings configured for specific web threats.

The firewall enforces the most severe action of the actions configured for detected URL categories of a given URL. For example, suppose example.com is categorized as real-time-detection, command-and-control, and shopping—categories with an alert, block, and allow action configured, respectively. The firewall will block the URL because block is the most severe action from the detected categories.

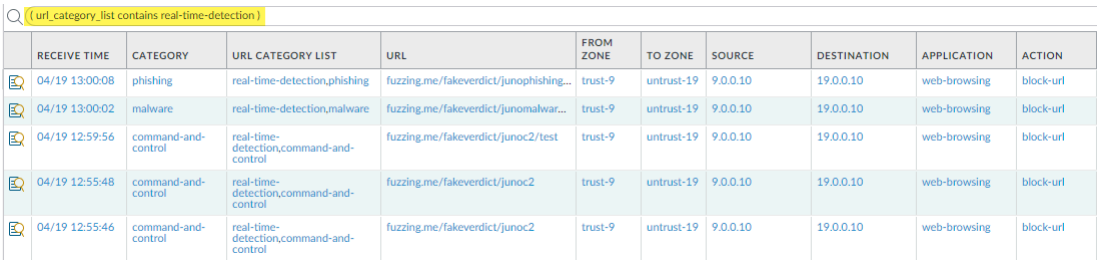
STEP 1 | Visit each of the following test URLs to verify that the Advanced URL Filtering service is properly categorizing URLs:

- **Malware**—urlfiltering.paloaltonetworks.com/test-real-time-detection-malware
- **Phishing**—urlfiltering.paloaltonetworks.com/test-real-time-detection-phishing
- **C2**—urlfiltering.paloaltonetworks.com/test-real-time-detection-command-and-control
- **Grayware**—urlfiltering.paloaltonetworks.com/test-real-time-detection-grayware
- **Benign (unknown)**—urlfiltering.paloaltonetworks.com/test-real-time-detection

STEP 2 | Monitor the activity on the firewall to verify that the tested URLs have been properly categorized as real-time-detection.

1. Select **Monitor** > **Logs** > **URL Filtering** and filter by `(url_category_list contains real-time-detection)` to view logs that have been analyzed using Advanced URL Filtering.

   Additional web page category matches are also displayed and corresponds to the categories as defined by PAN-DB.



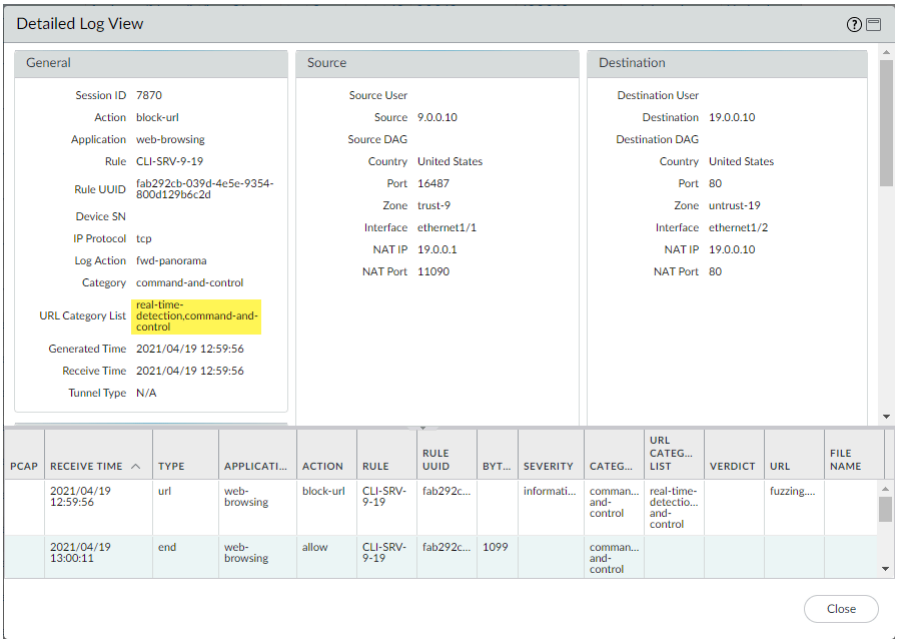
2. Take a detailed look at the logs to verify that each type of web threat is correctly analyzed and categorized.

   In the next example, the URL is categorized as having been analyzed in real-time and possessing qualities that define it as command-and-control (C2). Because the C2

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

category has a more severe action associated with it than real-time-detection (block as opposed to alert), the URL is categorized as command-and-control and blocked.



https://t.me/learningnets

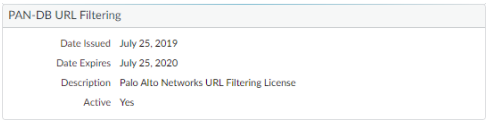
Configure URL Filtering Inline ML

To enable your URL Filtering inline ML configuration, attach the URL Filtering profile configured with the inline ML settings to a Security policy rule (see Set Up a Basic Security Policy).

URL Filtering inline ML is not currently supported on the VM-50 or VM50L virtual appliance.

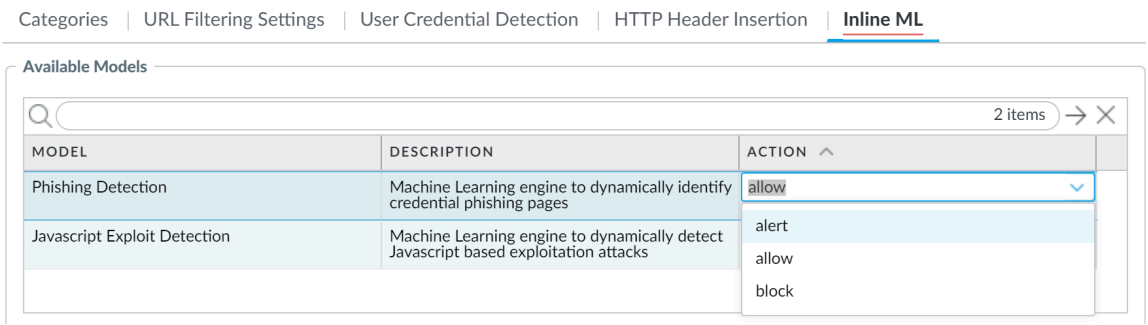
STEP 1 | To take advantage of URL Filtering inline ML, you must have an active Advanced URL Filtering or legacy URL Filtering subscription.

Verify that you have a URL Filtering subscription. To verify subscriptions for which you have currently-active licenses, select **Device** > **Licenses** and verify that the appropriate licenses display and are not expired. The image below shows the license entry for the legacy URL Filtering subscription.



STEP 2 | Create a new or update your existing URL Filtering profiles to use URL Filtering inline ML.

1. Select an existing **URL Filtering Profile** or **Add** a new one (**Objects** > **Security Profiles** > **URL Filtering**).
2. Select **Inline ML** and define a policy **Action** for each URL Filtering inline ML model. This enforces the selected policy action on a per model basis. Currently, there are two classification engines available: **Phishing** and **JavaScript Exploit**, one for each type of malicious webpage content.
   - **Block**—When the firewall detects a website with phishing content, the firewall generates a URL Filtering log entry.
   - **Alert**—The firewall allows access to the website and generates a URL Filtering log entry.
   - **Allow**—The firewall allows access to the website but does not generate a URL Filtering log entry.



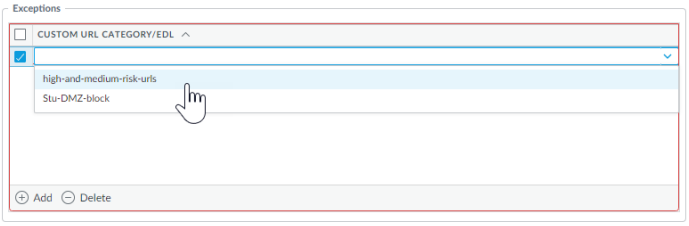
3. Click **OK** to exit the URL Filtering Profile dialog, then **Commit** your changes.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 3 | (Optional) Add URL exceptions to your URL Filtering profile if you encounter false-positives. You can add exceptions by specifying an external dynamic list from the URL Filtering profile or by adding a web page entry from the URL Filtering logs.

- Add an external dynamic list of URL exceptions.
  1. Select **Objects > Security Profiles > URL Filtering**.
  2. Select a URL Filtering profile for which you want to exclude specific URLs, then select **Inline ML**.
  3. **Add** a pre-existing URL-based external dynamic list. If none is available, create a new external dynamic list.



  4. Click **OK** to save the URL Filtering profile and **Commit** your changes.
- Add file exceptions from URL Filtering log entries.
  1. Select **Monitor > Logs > URL Filtering** and filter the logs for URL entries with an Inline ML Verdict of **malicious-javascript** or **phishing**. Select a URL Filtering log for a URL that you wish to create an exception for.
  2. Go to the **Detailed Log View** and scroll down to the **Details** pane, then select **Create Exception** located next to the **Inline ML Verdict**.



  3. Select a custom category for the URL exception and click **OK**.
  4. The new URL exception can be found in the list to which it was added, under **Objects > Custom Objects > URL Category**.

STEP 4 | (Optional) Verify the status of your firewall's connectivity to the inline ML cloud service.

Use the following CLI command on the firewall to view the connection status.

```
show mlav cloud-status
```

For example:

```
show mlav cloud-status

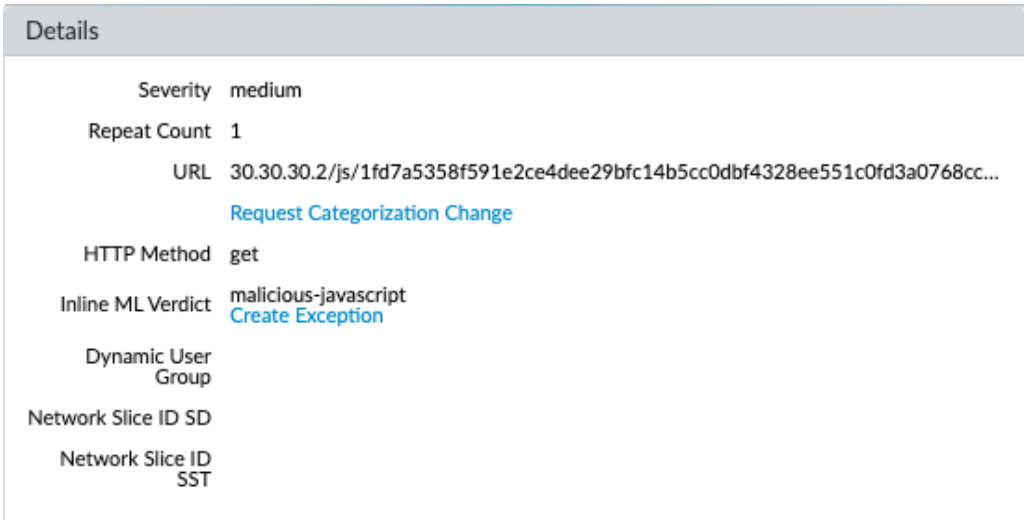
MLAV cloud
Current cloud server:          ml.service.paloaltonetworks.com
Cloud connection:              connected
```

If you are unable to connect to the inline ML cloud service, verify that the following domain is not being blocked: ml.service.paloaltonetworks.com.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

To view information about web pages that have been processed using URL Filtering inline ML, Filter the logs (**Monitor > Logs > URL Filtering**) based on **Inline ML Verdict**. Web pages that have been determined to contain threats are categorized with verdicts of either **phishing** or **malicious-javascript**. For example:



 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Monitor Web Activity

The ACC, URL filtering logs and reports show all user web activity for URL categories that are set to **alert**, **block**, **continue**, or **override**. By monitoring the logs, you can gain a better understanding of the web activity of your user base to determine a web access policy.

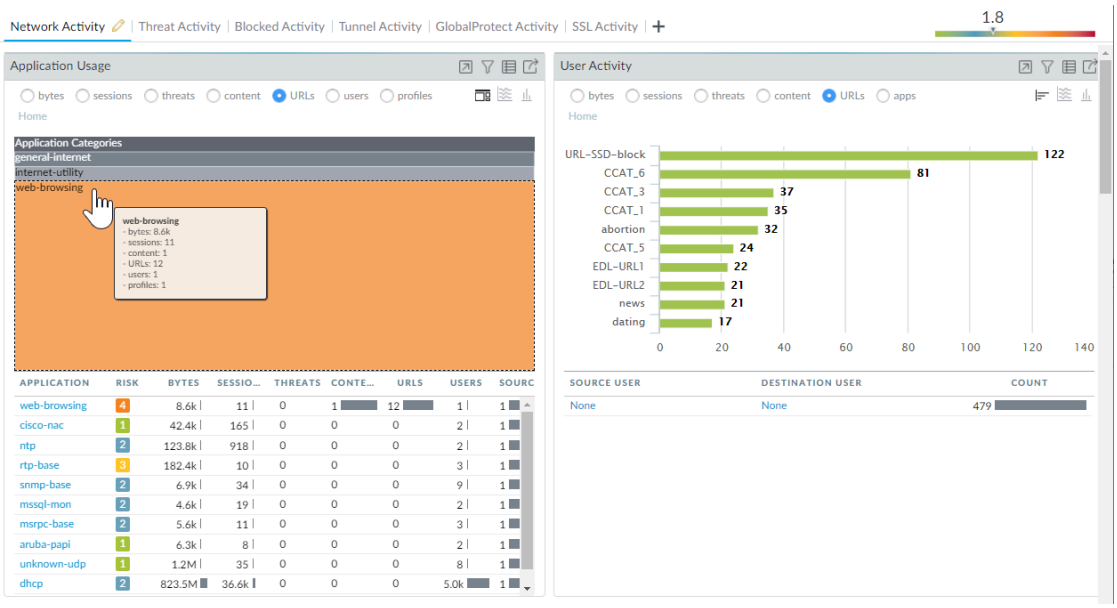
The following topics describe how to monitor web activity:

- Monitor Web Activity of Network Users
- View the User Activity Report
- Configure Custom URL Filtering Reports

Monitor Web Activity of Network Users

You can use the Application Command Center (ACC), URL filtering reports, and logs that are generated on the firewall to track user activity.

- For a quick view of the most common categories users access in your environment, check the **ACC** widgets. Most **Network Activity** widgets allow you to sort on URLs. For example, in the Application Usage widget, you can see that the networking category is the most accessed category, followed by encrypted tunnel, and ssl. You can also view the list of **Threat Activity** and **Blocked Activity** sorted on URLs.



View logs and configure log options:

https://t.me/learningnets

- From the ACC, you can jump directly to the logs ( ) or select **Monitor** > **Logs** > **URL Filtering**.

  The log action for each entry depends on the Site Access setting you defined for the corresponding category:

  - **Alert log**—In this example, the computer-and-internet-info category is set to alert.



	RECEIVE TIME	CATEGORY	URL	FROM ZONE	TO ZONE	SOURCE	SOURCE USER	DESTINATION	APPLICATION	ACTION
	2020/04/16 14:10:53	computer-and-internet-info	outlook.office36...	pm wifi	UNTRUST	[redacted]	[redacted]	[redacted]	outlook-web-online	alert

  - **Block log**—In this example, the insufficient-content category is set to continue. If the category had been set to block instead, the log Action would be block-url.



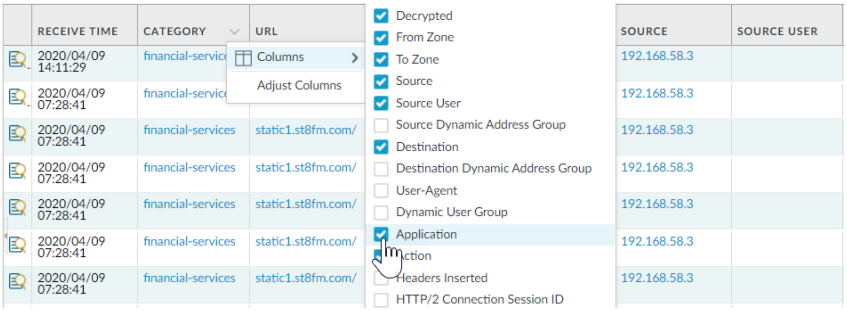
	RECEIVE TIME	CATEGORY	URL	FROM ZONE	TO ZONE	SOURCE	SOURCE USER	DESTINATION	APPLICATION	ACTION
	2020/04/08 18:47:49	insufficient-content	munchkin.mark...	pm wifi	UNTRUST	[redacted]	[redacted]	[redacted]	ssl	block-continue

  - **Alert log on encrypted website**—In this example, the category is private-ip-addresses and the application is web-browsing. This log also indicates that the firewall decrypted this traffic.



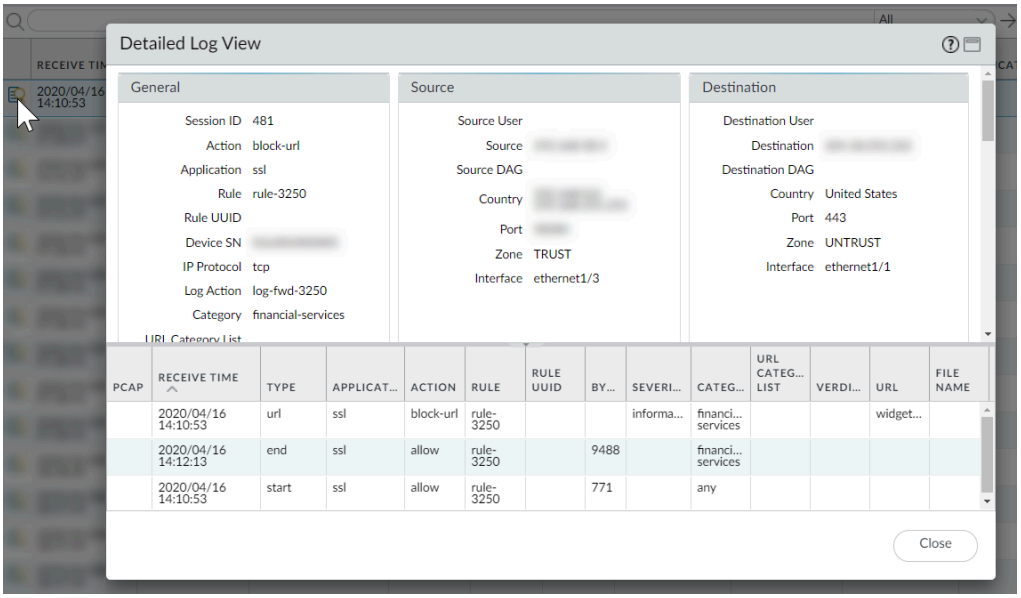
	RECEIVE TIME	CATEGORY	URL	DECRYPTED	FROM ZONE	TO ZONE	SOURCE	SOURCE USER	DESTINATION	APPLICATION	ACTION
	2020/04/09 14:11:29	private-ip-addresses	/Updates/Updat...	yes	TRUST	UNTRUST	192.168.58.3	[redacted]	[redacted]	web-browsing	alert

- You can also add several other columns to your URL Filtering log view, such as: to and from zone, content type, and whether or not a packet capture was performed. To modify what columns to display, click the down arrow in any column and select the attribute to display.



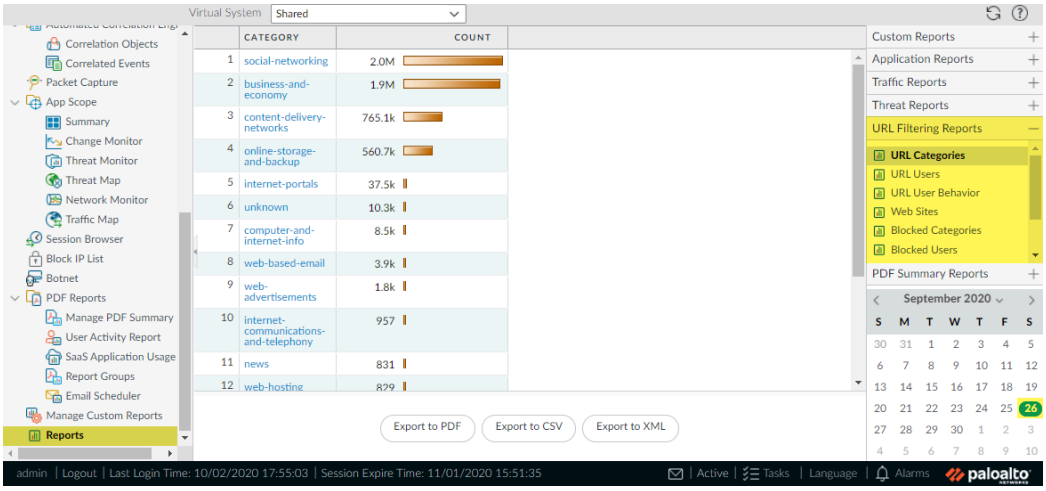
https://t.me/learningnets

- To view the complete log details and/or request a category change for the given URL that was accessed, click the log details icon in the first column of the log.



- Generate predefined URL filtering reports on URL categories, URL users, Websites accessed, Blocked categories, and more.

  Select **Monitor** > **Reports** and under the **URL Filtering Reports** section, select one of the reports. The reports cover the 24-hour period of the date you select on the calendar. You can also export the report to PDF, CSV, or XML.



View the User Activity Report

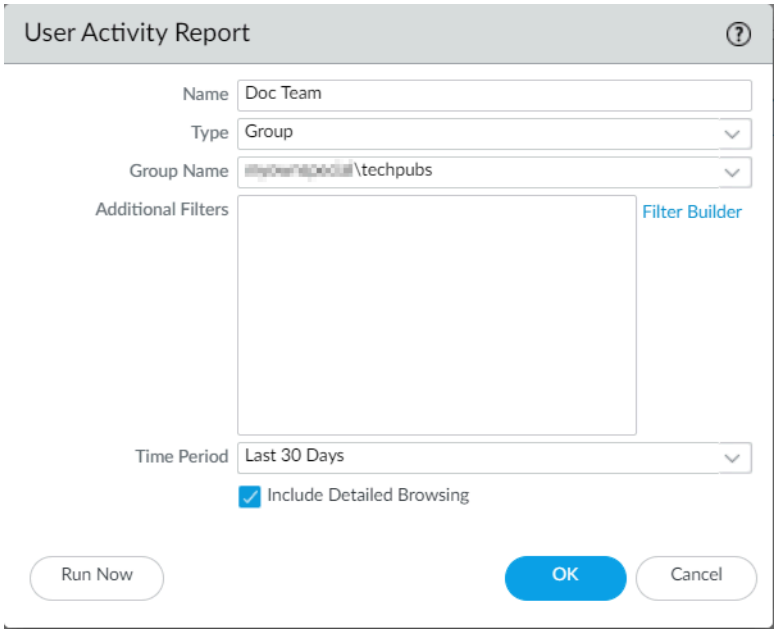
This report provides a quick method of viewing user or group activity and also provides an option to view browse time activity.

https://t.me/learningnets

STEP 1 | Configure a User Activity Report.

1. Select **Monitor** > **PDF Reports** > **User Activity Report**.
2. **Add** a report and enter a **Name** for it.
3. Select the report **Type**:
   - Select **User** to generate a report for one person.
   - Select **Group** for a group of users.

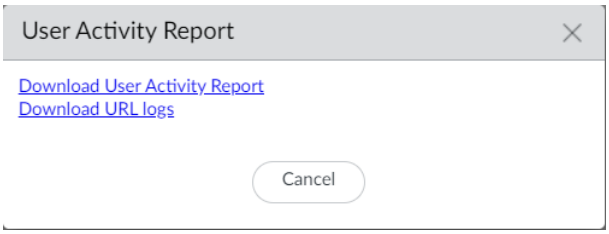
   *You must enable User-ID in order to be able to select user or group names. If User-ID is not configured, you can select the type **User** and enter the IP address of the user's computer.*
4. Enter the **Username/IP Address** for a user report or enter the group name for a user group report.
5. Select the time period. You can select an existing time period, or select **Custom**.
6. Select the **Include Detailed Browsing** check box, so browsing information is included in the report.



https://t.me/learningnets

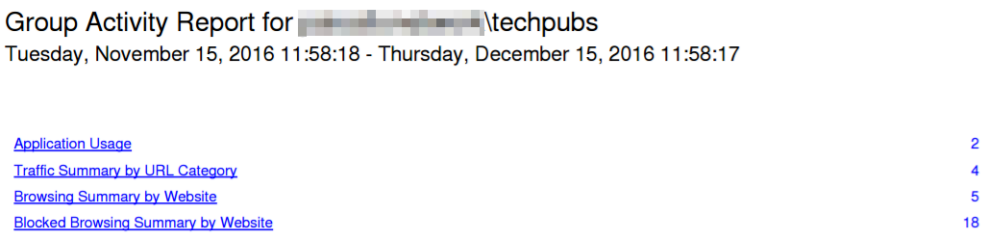
STEP 2 | Run the report.

1. Click **Run Now**.
2. When the firewall finishes generating report, click one of the links to download it:
   - Click **Download User Activity Report** to download a PDF version of the report.
   - Click **Download URL Logs** to download a CSV file of the corresponding log entries.

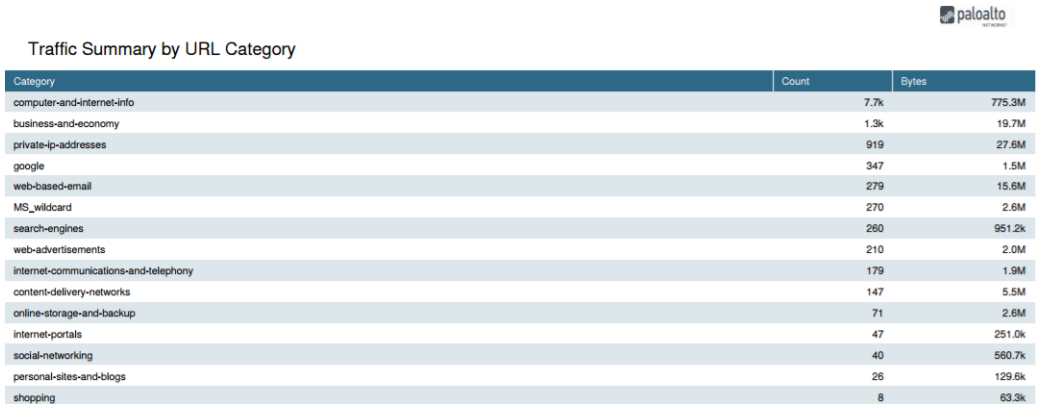


3. After downloading the report, click **Cancel**.
4. If you want to save the user activity report settings so you can run the same report again later, click **OK**; otherwise click **Cancel**.

STEP 3 | View the user activity report by opening the file that you downloaded. The PDF version of the report shows the user or group on which you based the report, the report time frame, and a table of contents:



STEP 4 | Click an item in the table of contents to view the report details. For example, click **Traffic Summary by URL Category** to view statistics for the selected user or group.



Category	Count	Bytes
computer-and-internet-info	7.7k	775.3M
business-and-economy	1.3k	19.7M
private-ip-addresses	919	27.6M
google	347	1.5M
web-based-email	279	15.6M
MS_wildcard	270	2.6M
search-engines	260	951.2k
web-advertisements	210	2.0M
internet-communications-and-telephony	179	1.9M
content-delivery-networks	147	5.5M
online-storage-and-backup	71	2.6M
internet-portals	47	251.0k
social-networking	40	560.7k
personal-sites-and-blogs	26	129.6k
shopping	8	63.3k

Configure Custom URL Filtering Reports

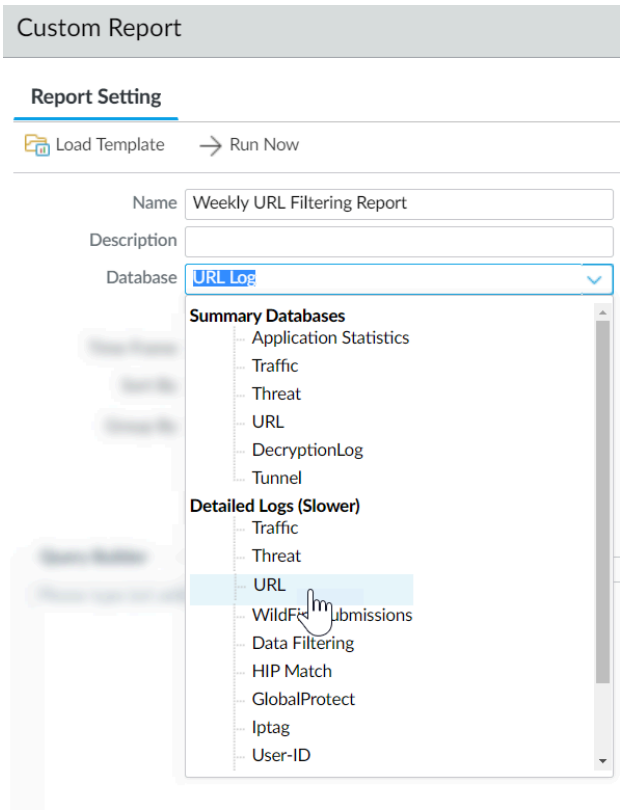
To generate a detailed report that you can schedule to run regularly, configure a custom URL Filtering report. You can choose any combination of URL Filtering log fields on which to base the report.

©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 1 | Add a new custom report.

1. Select **Monitor** > **Manage Custom Reports** and **Add** a report.
2. Give the report a unique **Name**, and optionally a **Description**.
3. Select the **Database** you want to use to generate the report. To generate a detailed URL Filtering report, select **URL** from the Detailed Logs section:

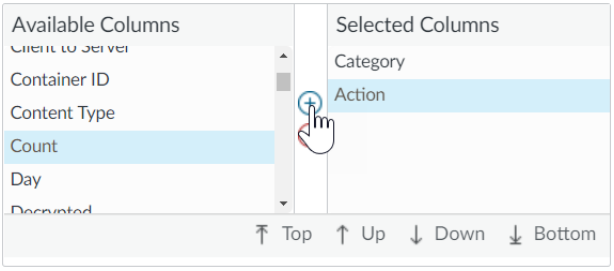


 ©2022 Palo Alto Networks, Inc.

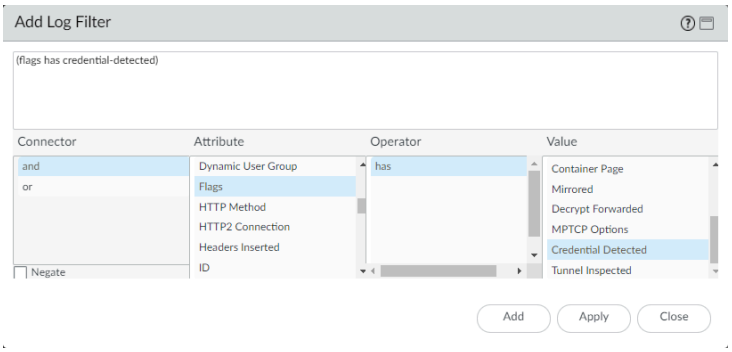
https://t.me/learningnets

STEP 2 | Configure report options.

1. Select a predefined **Time Frame** or select **Custom**.
2. Select the log columns to include in the report from the Available Columns list add them ( ) to the Selected Columns. For example, for a URL Filtering report you might select:
   - Action
   - App Category
   - Category
   - Destination Country
   - Source User
   - URL



3. If the firewall is enabled to prevent credential phishing, select the Attribute **Flags**, the Operator **has** and the Value **Credential Detected** to also include events in the report that record when a user submitted a valid corporate credential to a site.

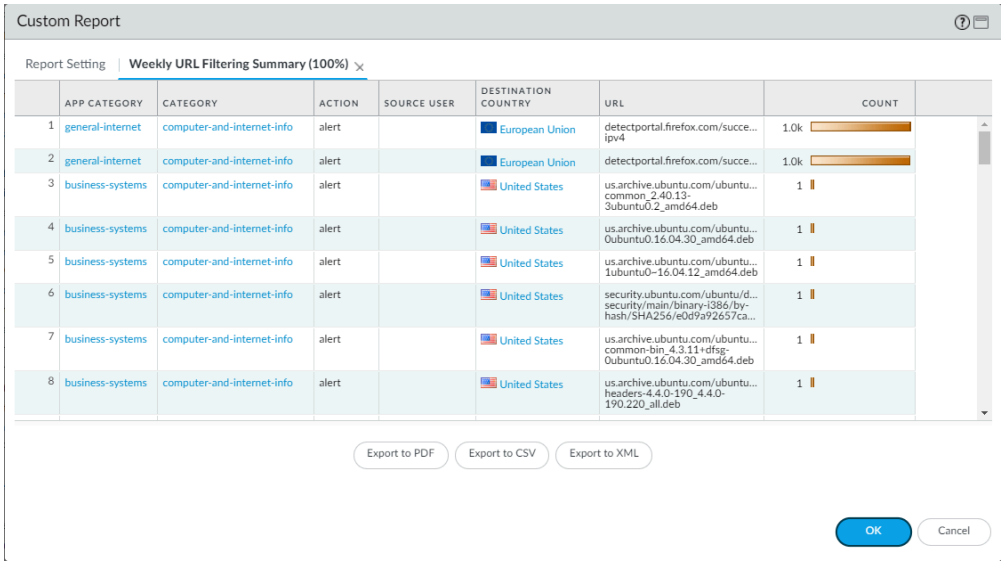


4. (Optional) Select a **Sort By** option to set the attribute to use to aggregate the report details. If you do not select an attribute to sort by, the report will return the first N number of results without any aggregation. Select a **Group By** attribute to use as an

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

anchor for grouping data. The following example shows a report with **Group By** set to **App Category** and **Sort By** set to a **Count** of **Top 5**.



STEP 3 | Run the report.

1. Click the **Run Now** icon to immediately generate the report, which opens in a new tab.
2. When you are done reviewing the report, go back to the **Report Setting** tab and either tune the settings and run the report again, or continue to the next step to schedule the report.
3. Select the **Schedule** check box to run the report once per day. This will generate a daily report that details web activity over the last 24 hours.

STEP 4 | **Commit** the configuration.

STEP 5 | View the custom report.

1. Select **Monitor** > **Reports**.
2. Expand the **Custom Reports** pane in the right column and select the report you want to view. The latest report displays automatically.
3. To view the report for a previous date, select the date from the calendar. You can also export the report to PDF, CSV, or XML.

https://t.me/learningnets

Log Only the Page a User Visits

A container page is the main page that a user accesses when visiting a website, but additional pages might be loaded along with the main page. If the **Log Container page only** option is enabled in a URL Filtering profile (**Objects** > **Security Profiles** > **URL Filtering**), only the main container page will be logged, not subsequent pages that may be loaded within the container page. Because URL filtering can potentially generate a lot of log entries, you may want to turn on this option, so log entries will only contain those URIs where the requested page file name matches the specific mime-types. The default set includes the following mime-types:

- application/pdf
- application/soap+xml
- application/xhtml+xml
- text/html
- text/plain
- text/xml



*If you enable the **Log container page only** option, there may not always be a correlated URL log entry for threats detected by antivirus or vulnerability protection.*

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Create a Custom URL Category

You can create a custom URL filtering object to specify exceptions to URL category enforcement, and to create a custom URL category based on multiple URL categories:

- **Define exceptions to URL category enforcement**—Create a custom list of URLs that you want to use as match criteria in a Security policy rule. This is a good way to specify exceptions to URL categories, where you'd like to enforce specific URLs differently than the URL category to which they belong.
- **Define a custom URL category based on multiple PAN-DB categories**—This allows you to target enforcement for websites that match a set of categories. The website or page must match *all* the categories defined as part of the custom category.

Follow these steps to create a custom URL category, and define how you'd like the firewall to enforce the custom URL category:

STEP 1 | Select **Objects** > **Custom Objects** > **URL Category**.

STEP 2 | **Add** or modify a custom URL category, and give the category a descriptive **Name**.

STEP 3 | Set the category **Type** to either **Category Match** or **URL List**:

- **URL List**—Add URLs that you want to enforce differently than the URL category to which they belong. Use this list type to define exceptions for URL category enforcement, or to define a list of URLs as belonging to a custom category. For details on how to populate this list, such as guidelines on how to use wildcards, see URL Category Exceptions.
- **Category Match**—Provide targeted enforcement for websites that match a set of categories. The website or page must match *all* the categories defined as part of the custom category.

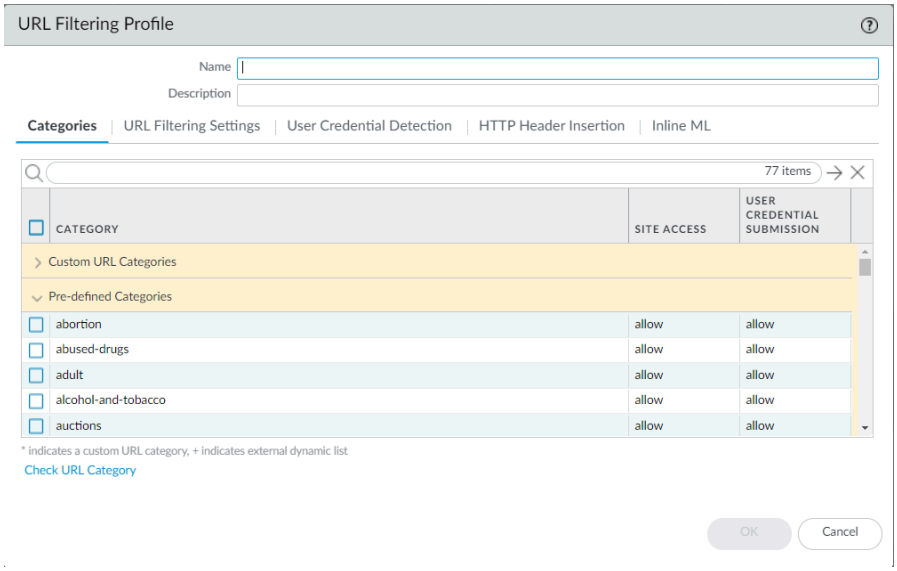
STEP 4 | Select **OK** to save the custom URL category.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 5 | Select **Objects** > **Security Profiles** > **URL Filtering** and **Add** or modify a URL Filtering profile.

Your new custom category will be listed under the **Custom URL Categories** drop down:



STEP 6 | Decide how you want to enforce **Site Access** and **User Credential Submissions** for the custom URL category. (To control the sites to which users can submit their corporate credentials, see Prevent Credential Phishing).

STEP 7 | Attach the URL Filtering profile to a Security policy rule, to enforce traffic that matches that rule.

Select **Policies** > **Security** > **Actions** and specify the Security policy rule to enforce traffic based on the URL Filtering profile you just updated. Make sure to **Commit** your changes.

*You can also use custom URL categories as Security policy match criteria. In this case, you do not need to define how the category should be enforced as part of a URL Filtering profile. After setting up the custom category, go directly to the Security policy rule to which you want to add the custom URL category (**Policies** > **Security**). Select **Service/URL Category** to use the custom URL category as match criteria for the rule.*

©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

URL Category Exceptions

You can exclude specific websites from URL category enforcement, ensuring that these websites are blocked or allowed regardless of the policy action associated with its URL categories. For example, you might block the social-networking URL category but allow access to LinkedIn. To create exceptions to URL category policy enforcement:

- Add the IP addresses or URLs of the sites you want to explicitly block or allow to a custom URL category list (**Objects** > **Custom Objects** > **URL Category**). Then, define site access for the custom URL category in a URL Filtering profile. Finally, attach the profile to a Security policy rule.

  *You can also use a custom URL category as match criteria in a Security policy rule (**Policies** > **Security**, and select **Service/URL Category**). The exception rule must be placed above any rules that block or allow the categories to which the URL exceptions belong.*

- Use an external dynamic list in a URL Filtering profile or as match criteria in a Security policy rule. The benefit to using an external dynamic list is that you can update the list without performing a configuration change or commit on the firewall.

The following guidelines describe how to populate URL category block and allow lists, or a text file that you're using as the source of an external dynamic list for URLs:

- Basic Guidelines For URL Category Exception Lists
- Wildcard Guidelines for URL Category Exception Lists
- URL Category Exception List—Wildcard Examples

Basic Guidelines For URL Category Exception Lists

- Enter the IP addresses or URLs of websites that you want to enforce separately from the associated URL category.
- List entries must be an exact match and are case-insensitive.
- Enter a string that is an exact match to the website (and possibly, specific subdomain) for which you want to control access, or use wildcard characters to allow an entry to match to multiple website subdomains. For details on using wildcard characters, review Wildcard Guidelines for URL Category Exception Lists.
- Omit `http` and `https` from URL entries.
- Each URL entry can be up to 255 characters in length.

Wildcard Guidelines for URL Category Exception Lists

You can use wildcards in URL category exception lists to easily configure a single entry to match to multiple website subdomains and pages, without having to specify exact subdomains and pages.

Follow these guidelines when creating wildcard entries:

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

- The following characters are considered token separators: . / ? & = ; +

  Every string separated by one or two of these characters is a token. Use wildcard characters as token placeholders, indicating that a specific token can contain any value.
- In place of a token, use either an asterisk (*) or a caret (^) to indicate a wildcard value.
- Wildcard characters must be the only character within a token. For example, www.gmail*.com would be invalid because the asterisk follows other characters. An entry can contain multiple wildcards, however.

How to Use Asterisk (*) and Caret (^) Wildcards

*	Use to indicate one or more variable subdomains. If you use *, the entry will match any additional subdomains, whether at the beginning or the end of the URL. Ex: • **`*.paloaltonetworks.com`** matches www.paloaltonetworks.com and www.paloaltonetworks.com.uk. • **`*.paloaltonetworks.com/`** matches www.paloaltonetworks.com but not www.paloaltonetworks.com.uk.
^	Use to indicate one variable subdomain. Ex: **`mail.^.com`** matches to mail.company.com but not mail.company.sso.com.

Do not create an entry with consecutive asterisk (*) wildcards or more than nine consecutive caret (^) wildcards—entries like these can affect firewall performance.

*For example, do not add an entry like **`mail.*.*.com`**; instead, depending on the range of websites you want to control access to, enter **`mail.*.com`** or **`mail.^.^.com`**. An entry like **`mail.*.com`** matches to a greater number of sites than **`mail.^.^.com`**; **`mail.*.com`** matches to sites with any number of subdomains and **`mail.^.^.com`** matches to sites with exactly two subdomains.*

URL Category Exception List—Wildcard Examples

The following table displays example URL list entries using wildcards and sites matching these entries.

URL Exception List Entry	Matching Sites
Example Set 1	

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

URL Exception List Entry	Matching Sites
*.company.com	eng.tools.company.com support.tools.company.com tools.company.com docs.company.com
^.company.com	tools.company.com docs.company.com
^.^.company.com	eng.tools.company.com support.tools.company.com
Example Set 2	
mail.google.*	mail.google.com mail.google.co.uk mail.google.example.org
mail.google.^	mail.google.com mail.google.info
mail.google.^.^	mail.google.co.uk mail.google.example.info
Example Set 3	
site.*.com	site.yourname.com site.abc.xyz.com
site.^.com	site.company.com site.example.com
site.^.^.com	site.a.b.com
site.com/*	site.com/photos site.com/blog/latest any site.com subdirectory

©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Use an External Dynamic List in a URL Filtering Profile

To protect your network from newly-discovered threats and malware, you can use external dynamic lists in URL Filtering profiles. External dynamic lists give you the ability to update the list without a configuration change or commit on the firewall. An external dynamic list is a text file that is hosted on an external web server. You can use this list to import URLs and enforce policy on these URLs. When the list is updated on the web server, the firewall retrieves the changes and applies policy to the modified list without requiring a commit on the firewall.

The firewall dynamically imports the list at the configured interval and enforces policy for the URLs (IP addresses or domains are ignored) in the list. For URL formatting guidelines, see URL Category Exceptions.

For more information, see External Dynamic List.

STEP 1 | Configure the firewall to access an external dynamic list.

- Ensure that the list does not include IP addresses or domain names; the firewall skips non-URL entries.
- Use the custom URL list guidelines to verify the list's formatting.
- Select **URL List** from the Type drop-down.

STEP 2 | Use the external dynamic list in a URL Filtering profile.

1. Select **Objects** > **Security Profiles** > **URL Filtering**.
2. **Add** or modify an existing URL Filtering profile.
3. **Name** the profile and, in the **Categories** tab, select the external dynamic list from the Category list.
4. Click Action to select a more granular action for the URLs in the external dynamic list.

   If a URL that is included in an external dynamic list is also included in a custom URL category, or block and allow list, the action specified in the custom category or the block and allow list will take precedence over the external dynamic list.
5. Click **OK**.
6. Attach the URL Filtering profile to a Security policy rule.
   1. Select **Policies** > **Security**.
   2. Select the **Actions** tab and, in the Profile Setting section, select the new profile in the **URL Filtering** drop-down.
   3. Click **OK** and **Commit**.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 3 | Test that the policy action is enforced.

1. View the external dynamic list entries and try to access a URL from the list.
2. Verify that the action you defined is enforced in the browser.
3. To monitor the activity on the firewall:
   1. Select **ACC** and add a URL Domain as a global filter to view the Network Activity and Blocked Activity for the URL you accessed.
   2. Select **Monitor** > **Logs** > **URL Filtering** to access the detailed log view.

STEP 4 | Verify whether entries in the external dynamic list were ignored or skipped.

In a list of type URL, the firewall skips non-URL entries as invalid and ignores entries that exceed the maximum limit for the firewall model.



*To check whether you have reached the limit for an external dynamic list type, select **Objects** > **External Dynamic Lists** and click **List Capacities**.*

Use the following CLI command on a firewall to review the details for a list.

```
request system external-list show type url name <list_name>
```

For example:

```
request system external-list show type url name My_URL_List
vsys5/My_URL_List:
Next update at: Tue Jan 3 14:00:00 2017
Source: http://example.com/My_URL_List.txt
Referenced: Yes
Valid: Yes
Auth-Valid: Yes

Total valid entries: 3
Total invalid entries: 0
Valid urls:
www.URL1.com
www.URL2.com
www.URL3.com
```

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Allow Password Access to Certain Sites

In some cases, there may be URL categories that you want to block but allow certain individuals to access on occasion. In this case, you would set the category action to **override** and define a URL admin override password in the firewall Content-ID configuration. Users will be required to provide the override password before they can access sites in these categories. Use the following procedure to configure URL admin override:

STEP 1 | Set the URL admin override password.

1. Select **Device** > **Setup** > **Content ID**.
2. In the **URL Admin Override** section, click **Add**.
3. In the **Location** field, select the virtual system to which this password applies.
4. Enter the **Password** and **Confirm Password**.
5. Select an **SSL/TLS Service Profile**. The profile specifies the certificate that the firewall presents to the user if the site with the override is an HTTPS site. For details, see Configure an SSL/TLS Service Profile.
6. Select the **Mode** for prompting the user for the password:
   - **Transparent**—The firewall intercepts the browser traffic destined for site in a URL category you have set to override and impersonates the original destination URL, issuing an HTTP 302 to prompt for the password, which applies on a per-vsys level.

     The client browser will display certificate errors if it does not trust the certificate.
   - **Redirect**—The firewall intercepts HTTP or HTTPS traffic to a URL category set to override and redirects the request to a Layer 3 interface on the firewall using an HTTP 302 redirect in order to prompt for the override password. If you select this option, you must provide the **Address** (IP address or DNS hostname) to which to redirect the traffic.
7. Click **OK**.

STEP 2 | (Optional) Set a custom override period.

1. Edit the URL Filtering section.
2. To change the amount of time users can browse to a site in a category for which they have successfully entered the override password, enter a new value in the **URL Admin Override Timeout** field. By default, users can access sites within the category for 15 minutes without re-entering the password.
3. To change the amount of time users are blocked from accessing a site set to override after three failed attempts to enter the override password, enter a new value in the **URL Admin Lockout Timeout** field. By default, users are blocked for 30 minutes.
4. Click **OK**.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 3 | (Redirect mode only) Create a Layer 3 interface to which to redirect web requests to sites in a category configured for override.

1. Create a management profile to enable the interface to display the URL Filtering Continue and Override Page response page:
   1. Select **Network** > **Interface Mgmt** and click **Add**.
   2. Enter a **Name** for the profile, select **Response Pages**, and then click **OK**.
2. Create the Layer 3 interface. Be sure to attach the management profile you just created (on the **Advanced** > **Other Info** tab of the Ethernet Interface dialog).

STEP 4 | (Redirect mode only) To transparently redirect users without displaying certificate errors, install a certificate that matches the IP address of the interface to which you are redirecting web requests to a site in a URL category configured for override.You can either generate a self-signed certificate or import a certificate that is signed by an external CA.

To use a self-signed certificate, you must first create a root CA certificate and then use that CA to sign the certificate you will use for URL admin override as follows:

1. To create a root CA certificate, select **Device** > **Certificate Management** > **Certificates** > **Device Certificates** and then click **Generate**. Enter a **Certificate Name**, such as RootCA. Do not select a value in the **Signed By** field (this is what indicates that it is self-signed). Make sure you select the **Certificate Authority** check box and then click **Generate** the certificate.
2. To create the certificate to use for URL admin override, click **Generate**. Enter a **Certificate Name** and enter the DNS hostname or IP address of the interface as the **Common Name**. In the **Signed By** field, select the CA you created in the previous step. Add an IP address attribute and specify the IP address of the Layer 3 interface to which you will be redirecting web requests to URL categories that have the override action.
3. **Generate** the certificate.
4. To configure clients to trust the certificate, select the CA certificate on the **Device Certificates** tab and click **Export**. You must then import the certificate as a trusted root CA into all client browsers, either by manually configuring the browser or by adding the certificate to the trusted roots in an Active Directory Group Policy Object (GPO).

STEP 5 | Specify which URL categories require an override password to enable access.

1. Select **Objects** > **URL Filtering** and either select an existing URL Filtering profile or **Add** a new one.
2. On the **Categories** tab, set the Action to **override** for each category that requires a password.
3. Complete any remaining sections on the URL Filtering profile and then click **OK** to save the profile.

STEP 6 | Apply the URL Filtering profile to the Security policy rule(s) that allows access to the sites requiring password override for access.

1. Select **Policies** > **Security** and select the appropriate Security policy to modify it.
2. Select the **Actions** tab and in the **Profile Setting** section, click the drop-down for **URL Filtering** and select the profile.
3. Click **OK** to save.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 7 | Save the configuration.

Click **Commit**.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Prevent Credential Phishing

Phishing sites are sites that attackers disguise as legitimate websites with the intent to steal user information, especially the credentials that provide access to your network. When a phishing email enters a network, it takes just a single user to click the link and enter credentials to set a breach into motion. You can detect and prevent in-progress phishing attacks, thereby preventing credential theft, by controlling sites to which users can submit corporate credentials based on the site's URL category. This allows you to block users from submitting credentials to untrusted sites while allowing credential submissions to corporate and sanctioned sites.

Credential phishing prevention works by scanning username and password submissions to websites and comparing those submissions against valid corporate credentials. You can choose what websites you want to either allow or block corporate credential submissions to based on the URL category of the website. When the firewall detects a user attempting to submit credentials to a site in a category you have restricted, it either displays a block response page that prevents the user from submitting credentials or presents a continue page that warns users against submitting credentials to sites in certain URL categories but still allows them to continue with the submission. You can customize these block pages to educate users against reusing corporate credentials, even on legitimate, non-phishing sites.

To enable credential phishing prevention, you must configure both User-ID to detect when users submit valid corporate credentials to a site (as opposed to personal credentials) and URL Filtering to specify the URL categories in which you want to prevent users from entering their corporate credentials. The following topics describe the different methods you can use to detect credential submissions and provide instructions for configuring credential phishing protection.

- Methods to Check for Corporate Credential Submissions
- Configure Credential Detection with the Windows-based User-ID Agent
- Set Up Credential Phishing Prevention

Methods to Check for Corporate Credential Submissions

Before you set up credential phishing prevention, decide which method you want the firewall to use to check if valid corporate credentials have been submitted to a web page.

Method to Check Submitted Credentials	User-ID Configuration Requirements	How does this method detect corporate usernames and/or passwords that users submit to websites?
Group Mapping	Group Mapping configuration on the firewall	The firewall checks to determine if the username a user submits to a restricted site matches any valid corporate username. To do this, the firewall matches the submitted username to the list of usernames in its user-to-group mapping table to detect when users submit corporate usernames to sites in a restricted category. This method only checks for corporate username submissions based on LDAP group membership, which

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Method to Check Submitted Credentials	User-ID Configuration Requirements	How does this method detect corporate usernames and/or passwords that users submit to websites?
		makes it simple to configure, but more prone to false positives.
IP User Mapping	IP address-to- username mappings identified through User Mapping, GlobalProtect, or Authentication Policy and Authentication Portal.	The firewall checks to determine if the username a user submits to a restricted site maps to the IP address of the login username. To do this, the firewall matches the IP address of the login username and the username submitted to a web site to its IP address-to-user mapping table to detect when users submit their corporate usernames to sites in a restricted category. Because this method matches the IP address of the login username associated with the session against the IP address-to-username mapping table, it is an effective method for detecting corporate username submissions, but it does not detect corporate password submission. If you want to detect corporate username and password submission, you must use the Domain Credential Filter method.
Domain Credential Filter	Windows User-ID agent configured with the User-ID credential service add-on - AND - IP address-to- username mappings identified through User Mapping, GlobalProtect, or Authentication Policy and Authentication Portal.	The firewall checks to determine if the username and password a user submits match the same user's corporate username and password. To do this, the firewall must be able to match credential submissions to valid corporate usernames and passwords and verify that the username submitted maps to the IP address of the login username as follows: • **To detect corporate usernames and passwords**—The firewall retrieves a secure bit mask, called a *bloom filter*, from a Windows User-ID agent equipped with the User-ID credential service add-on. This add-on service scans your directory for usernames and password hashes and deconstructs them into a secure bit mask (the bloom filter) and delivers it to the Windows User-ID agent. The firewall retrieves the bloom filter from the Windows User-ID agent at regular intervals. Whenever it detects a user submitting credentials to a restricted category, it reconstructs the bloom filter and looks for a matching username and password hash. The firewall can only connect to one Windows User-ID agent running the User-ID credential service add-on.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Method to Check Submitted Credentials	User-ID Configuration Requirements	How does this method detect corporate usernames and/or passwords that users submit to websites?
		• **To verify that the credentials belong to the login username**—The firewall looks for a mapping between the IP address of the login username and the detected username in its IP address-to-username mapping table. To learn more how the domain credential method works and the requirements for enabling this type of detection, see Configure Credential Detection with the Windows-based User-ID Agent.

Configure Credential Detection with the Windows User-ID Agent

Domain Credential Filter detection enables the firewall to detect passwords submitted to web pages. This credential detection method requires the Windows User-ID agent and the User-ID credential service, an add-on to the User-ID agent, to be installed on a *read-only domain controller (RODC)*.

The Domain Credential Filter detection method is supported with the Windows User-ID agent only. You cannot use the PAN-OS integrated User-ID agent to configure this method of credential detection.

An RODC is a Microsoft Windows server that maintains a read-only copy of an Active Directory database that a domain controller hosts. When the domain controller is located at a corporate headquarters, for example, RODCs can be deployed in remote network locations to provide local authentication services. Installing the User-ID agent on an RODC can be useful for a few reasons: access to the domain controller directory is not required to enable credential detection and you can support credential detection for a limited or targeted set of users. Because the directory the RODC hosts is read-only, the directory contents remain secure on the domain controller.

Because you must install the Windows User-ID agent on the RODC for credential detection, as a best practice deploy a separate agent for this purpose. Do not use the User-ID agent installed on the RODC to map IP addresses to users.

After you install the User-ID agent on an RODC, the User-ID credential service runs in the background and scans the directory for the usernames and password hashes of group members that are listed in the RODC password replication policy (PRP)—you can define who you want to be on this list. The User-ID credential service then takes the collected usernames and password hashes and deconstructs the data into a type of bit mask called a *bloom filter*. Bloom filters are compact data structures that provide a secure method to check if an element (a username or a password hash) is a member of a set of elements (the sets of credentials you have approved for replication to the RODC). The User-ID credential service forwards the bloom filter to the Windows User-ID agent; the firewall retrieves the latest bloom filter from the User-ID agent at regular intervals and uses it to detect usernames and password hash submissions. Depending on your settings, the firewall then blocks, alerts, or allows on valid password submissions to web

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

pages, or displays a response page to users warning them of the dangers of phishing, but allowing them to continue with the submission.

Throughout this process, the User-ID agent does not store or expose any password hashes, nor does it forward password hashes to the firewall. Once the password hashes are deconstructed into a bloom filter, there is no way to recover them.

STEP 1 | Configure user mapping using the Windows User-ID agent.

> *To enable credential detection, you must install the Windows User-ID agent on an RODC. Refer to the Compatibility Matrix for a list of supported servers. Install a separate User-ID agent for this purpose.*

Important items to remember when setting up User-ID to enable Domain Credential Filter detection:

- Because the effectiveness of credential phishing detection is dependent on your RODC setup, make sure that you also review best practices and recommendations for RODC Administration.
- Download User-ID software updates:
  - User-ID Agent Windows installer—UaInstall-x.x.x-x.msi.
  - User-ID Agent Credential Service Windows installer—UaCredInstall64-x.x.x-x.msi.
- Install the User-ID agent and the User Agent Credential service on an RODC using an account that has privileges to read Active Directory via LDAP (the User-ID agent also requires this privilege).
  - The User-ID Agent Credential Service requires permission to log on with the local system account. For more information, refer to Create a Dedicated Service Account for the User-ID Agent.
  - The service account must be a member of the local administrator group on the RODC. For more information, refer to the following link.

STEP 2 | Enable the User-ID agent and the User Agent Credential service (which runs in the background to scan permitted credentials) to share information.

1. On the RODC server, launch the User-ID Agent.
2. Select **Setup** and edit the Setup section.
3. Select the **Credentials** tab. This tab only displays if you have already installed the User-ID Agent Credential Service.
4. Select **Import from User-ID Credential Agent**. This enables the User-ID agent to import the bloom filter that the User-ID credential agent creates to represent users and the corresponding password hashes.
5. Click **OK**, **Save** your settings, and **Commit**.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 3 | In the RODC directory, define the group of users for which you want to support credential submission detection.

- Confirm that the groups that should receive credential submission enforcement are added to the Allowed RODC Password Replication Group.
- Check that none of the groups in the Allowed RODC Password Replication Group are also in the Denied RODC Password Replication Group by default. Groups listed in both will not be subject to credential phishing enforcement.

STEP 4 | Continue to the next task.

Set up credential phishing prevention on the firewall.

Set Up Credential Phishing Prevention

After you have decided which of the methods to detect corporate credential submissions you want to use, take the following steps to enable the firewall to detect when users submit corporate credentials to web pages and either alert on this action, block the credential submission, or require users to acknowledge the dangers of phishing before continuing with the submission.



Before enabling credential phishing prevention, verify that the Primary Username that you configure on the firewall uses the sAMAccountName attribute. Credential phishing prevention does not support alternate attributes.

STEP 1 | If you have not done so already, enable User-ID.

Each of the methods to check for corporate credential submissions requires a different User-ID configuration:

- If you plan to use the group mapping method, which detects whether a user is submitting a valid corporate username, map users to groups.
- If you plan to use the IP user mapping method, which detects whether a user is submitting a valid corporate username and that username is the same as the login username, map IP addresses to users.
- If you plan to use the domain credential filter method, which detects whether a user is submitting a valid username and password and that those credentials belong to the logged-in user, configure credential detection with the Windows-based User-ID agent and map IP addresses to users.

STEP 2 | If you have not done so already, configure a best practice URL Filtering profile to ensure protection against URLs that have been observed hosting malware or exploitive content.

1. Select **Objects** > **Security Profiles** > **URL Filtering** and **Add** or modify a URL Filtering profile.
2. Block access to all known dangerous URL categories: malware, phishing, dynamic-dns, unknown, command-and-control, extremism, copyright-infringement, proxy-avoidance-and-anonymizers, newly-registered-domain, grayware, and parked.

STEP 3 | Add a Decryption policy rule to decrypt the traffic you want to monitor for user credential submissions.

https://t.me/learningnets

STEP 4 | Configure the URL Filtering profile to detect corporate credential submissions to websites that are in allowed URL categories.

The firewall does not check credential submissions for trusted sites, even if you enable the checks for the URL categories for these sites, to provide best performance. The trusted sites represent sites where Palo Alto Networks has not observed any malicious or phishing attacks. Updates for this trusted sites list are delivered through Application and Threat content updates. For a list of App-IDs that are exempt from credential detection, see Trusted App-IDs That Skip Credential Submission Detection *on* live.paloaltonetworks.com.

1. Select **User Credential Detection**.
2. Select one of the user credential detection methods to web pages from the **User Credential Detection** drop-down:

   Confirm that the format for the primary username is the same as the username format that the User-ID source provides.

   - **Use IP User Mapping—**Checks for valid corporate username submissions and verifies that the login username maps to the source IP address of the session. To do this, the firewall matches the submitted username and source IP address of the session against its IP-address-to-username mapping table. To use this method you can use any of the user mapping methods described in map IP addresses to users.
   - **Use Domain Credential Filter**—Checks for valid corporate usernames and password submissions and verifies that the username maps to the IP address of the logged in user. See Configure Credential Detection with the Windows-based User-ID Agent for instructions on how to set up User-ID to enable this method.
   - **Use Group Mapping**—Checks for valid username submissions based on the user-to-group mapping table populated when you configure the firewall to map users to groups.

     With group mapping, you can apply credential detection to any part of the directory, or for specific groups that have access to your most sensitive applications, such as IT.

     This method is prone to false positives in environments that do not have uniquely structured usernames. Because of this, you should only use this method to protect your high-value user accounts.
3. Set the **Valid Username Detected Log Severity** the firewall uses to log detection of corporate credential submissions. By default, the firewall logs these events as medium severity.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 5 | Block (or alert) on credential submissions to allowed sites.

1. Select **Categories**.
2. For each Category to which **Site Access** is allowed, select how you want to treat **User Credential Submissions**:
   - **alert**—Allow users to submit credentials to the website, but generate a URL Filtering log each time a user submits credentials to sites in this URL category.
   - **allow**—(default) Allow users to submit credentials to the website.
   - **block**—Block users from submitting credentials to the website. When a user tries to submit credentials, the firewall displays the Anti-Phishing Block Page, preventing the credential submission.
   - **continue**—Present the Anti-Phishing Continue Page response page to users when they attempt to submit credentials. Users must select Continue on the response page to continue with the submission.
3. Select **OK** to save the URL Filtering profile.

STEP 6 | Apply the URL Filtering profile with the credential detection settings to your Security policy rules.

1. Select **Policies** > **Security** and **Add** or modify a Security policy rule.
2. On the **Actions** tab, set the **Profile Type** to **Profiles**.
3. Select the new or updated **URL Filtering** profile to attach it to the Security policy rule.
4. Select **OK** to save the Security policy rule.

STEP 7 | **Commit** the configuration.

STEP 8 | Monitor credential submissions the firewall detects.



*Select **ACC** > **Hosts Visiting Malicious URLs** to see the number of users who have visited malware and phishing sites.*

Select **Monitor** > **Logs** > **URL Filtering**.

The new **Credential Detected** column indicates events where the firewall detected a HTTP post request that included a valid credential:

To display this column, hover over any column header and click the arrow to select the columns you'd like to display.

Log entry details also indicate credential submissions:

STEP 9 | Validate and troubleshoot credential submission detection.

- Use the following CLI command to view credential detection statistics:

```
> show user credential-filter statistics
```

The output for this command varies depending on the method configured for the firewall to detect credential submissions. For example, if the Domain Credential Filter method is

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

configured in any URL Filtering profile, a list of User-ID agents that have forwarded a bloom filter to the firewall is displayed, along with the number of credentials contained in the bloom filter.

- (Group Mapping method only) Use the following CLI command to view group mapping information, including the number of URL Filtering profiles with Group Mapping credential detection enabled and the usernames of group members that have attempted to submit credentials to a restricted site.

  ```
  > show user group-mapping statistics
  ```

- (Domain Credential Filter method only) Use the following CLI command to see all Windows-based User-ID agents that are sending mappings to the firewall:

  ```
  > show user user-id-agent state all
  ```

  The command output now displays bloom filter counts that include the number of bloom filter updates the firewall has received from each agent, if any bloom filter updates failed to process, and how many seconds have passed since the last bloom filter update.

- (Domain Credential Filter method only) The Windows-based User-ID agent displays log messages that reference BF (bloom filter) pushes to the firewall. In the User-ID agent interface, select **Monitoring** > **Logs**.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Safe Search Enforcement

Many search engines have a safe search setting that filters out adult images and videos in search query return traffic. You can enable the firewall to block search results if the end user is not using the strictest safe search settings, and you can also transparently enable safe search for your users. The firewall supports safe search enforcement for the following search providers: Google, Yahoo, Bing, Yandex, and YouTube. Consider that safe search is a best-effort setting and service providers do not guarantee that it works with every website, and search providers classify sites as safe or unsafe (not Palo Alto Networks).

To use this feature you must enable the **Safe Search Enforcement** option in a URL Filtering profile and attach it to a Security policy rule. The firewall then blocks any matching search query return traffic that is not using the strictest safe search settings. There are two methods to enforce safe search:

- Block Search Results when Strict Safe Search is not Enabled—When an end user attempts to perform a search without first enabling the strictest safe search settings, the firewall blocks the search query results and displays the URL Filtering Safe Search Block Page. By default, this page will provide a URL to the search provider settings for configuring safe search.
- Transparently Enable Safe Search for Users—When an end user attempts to perform a search without first enabling the strict safe search settings, the firewall blocks the search results with an HTTP 503 status code and redirects the search query to a URL that includes the safe search parameters. You enable this functionality by importing a new URL Filtering Safe Search Block Page containing the JavaScript for rewriting the search URL to include the strict safe search parameters. In this configuration, users will not see the block page, but will instead be automatically redirected to a search query that enforces the strictest safe search options. This safe search enforcement method is supported for Google, Yahoo, and Bing searches.

As safe search settings differ by search provider, get started by reviewing the different safe search implementations. Then, enforce safe search in one of two ways: block search results when safe search is disabled, or transparently enable safe search for your users.

- Safe Search Settings for Search Providers
- Block Search Results when Strict Safe Search is not Enabled
- Transparently Enable Safe Search for Users

Safe Search Settings for Search Providers

Safe search settings differ for each search provider—review the following settings to learn more.

Search Provider	Safe Search Setting Description
Google/YouTube	Offers safe search on individual computers or network-wide through Google's safe search virtual IP address: **Safe Search Enforcement for Google Searches on Individual Computers** In the Google Search Settings, the **Filter explicit results** setting enables safe search functionality. When

©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Search Provider	Safe Search Setting Description
	enabled, the setting is stored in a browser cookie as FF= and passed to the server each time the user performs a Google search. Appending **safe=active** to a Google search query URL also enables the strictest safe search settings. **Safe Search Enforcement for Google and YouTube Searches using a Virtual IP Address** Google provides servers thatLock SafeSearch (forcesafesearch.google.com) settings in every Google and YouTube search. By adding a DNS entry for **www.google.com** and **www.youtube.com** (and other relevant Google and YouTube country subdomains) that includes a CNAME record pointing to **forcesafesearch.google.com** to your DNS server configuration, you can ensure that all users on your network are using strict safe search settings every time they perform a Google or YouTube search. Keep in mind, however, that this solution is not compatible with Safe Search Enforcement on the firewall. Therefore, if you are using this option to force safe search on Google, the best practice is to block access to other search engines on the firewall by creating custom URL categories and adding them to the block list in the URL Filtering profile.

https://t.me/learningnets

Search Provider	Safe Search Setting Description
	 • *PAN-OS supports safe search enforcement for YouTube through HTTP header insertion. HTTP header insertion is not currently supported for HTTP/2. To enforce safe search for YouTube, App-ID and HTTP/2 Inspection downgrade HTTP/2 connections to HTTP/1.1 using the **Strip ALPN** feature in the appropriate decryption profile.* • *If you plan to use the Google Lock SafeSearch solution, consider configuring DNS Proxy (**Network** > **DNS Proxy**) and setting the inheritance source as the Layer 3 interface on which the firewall receives DNS settings from service provider via DHCP. You would configure the DNS proxy with **Static Entries** for www.google.com and www.youtube.com, using the local IP address for the forcesafesearch.google.com server.*
Yahoo	Offers safe search on individual computers only. The Yahoo Search Preferences includes three SafeSearch settings: **Strict**, **Moderate**, or **Off**. When enabled, the setting is stored in a browser cookie as `vm=` and passed to the server each time the user performs a Yahoo search. Appending `vm=r` to a Yahoo search query URL also enables the strictest safe search settings.  *When performing a search on Yahoo Japan (**`yahoo.co.jp`**) while logged into a Yahoo account, end users must also enable the **SafeSearch Lock** option.*
Bing	Offers safe search on individual computers or through their Bing in the Classroom program. The Bing Settings include three SafeSearch settings: **Strict**, **Moderate**, or **Off**. When enabled, the setting is stored in a browser cookie as `adtl=` and passed to the server each time the user performs a Bing search. Appending **`adlt=strict`** to a Bing search query URL also enables the strictest safe search settings.

https://t.me/learningnets

Search Provider	Safe Search Setting Description
	The Bing SSL search engine does not enforce the safe search URL parameters and you should therefore consider blocking Bing over SSL for full safe search enforcement.

Block Search Results when Strict Safe Search is not Enabled

By default, if you enable safe search enforcement and a user searches without the strictest settings, the firewall blocks the search query results and displays the URL Filtering Safe Search Block Page in the browser. The response page provides a link to the search settings page for the corresponding search provider so that end users can adjust their settings. You can customize the URL Filtering response page to meet your organization's specific needs. If you plan to use this default method to enforce safe search, you should communicate the policy to your end users before implementing it. Alternatively, you can transparently enable safe search for users so that they don't have to configure the settings manually.

Safe Search Settings for Search Providers summarizes how each search provider implements safe search.

STEP 1 | Enable Safe Search Enforcement in the URL Filtering profile.

1. Select **Objects** > **Security Profiles** > **URL Filtering**.
2. Select an existing profile to modify, or clone the default profile to create a new profile.
3. On the **Settings** tab, select the **Safe Search Enforcement** option.
4. (Optional) Restrict users to specific search engines:
   1. On the **Categories** tab, set the **search-engines** category to **block**.
   2. For each search engine that you want end users to be able to access, enter the web address in the **Allow List** text box. For example, to allow users access to Google and Bing searches only, you would enter the following:

      `www.google.com`

      `www.bing.com`
5. Configure other settings as necessary to:
   - Define site access for each URL category.
   - Define block and allow lists to specify websites that should always be blocked or allowed, regardless of URL category.
6. Click **OK** to save the profile.

STEP 2 | Add the URL Filtering profile to the Security policy rule that allows traffic from clients in the trust zone to the Internet.

1. Select **Policies** > **Security** and select a rule to which to apply the URL Filtering profile that you just enabled for Safe Search Enforcement.
2. On the **Actions** tab, select the **URL Filtering** profile.
3. Click **OK** to save the Security policy rule.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 3 | Enable SSL Forward Proxy decryption.

Because most search engines encrypt their search results, you must enable SSL Forward Proxy decryption so that the firewall can inspect the search traffic and detect the safe search settings.

1. Add a custom URL category for the search sites:
   1. Select **Objects** > **Custom Objects** > **URL Category** and **Add** a custom category.
   2. Enter a **Name** for the category, such as SearchEngineDecryption.
   3. **Add** the following to the Sites list:

      `www.bing.*`

      `www.google.*`

      `search.yahoo.*`
   4. Click **OK** to save the custom URL category object.
2. Follow the steps to configure SSL Forward Proxy.
3. On the **Service/URL Category** tab in the decryption policy rule, **Add** the custom URL category you just created and then click **OK**.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 4 | (Recommended) Block Bing search traffic running over SSL.

Because the Bing SSL search engine does not adhere to the safe search settings, for full safe search enforcement, you must deny all Bing sessions that run over SSL.

1. Add a custom URL category for Bing:
   1. Select **Objects** > **Custom Objects** > **URL Category** and **Add** a custom category.
   2. Enter a **Name** for the category, such as EnableBingSafeSearch.
   3. **Add** the following to the Sites list:

      `www.bing.com/images/*`

      `www.bing.com/videos/*`
   4. Click **OK** to save the custom URL category object.
2. Create another URL Filtering profile to block the custom category you just created:
   1. Select **Objects** > **Security Profiles** > **URL Filtering**.
   2. **Add** a new profile and give it a descriptive **Name**.
   3. Locate the custom category in the Category list and set it to **block**.
   4. Click **OK** to save the URL Filtering profile.
3. Add a Security policy rule to block Bing SSL traffic:
   1. Select **Policies** > **Security** and **Add** a policy rule that allows traffic from your trust zone to the Internet.
   2. On the **Actions** tab, attach the URL Filtering profile you just created to block the custom Bing category.
   3. On the **Service/URL Category** tab **Add** a **New Service** and give it a descriptive **Name**, such as bingssl.
   4. Select **TCP** as the **Protocol** and set the **Destination Port** to 443.
   5. Click **OK** to save the rule.
   6. Use the **Move** options to ensure that this rule is below the rule that has the URL Filtering profile with Safe Search Enforcement enabled.

STEP 5 | Save the configuration.

Click **Commit**.

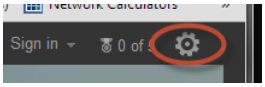
 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

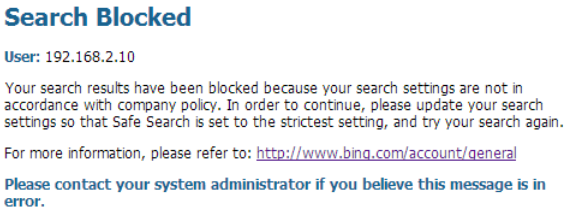
STEP 6 | Verify the Safe Search Enforcement configuration.

This verification step only works if you are using block pages to enforce safe search. If you are using transparent safe search enforcement, the firewall block page will invoke a URL rewrite with the safe search parameters in the query string.

1. From a computer that is behind the firewall, disable the strict search settings for one of the supported search providers. For example, on bing.com, click the **Preferences** icon on the Bing menu bar.



2. Set the **SafeSearch** option to **Moderate** or **Off** and click **Save**.
3. Perform a Bing search and verify that the URL Filtering Safe Search Block page displays instead of the search results:



4. Use the link in the block page to go to the search settings for the search provider and set the safe search setting back to the strictest setting (**Strict** in the case of Bing) and then click **Save**.
5. Perform a search again from Bing and verify that the filtered search results display instead of the block page.

Transparently Enable Safe Search for Users

If you want to enforce filtering of search query results with the strictest safe search filters, but you don't want your end users to have to manually configure the settings, you can enable transparent safe search enforcement as follows. This functionality is supported on Google, Yahoo, and Bing search engines only and requires Content Release version 475 or later.

STEP 1 | Make sure the firewall is running Content Release version 475 or later.

1. Select **Device** > **Dynamic Updates**.
2. Check the **Applications and Threats** section to determine what update is currently running.
3. If the firewall is not running the required update or later, click **Check Now** to retrieve a list of available updates.
4. Locate the required update and click **Download**.
5. After the download completes, click **Install**.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 2 | Enable Safe Search Enforcement in the URL Filtering profile.

1. Select **Objects** > **Security Profiles** > **URL Filtering**.
2. Select an existing profile to modify, or clone the default profile to create a new one.
3. On the **Settings** tab, select the **Safe Search Enforcement** check box to enable it.
4. (Optional) Allow access to specific search engines only:
   1. On the **Categories** tab, set the **search-engines** category to **block**.
   2. For each search engine that you want end users to be able to access, enter the web address in the **Allow List** text box. For example, to allow users access to Google and Bing searches only, you would enter the following:

      `www.google.com`

      `www.bing.com`
5. Configure other settings as necessary to:
   - Define site access for each URL category.
   - Define block and allow lists to specify websites that should always be blocked or allowed, regardless of URL category.
6. Click **OK** to save the profile.

STEP 3 | Add the URL Filtering profile to the Security policy rule that allows traffic from clients in the trust zone to the Internet.

1. Select **Policies** > **Security** and select a rule to which to apply the URL Filtering profile that you just enabled for Safe Search Enforcement.
2. On the **Actions** tab, select the **URL Filtering** profile.
3. Click **OK** to save the Security policy rule.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 4 | (Recommended) Block Bing search traffic running over SSL.

Because the Bing SSL search engine does not adhere to the safe search settings, for full safe search enforcement, you must deny all Bing sessions that run over SSL.

1. Add a custom URL category for Bing:
   1. Select **Objects** > **Custom Objects** > **URL Category** and **Add** a custom category.
   2. Enter a **Name** for the category, such as EnableBingSafeSearch.
   3. **Add** the following to the Sites list:

      www.bing.com/images/*

      www.bing.com/videos/*
   4. Click **OK** to save the custom URL category object.
2. Create another URL Filtering profile to block the custom category you just created:
   1. Select **Objects** > **Security Profiles** > **URL Filtering**.
   2. **Add** a new profile and give it a descriptive **Name**.
   3. Locate the custom category you just created in the Category list and set it to **block**.
   4. Click **OK** to save the URL Filtering profile.
3. **Add** a Security policy rule to block Bing SSL traffic:
   1. Select **Policies** > **Security** and **Add** a policy rule that allows traffic from your trust zone to the Internet.
   2. On the **Actions** tab, attach the URL Filtering profile you just created to block the custom Bing category.
   3. On the **Service/URL Category** tab **Add** a **New Service** and give it a descriptive **Name**, such as bingssl.
   4. Select **TCP** as the **Protocol**, set the **Destination Port** to **443**.
   5. Click **OK** to save the rule.
   6. Use the **Move** options to ensure that this rule is below the rule that has the URL Filtering profile with safe search enforcement enabled.

STEP 5 | Edit the URL Filtering Safe Search Block Page, replacing the existing code with the JavaScript for rewriting search query URLs to enforce safe search transparently.

1. Select **Device** > **Response Pages** > **URL Filtering Safe Search Block Page**.
2. Select **Predefined** and then click **Export** to save the file locally.
3. Use an HTML editor and replace all of the existing block page text with the following text and then save the file.

```
<html>
  <head>
    <title>Search Blocked</title>
    <meta http-equiv="Content-Type" content="text/html;
 charset=utf-8">
    <meta http-equiv="pragma" content="no-cache">
    <meta name="viewport" content="initial-scale=1.0">
    <style>
      #content {
```

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

```
      border:3px solid#aaa;
      background-color:#fff;
      margin:1.5em;
      padding:1.5em;
      font-family:Tahoma,Helvetica,Arial,sans-serif;
      font-size:1em;
      }
      h1 {
      font-size:1.3em;
      font-weight:bold;
      color:#196390;
      }
      b {
      font-weight:normal;
      color:#196390;
      }
    </style>
  </head>
  <body bgcolor="#e7e8e9">
    <div id="content">
      <h1>Search Blocked</h1>
      <p>
        <b>User:</b>
        <user/>
      </p>
      <p>Your search results have been blocked because your
search settings are not in accordance with company policy.
In order to continue, please update your search settings so
that Safe Search is set to the strictest setting. If you are
currently logged into your account, please also lock Safe
Search and try your search again.</p>
      <p>
        For more information, please refer to:
        <a href="<ssurl/>">
          <ssurl/>
        </a>
      </p>
      <p id="java_off"> Please enable JavaScript in your
browser.<br></p>
      <p><b>Please contact your system administrator if you
believe this message is in error.</b></p>
    </div>
  </body>
  <script>
    // Grab the URL that's in the browser.
    var s_u = location.href;
    //bing
    // Matches the forward slashes in the beginning, anything,
then ".bing." then anything followed by a non greedy slash.
Hopefully the first forward slash.
    var b_a = /^.*\/\/(.+\.bing\..+?)\//.exec(s_u);
    if (b_a) {
       s_u = s_u + "&adlt=strict";
        window.location.replace(s_u);
        document.getElementById("java_off").innerHTML = 'You
are being redirected to a safer search!';
```

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

```
    }
    //google
    // Matches the forward slashes in the beginning, anything,
 then ".google." then anything followed by a non greedy slash.
 Hopefully the first forward slash.
    var g_a = /^.*\/\/(.+\.google\..+?)\//.exec(s_u);
    if (g_a) {
        s_u = s_u.replace(/&safe=off/ig,"");
        s_u = s_u + "&safe=active";
        window.location.replace(s_u);
        document.getElementById("java_off").innerHTML = 'You
 are being redirected to a safer search!';    }
    //yahoo
    // Matches the forward slashes in the beginning, anything,
 then ".yahoo."" then anything followed by a non greedy slash.
 Hopefully the first forward slash.
    var y_a = /^.*\/\/(.+\.yahoo\..+?)\//.exec(s_u);
    if (y_a) {
        s_u = s_u.replace(/&vm=p/ig,"");
        s_u = s_u + "&vm=r";
        window.location.replace(s_u);
        document.getElementById("java_off").innerHTML = 'You
 are being redirected to a safer search!';
    }
    document.getElementById("java_off").innerHTML = ' ';
  </script>
</html>
```

STEP 6 | Import the edited URL Filtering Safe Search Block page onto the firewall.

1. To import the edited block page, select **Device** > **Response Pages** > **URL Filtering Safe Search Block Page**.
2. Click **Import** and then enter the path and filename in the **Import File** field or **Browse** to locate the file.
3. (Optional) Select the virtual system on which this login page will be used from the **Destination** drop-down or select **shared** to make it available to all virtual systems.
4. Click **OK** to import the file.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 7 | Enable SSL Forward Proxy decryption.

Because most search engines encrypt their search results, you must enable SSL forward proxy decryption so that the firewall can inspect the search traffic and detect the safe search settings.

1. Add a custom URL category for the search sites:
    1. Select **Objects** > **Custom Objects** > **URL Category** and **Add** a custom category.
    2. Enter a **Name** for the category, such as SearchEngineDecryption.
    3. **Add** the following to the Sites list:

        `www.bing.*`

        `www.google.*`

        `search.yahoo.*`
    4. Click **OK** to save the custom URL category object.
2. Follow the steps to configure SSL Forward Proxy.
3. On the **Service/URL Category** tab in the Decryption policy rule, **Add** the custom URL category you just created and then click **OK**.

STEP 8 | Save the configuration.

Click **Commit**.

 ©2022 Palo Alto Networks, Inc.

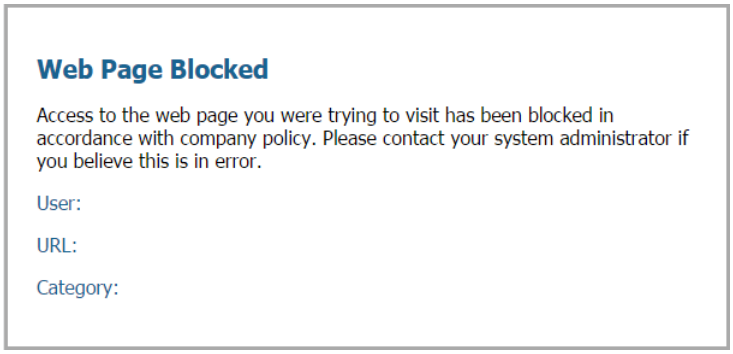
https://t.me/learningnets

URL Filtering Response Pages

The firewall provides three predefined response pages that display by default when a user attempts to browse to a site in a category that is configured with one of the block actions in the URL Filtering profile (block, continue, or override) or when Container Pages is enabled:

- **URL Filtering and Category Match Block Page**

  Access blocked by a URL Filtering profile or because the URL category is blocked by a Security policy rule.



- **URL Filtering Continue and Override Page**

  Page with initial block policy that allows users to bypass the block by clicking **Continue**. With URL Admin Override enabled (Allow Password Access to Certain Sites), after clicking **Continue**, the user must supply a password to override the policy that blocks the URL.



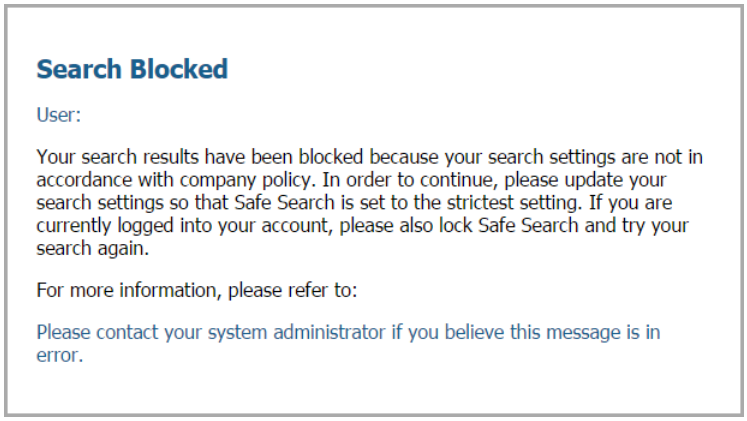
- **URL Filtering Safe Search Block Page**

  Access blocked by a Security policy rule with a URL Filtering profile that has the Safe Search Enforcement option enabled (see Safe Search Enforcement). The user will see this page if a

 ©2022 Palo Alto Networks, Inc.

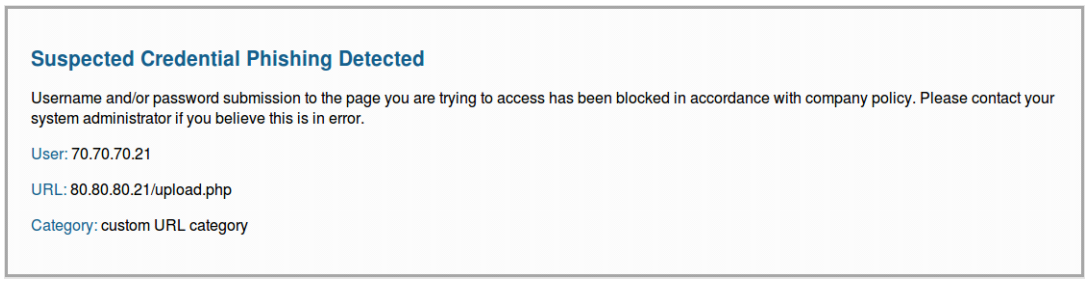
https://t.me/learningnets

search is performed using Google, Bing, Yahoo, or Yandex and their browser or search engine account setting for Safe Search is not set to strict.



- **Anti Phishing Block Page**

  This page displays to users when they attempt to enter corporate credentials (usernames or passwords) on a web page in a category for which credential submissions are blocked. The user can continue to access the site but remains unable to submit valid corporate credentials to any associated web forms. To control the sites to which users can submit corporate credentials, the firewall must be configured with User-ID and enabled to prevent credential phishing based on URL category.



- **Anti Phishing Continue Page**

  This page warns users against submitting credentials (usernames and passwords) to a web site. Warning users against submitting credentials can help to discourage them from reusing corporate credentials and to educate them about possible phishing attempts. They must select Continue to proceed to credentials on the site. To control the sites to which users can submit corporate credentials, the firewall must be configured with User-ID and enabled to prevent credential phishing based on URL category.



https://t.me/learningnets

You can either use the predefined pages, or you can customize the URL Filtering response pages to communicate your specific acceptable use policies and/or corporate branding. In addition, you can use the URL Filtering response page variables for substitution at the time of the block event or add one of the supported response page references to external images, sounds, or style sheets.

The browser will not display response pages if you have enabled your firewall to inspect SSL/TLS handshakes.

Table 2: URL Filtering Response Page Variables

Variable	Usage
<user/>	The firewall replaces the variable with the username (if available via User-ID) or IP address of the user when displaying the response page.
<url/>	The firewall replaces the variable with the requested URL when displaying the response page.
<category/>	The firewall replaces the variable with the URL filtering category of the blocked request.
<pan_form/>	HTML code for displaying the **Continue** button on the URL Filtering Continue and Override page.

You can also add code that triggers the firewall to display different messages depending on what URL category the user is attempting to access. For example, the following code snippet from a response page specifies to display Message 1 if the URL category is games, Message 2 if the category is travel, or Message 3 if the category is kids:

```
var cat = "<category/>";
switch(cat)
{
 case 'games':
 document.getElementById("warningText").innerHTML = "Message 1";
 break;
 case 'travel':
 document.getElementById("warningText").innerHTML = "Message 2";
 break;
 case 'kids':
 document.getElementById("warningText").innerHTML = "Message 3";
 break;
}
```

Only a single HTML page can be loaded into each virtual system for each type of block page. However, other resources such as images, sounds, and cascading style sheets (CSS files) can be loaded from other servers at the time the response page is displayed in the browser. All references must include a fully qualified URL.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Table 3: Response Page References

Reference Type	Example HTML Code
Image	`<img src="http://virginiadot.org/images/Stop-Sign-gif.gif">`
Sound	`<embed src="http://simplythebest.net/sounds/WAV/WAV_files/ movie_WAV_files/ do_not_go.wav" volume="100" hidden="true" autostart="true">`
Style Sheet	`<link href="http://example.com/style.css" rel="stylesheet" type="text/css" />`
Hyperlink	`<a href="http://en.wikipedia.org/wiki/Acceptable_use_policy">View Corporate Policy</a>`

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Customize the URL Filtering Response Pages

The firewall provides predefined URL Filtering response pages that display by default when:

- A user attempts to browse to a site in a category with restricted access.
- A user submits valid corporate credentials to a site for which credential detection is enabled (Prevent Credential Phishing based on URL category).
- Log Only the Page a User Visits blocks a search attempt.

However, you can create your own custom response pages with your corporate branding, acceptable use policies, and links to your internal resources.

Custom response pages larger than the maximum supported size are not decrypted or displayed to users. In PAN-OS 8.1.2 and earlier PAN-OS 8.1 releases, custom response pages on a decrypted site cannot exceed 8,191 bytes; the maximum size is increased to 17,999 bytes in PAN-OS 8.1.3 and later releases.

STEP 1 | Export the default response page(s).

1. Select **Device** > **Response Pages**.
2. Select the link for the URL Filtering response page you want to modify.
3. Click the response page (predefined or shared) and then click the **Export** link and save the file to your desktop.

STEP 2 | Edit the exported page.

1. Using the HTML text editor of your choice, edit the page:
   - If you want the response page to display custom information about the specific user, URL, or category that was blocked, add one or more of the supported response page variables.
   - If you want to include custom images (such as your corporate logo), a sound, or style sheet, or link to another URL, for example to a document detailing your acceptable web use policy, include one or more of the supported response page references.
2. Save the edited page with a new filename. Make sure that the page retains its UTF-8 encoding. For example, in Notepad you would select **UTF-8** from the **Encoding** drop-down in the Save As dialog.

STEP 3 | Import the customized response page.

1. Select **Device** > **Response Pages**.
2. Select the link that corresponds to the URL Filtering response page you edited.
3. Click **Import** and then enter the path and filename in the **Import File** field or **Browse** to locate the file.
4. (Optional) Select the virtual system on which this login page will be used from the **Destination** drop-down or select **shared** to make it available to all virtual systems.
5. Click **OK** to import the file.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 4 | Save the new response page(s).

Commit the changes.

STEP 5 | Verify that the new response page displays.

From a browser, go to the URL that will trigger the response page. For example, to see a modified URL Filtering and Category Match response page, browse to URL that your URL Filtering policy is set to block.

The firewall uses the following ports to display the URL Filtering response pages:

- **HTTP**—6080
- **Default TLS with firewall certificate**—6081
- **Custom SSL/TLS profile**—6082

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

HTTP Header Logging

URL filtering provides visibility and control over web traffic on your network. For improved visibility into web content, you can configure the URL Filtering profile to log HTTP header attributes included in a web request. When a client requests a web page, the HTTP header includes the user agent, referer, and x-forwarded-for fields as attribute-value pairs and forwards them to the web server. When enabled for logging HTTP headers, the firewall logs the following attribute-value pairs in the URL Filtering logs.



You can also use HTTP headers to manage access to SaaS applications. You don't need a URL Filtering license to do this, but you must use a URL Filtering profile to turn this feature on.

Attribute	Description
User-Agent	The web browser that the user used to access the URL, for example, Internet Explorer. This information is sent in the HTTP request to the server.
Referer	The URL of the web page that linked the user to another web page; it is the source that redirected (referred) the user to the web page that is being requested.
X-Forwarded-For (XFF)	The option in the HTTP request header field that preserves the IP address of the user who requested the web page. If you have a proxy server on your network, the XFF allows you to identify the IP address of the user who requested the content, instead of only recording the proxy server's IP address as source IP address that requested the web page.
Headers Inserted	The type of header and the text of the header that the firewall inserts.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Request to Change the Category for a URL

If you think that a URL is not categorized accurately, you can request for us to categorize it differently. Submit a change request directly in the firewall, or use Test A Site. A change request triggers PAN-DB—the URL Filtering cloud—to do an immediate analysis of the URL for which you're suggesting a category change. If PAN-DB validates that the new category suggestion is accurate, the change request is approved. If PAN-DB does not find the new category suggestion to be accurate, the change request is then reviewed by human editors from the Palo Alto Networks threat research and data science teams.

After you've submitted a change request, you'll receive an email from us confirming that we've received your request. When we've completed our investigation, you'll receive a second email confirming the results.

You cannot request to change the risk category a URL receives (**high risk**, **medium risk**, or **low risk**), or for URLs categorized as **insufficient content** or **newly-registered domains**.

- Make a Change Request Online
- Make a Bulk Change Request
- Make a Change Request From the Firewall

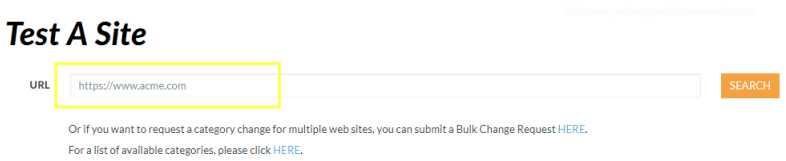
Make a Change Request Online

Visit Palo Alto Networks URL Filtering Test A Site to make a change request online.

STEP 1 | Go to Test A Site.

You do not need to log in to submit a change request, though you will need to provide your email as part of the change request form. If you decide not to log in, you'll need to take a CAPTCHA test to confirm that you're a human being (log in to avoid the CAPTCHA test).

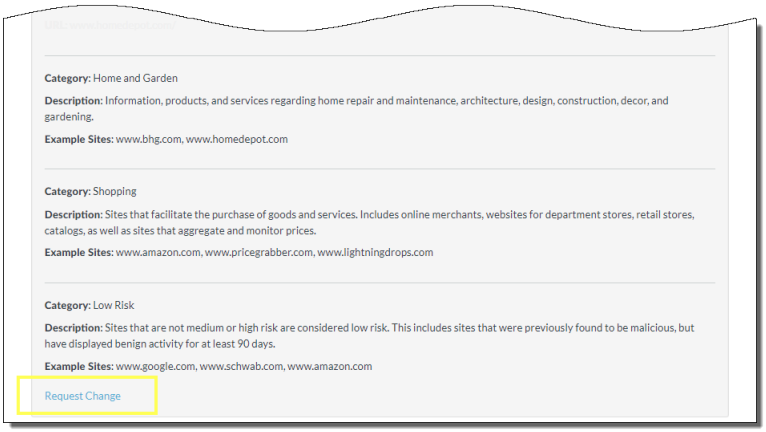
STEP 2 | Enter a URL to check its categories:



 ©2022 Palo Alto Networks, Inc.

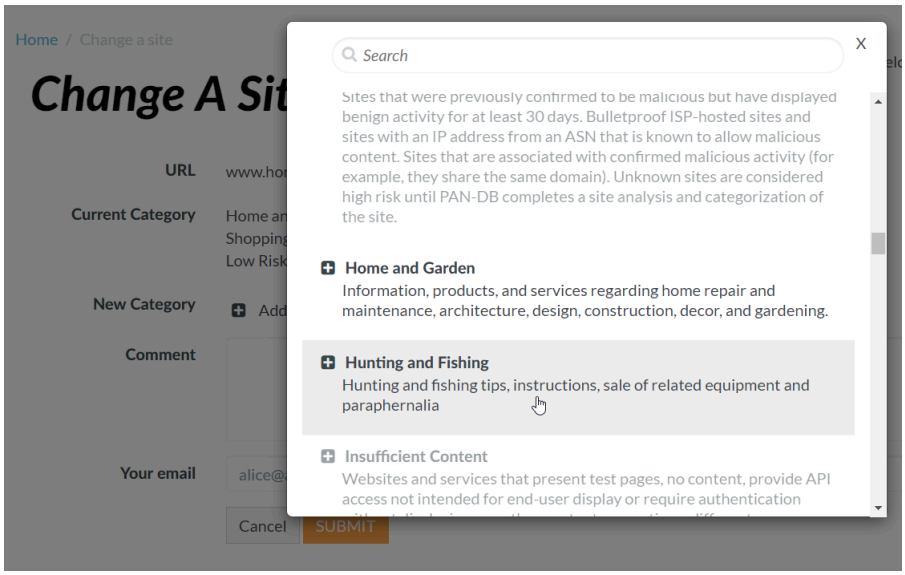
https://t.me/learningnets

STEP 3 | Review the URL categories, and if you don't think that they're accurate, select **Request Change**.



STEP 4 | Continue to populate and submit the change request form.

Include at least one (and up to two) new category suggestions, and leave an (optional) comment to tell us more about your suggestion.



Make a Bulk Change Request

You can also use Test A Site to make a bulk change request, where you want to submit change requests for multiple URLs at a single time.

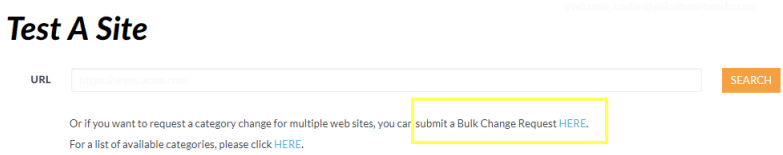
STEP 1 | Go to Test A Site.

You don't need to log in to make a change request; however, you'll need to provide your email as part of completing the change request form. If you decide not to log in, you'll need to take a CAPTCHA test to confirm that you're a human being (log in to avoid the CAPTCHA test).

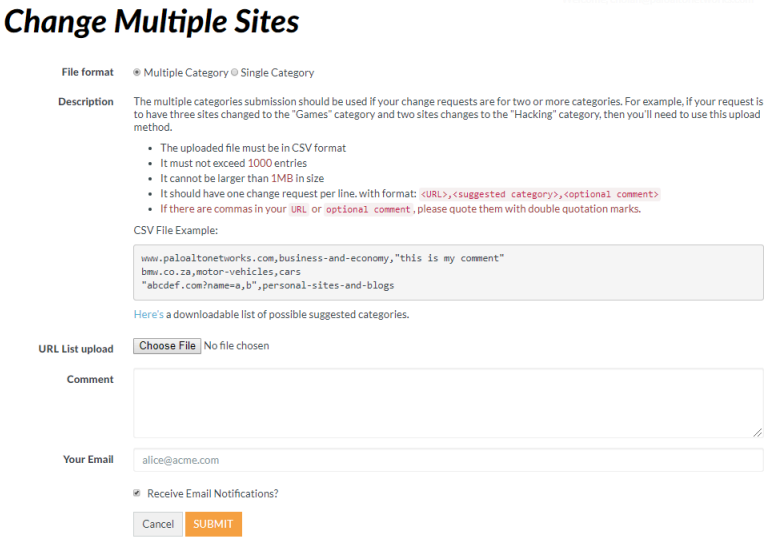
 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 2 | Choose the option to submit a bulk change request:

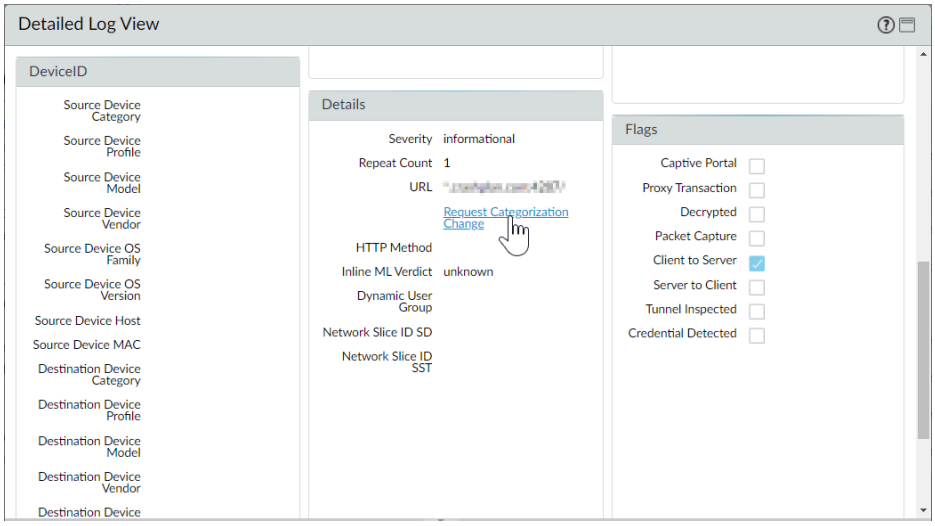


STEP 3 | Complete and submit the bulk change request form.



Make a Change Request from the Firewall

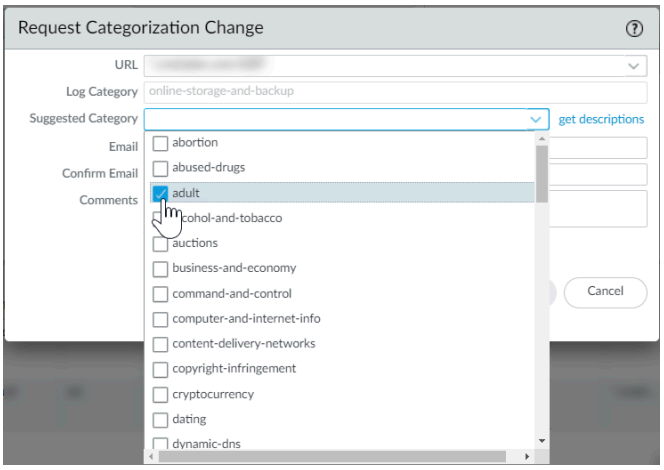
You can also submit a URL category change request directly from the firewall. In the URL Filtering logs, the details for each log entry include an option to **Request Categorization Change** (**Monitor** > **Logs** > **URL Filtering**).



From here you can complete the request form, and submit it.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets



 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Troubleshoot URL Filtering

The following topics provide troubleshooting guidelines for diagnosing and resolving common URL filtering problems.

- Problems Activating Advanced URL Filtering
- PAN-DB Cloud Connectivity Issues
- URLs Classified as Not-Resolved
- Incorrect Categorization

Problems Activating Advanced URL Filtering

Use the following workflow to troubleshoot Advanced URL Filtering activation issues.

STEP 1 | Access the PAN-OS CLI.

STEP 2 | Verify whether Advanced URL Filtering has been activated by running the following command:

```
show system setting url-database
```

If the response is `paloaltonetworks`, PAN-DB, the Palo Alto Networks URL filtering database, is the active vendor.

STEP 3 | Verify that the firewall has a valid Advanced URL Filtering license by running the following command:

```
request license info
```

You should see the license entry `Feature: Advanced URL Filtering`. If the license is not installed, you will need to obtain and install a license. See Configure URL Filtering.

STEP 4 | Check the PAN-DB cloud connection status.

PAN-DB Cloud Connectivity Issues

To help ensure connectivity to PAN-DB cloud, create a dedicated Security policy rule to allow all Palo Alto Management Service traffic. This will avoid management traffic from being classified as `not-resolved` and prevent the traffic from being blocked when routed through the dataplane.

To check connectivity between the firewall and the PAN-DB cloud:

```
show url-cloud status
```

If the cloud is accessible, the expected response is similar to the following:

```
show url-cloud status
PAN-DB URL Filtering
License :                          valid
```


©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

```
Current cloud server :
 serverlist.urlcloud.paloaltonetworks.com
Cloud connection :                  connected
Cloud mode :                        public
URL database version - device :     20200624.20296
URL database version - cloud :      20200624.20296  ( last update time
 2020/06/24 12:39:19 )
URL database status :               good
URL protocol version - device :     pan/2.0.0
URL protocol version - cloud :      pan/2.0.0
Protocol compatibility status :     compatible
```

If the cloud is not accessible, the expected response is similar to the following:

```
show url-cloud status
PAN-DB URL Filtering
License :                           valid
Cloud connection :                  not connected
URL database version - device :     0000.00.00.000
URL protocol version - device :     pan/0.0.2
```

Use the following checklist to identify and resolve connectivity issues:

- Does the PAN-DB URL Filtering license field shows as invalid? Obtain and install a valid PAN-DB license.
- Does the URL protocol version show as not compatible? Upgrade PAN-OS to the latest version.
- Can you ping the PAN-DB cloud server from the firewall? Run the following command to check:

  ```
  ping source <ip-address> host
   serverlist.urlcloud.paloaltonetworks.com <
  ```

  For example, if your management interface IP address is 10.1.1.5, run the following command:

  ```
  ping source 10.1.1.5 host serverlist.urlcloud.paloaltonetworks.com
  ```

- Is the firewall in an HA configuration? Verify that the HA state of the firewalls is in the active, active-primary, or active-secondary state. Access to the PAN-DB cloud will be blocked if the firewall is in a different state. Run the following command on each firewall in the pair to see the state:

  ```
  show high-availability state
  ```

If you still have problems with connectivity between the firewall and the PAN-DB cloud, contact Palo Alto Networks support.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

URLs Classified as Not-Resolved

URLs are classified as Not-resolved when your connection to the PAN-DB URL filtering cloud service is disrupted, resulting in failed URL lookups. The cloud connection status and URL classification does not apply to expired subscription licenses or unlicensed users.

Use the following workflow to troubleshoot why some or all of the URLs being identified by PAN-DB are classified as Not-resolved:

STEP 1 | Check the PAN-DB cloud connection by running the following command:

```
show url-cloud status
```

The Cloud connection: field should show `connected`. If you see anything other than `connected`, any URL that do not exist in the management plane cache will be categorized as `not-resolved`. To resolve this issue, see PAN-DB Cloud Connectivity Issues.

STEP 2 | If the cloud connection status shows `connected`, check the current utilization of the firewall. If firewall utilization is spiking, URL requests may be dropped (may not reach the management plane), and will be categorized as `not-resolved`.

To view system resources, run the following command and view the `%CPU` and `%MEM` columns:

```
show system resources
```

You can also view system resources on the System Resources widget on the **Dashboard** in the web interface.

STEP 3 | If the problem persists, contact Palo Alto Networks support.

Incorrect Categorization

Sometimes you may come across a URL that you believe is categorized incorrectly. Use the following workflow to determine the URL categorization for a site and request a category change, if appropriate.

STEP 1 | Verify the category in the dataplane by running the following command:

```
show running url <URL>
```

For example, to view the category for the Palo Alto Networks website, run the following command:

```
show running url paloaltonetworks.com
```

If the URL stored in the dataplane cache has the correct category (computer-and-internet-info in this example), then the categorization is correct and no further action is required. If the category is not correct, continue to the next step.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 2 | Verify if the category in the management plane by running the command:

```
test url-info-host <URL>
```

For example:

```
test url-info-host paloaltonetworks.com
```

If the URL stored in the management plane cache has the correct category, remove the URL from the dataplane cache by running the following command:

```
clear url-cache url <URL>
```

The next time the firewall requests the category for this URL, the request will be forwarded to the management plane. This will resolve the issue and no further action is required. If this does not solve the issue, go to the next step to check the URL category on the cloud systems.

STEP 3 | Verify the category in the cloud by running the following command:

```
test url-info-cloud <URL>
```

STEP 4 | If the URL stored in the cloud has the correct category, remove the URL from the dataplane and the management plane caches.

Run the following command to delete a URL from the dataplane cache:

```
clear url-cache url <URL>
```

Run the following command to delete a URL from the management plane cache:

```
delete url-database url <URL>
```

The next time the firewall queries for the category of the given URL, the request will be forwarded to the management plane and then to the cloud. This should resolve the category lookup issue. If problems persist, see the next step to submit a categorization change request.

STEP 5 | To submit a change request from the web interface, go to the URL log and select the log entry for the URL you would like to have changed.

STEP 6 | Click the **Request Categorization** change link and follow instructions. You can also request a category change from the Palo Alto Networks Test A Site website by searching for the URL and then clicking the **Request Change** icon. To view a list of all available categories

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

with descriptions of each category, refer to https://urlfiltering.paloaltonetworks.com/CategoryList.aspx.

If your change request is approved, you will receive an email notification. You then have two options to ensure that the URL category is updated on the firewall:

- Wait until the URL in the cache expires and the next time the URL is accessed by a user, the new categorization update will be put in the cache.
- Run the following command to force an update in the cache:

```
request url-filtering update url <URL>
```

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

PAN-DB Private Cloud

The PAN-DB private cloud is an on-premises solution for organizations that restrict the usage of cloud services. With this on-premises solution, you can deploy one or more M-600 appliances as PAN-DB servers within your network or data center. The firewalls query the PAN-DB private cloud to perform URL lookups, instead of accessing the PAN-DB public cloud.

The process for performing URL lookups, in both the private and the public cloud is the same for the firewalls on the network. By default, the firewall is configured to access the public PAN-DB cloud. If you deploy a PAN-DB private cloud, you must configure the firewalls with a list of IP addresses or FQDNs to access the server(s) in the private cloud.

Firewalls running PAN-OS 5.0 or later versions can communicate with the PAN-DB private cloud.

When you set up the PAN-DB private cloud, you can either configure the M-600 appliance(s) to have direct internet access or keep it completely offline. Because the M-600 appliance requires database and content updates to perform URL lookups, if the appliance does not have an active internet connection, you must manually download the updates to a server on your network and then, import the updates using SCP into each M-600 appliance in the PAN-DB private cloud. In addition, the appliances must be able to obtain the seed database and any other regular or critical content updates for the firewalls that it services.

To authenticate the firewalls that connect to the PAN-DB private cloud, a set of default server certificates are packaged with the appliance; you cannot import or use another server certificate for authenticating the firewalls. If you change the hostname on the M-600 appliance, the appliance automatically generates a new set of certificates to authenticate the firewalls.

- M-600 Appliance for PAN-DB Private Cloud
- Set Up the PAN-DB Private Cloud

M-600 Appliance for PAN-DB Private Cloud

To deploy a PAN-DB private cloud, you need one or more M-600 appliances. The M-600 appliance ships in Panorama mode, and to be deployed as PAN-DB private cloud you must set it up to operate in PAN-URL-DB mode. In the PAN-URL-DB mode, the appliance provides URL categorization services for enterprises that do not want to use the PAN-DB public cloud.

The M-600 appliance when deployed as a PAN-DB private cloud uses two ports- MGT (Eth0) and Eth1; Eth2 is not available for use. The management port is used for administrative access to the appliance and for obtaining the latest content updates from the PAN-DB public cloud or from a server on your network. For communication between the PAN-DB private cloud and the firewalls on the network, you can use the MGT port or Eth1.

The M-200 appliance cannot be deployed as a PAN-DB private cloud.

The M-600 appliance in PAN-URL-DB mode:

- Does not have a web interface, it only supports a command line interface (CLI).
- Cannot be managed by Panorama.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

- Cannot be deployed in a high availability pair.
- Does not require a URL Filtering license. The firewalls, must have a valid PAN-DB URL Filtering license to connect with and query the PAN-DB private cloud.
- Ships with a set of default server certificates that are used to authenticate the firewalls that connect to the PAN-DB private cloud. You cannot import or use another server certificate for authenticating the firewalls. If you change the hostname on the M-600 appliance, the appliance automatically generates a new set of certificates to authenticate the firewalls that it services.
- Can be reset to Panorama mode only. If you want to deploy the appliance as a Dedicated Log Collector, switch to Panorama mode and then set it in Log Collector mode.

Table 4: Differences Between the PAN-DB Public Cloud and PAN-DB Private Cloud

Differences	PAN-DB Public Cloud	PAN-DB Private Cloud
Content and Database Updates	Content (regular and critical) updates and full database updates are published multiple times during the day. The PAN-DB public cloud updates the URL categories malware and phishing every five minutes. The firewall checks for critical updates whenever it queries the cloud servers for URL lookups.	Content updates and full URL database updates are available once a day during the work week.
URL Categorization Requests	Submit URL categorization change requests using the following options: • Palo Alto Networks Test A Site website. • URL Filtering profile setup page on the firewall. • URL Filtering log on the firewall.	Submit URL categorization change requests only using the Palo Alto Networks Test A Site website.
Unresolved URL Queries	If the firewall cannot resolve a URL query, the request is sent to the servers in the public cloud.	If the firewall cannot resolve a query, the request is sent to the M-600 appliance(s) in the PAN-DB private cloud. If there is no match for the URL, the PAN-DB private cloud sends a category *unknown* response to the firewall; the request is not sent to the public cloud unless you have configured the M-600 appliance to access the PAN-DB public cloud. If the M-600 appliance(s) that constitute your PAN-DB private cloud is configured to be completely offline,

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Differences	PAN-DB Public Cloud	PAN-DB Private Cloud
		it does not send any data or analytics to the public cloud.

Set Up the PAN-DB Private Cloud

To deploy one or more M-600 appliances as a PAN-DB private cloud within your network or data center, you must complete the following tasks:

- Configure the PAN-DB Private Cloud
- Configure the Firewalls to Access the PAN-DB Private Cloud
- Configure Authentication with Custom Certificates on the PAN-DB Private Cloud

Configure the PAN-DB Private Cloud

STEP 1 | Rack mount the M-600 appliance.

Refer to the M-600 Hardware Reference Guide for instructions.

STEP 2 | Register the M-600 appliance.

For instructions on registering the M-600 appliance, see Register the Firewall.

STEP 3 | Perform initial configuration of the M-600 Appliance.

The M-600 appliance in PAN-DB mode uses two ports- MGT (Eth0) and Eth1; Eth2 is not used in PAN-DB mode. The management port is used for administrative access to the appliance and for obtaining the latest content updates from the PAN-DB public cloud. For communication between the appliance (PAN-DB server) and the firewalls on the network, you can use the MGT port or Eth1.

1. Connect to the M-600 appliance in one of the following ways:
   - Attach a serial cable from a computer to the Console port on the M-600 appliance and connect using a terminal emulation software (9600-8-N-1).
   - Attach an RJ-45 Ethernet cable from a computer to the MGT port on the M-600 appliance. From a browser, go to https://192.168.1.1. Enabling access to this

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

URL might require changing the IP address on the computer to an address in the 192.168.1.0 network (for example, 192.168.1.2).

2. When prompted, log in to the appliance. Log in using the default username and password (admin/admin). The appliance will begin to initialize.
3. Configure network access settings including the IP address for the MGT interface:

```
set deviceconfig system ip-address  <server-IP>
 netmask <netmask> default-gateway <gateway-IP> dns-setting
 servers primary <DNS-IP>
```

where *<server-IP>* is the IP address you want to assign to the management interface of the server, *<netmask>* is the subnet mask, *<gateway-IP>* is the IP address of the network gateway, and *<DNS-IP>* is the IP address of the primary DNS server.

4. Configure network access settings including the IP address for the Eth1 interface:

```
set deviceconfig system eth1 ip-address  <server-IP>
 netmask <netmask> default-gateway <gateway-IP> dns-setting
 servers primary <DNS-IP>
```

where *<server-IP>* is the IP address you want to assign to the data interface of the server, *<netmask>* is the subnet mask, *<gateway-IP>* is the IP address of the network gateway, and *<DNS-IP>* is the IP address of the DNS server.

5. Save your changes to the PAN-DB server.

   commit

STEP 4 | Switch to PAN-DB private cloud mode.

1. To switch to PAN-DB mode, use the CLI command:

```
request system system-mode pan-url-db
```

   You can switch from Panorama mode to PAN-DB mode and back; and from Panorama mode to Log Collector mode and back. Switching directly from PAN-DB mode to Log Collector mode or vice versa is not supported. When switching operational mode, a data reset is triggered. With the exception of management access settings, all existing configuration and logs will be deleted on restart.

2. Use the following command to verify that the mode is changed:

```
show pan-url-cloud-status
hostname: M-600
ip-address: 1.2.3.4
netmask: 255.255.255.0
default-gateway: 1.2.3.1
ipv6-address: unknown
ipv6-link-local-address: fe80:00/64
ipv6-default-gateway:
mac-address: 00:56:90:e7:f6:8e
```

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

```
time: Mon Apr 27 13:43:59 2015
uptime: 10 days, 1:51:28
family: m
model: M-600
serial: 0073010000xxx
sw-version: 7.0.0
app-version: 492-2638
app-release-date: 2015/03/19  20:05:33
av-version: 0
av-release-date: unknown
wf-private-version: 0
wf-private-release-date: unknown
logdb-version: 7.0.9
platform-family: m
pan-url-db: 20150417-220
system-mode: Pan-URL-DB
operational-mode: normal
```

3. Use the following command to check the version of the cloud database on the appliance:

```
show pan-url-cloud-status
Cloud status:            Up
URL database version:    20150417-220
```

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 5 | Install content and database updates.

The appliance only stores the currently running version of the content and one earlier version.

Pick one of the following methods of installing the content and database updates:

- If the PAN-DB server has direct Internet access use the following commands:
  1. To check whether a new version is published use:

     request pan-url-db upgrade check
  2. To check the version that is currently installed on your server use:

     request pan-url-db upgrade info
  3. To download and install the latest version:
     - **request pan-url-db upgrade download latest**
     - **request pan-url-db upgrade install** *<version latest | file>*
  4. To schedule the M-600 appliance to automatically check for updates:

     set deviceconfig system update-schedule pan-url-db recurring weekly action download-and-install day-of-week *<day of week>* at *<hr:min>*
- If the PAN-DB server is offline, access the Palo Alto Networks Customer Support web site to download and save the content updates to an SCP server on your network. You can then import and install the updates using the following commands:
  - **scp import pan-url-db remote-port** *<port-number>* from username@host:path
  - **request pan-url-db upgrade install file** *<filename>*

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 6 | Set up administrative access to the PAN-DB private cloud.

The appliance has a default `admin` account. Any additional administrative users that you create can either be superusers (with full access) or superusers with read-only access.

PAN-DB private cloud does not support the use of RADIUS VSAs. If the VSAs used on the firewall or Panorama are used for enabling access to the PAN-DB private cloud, an authentication failure will occur.

- To set up a local administrative user on the PAN-DB server:
  1. **configure**
  2. **set mgt-config users** *<username>* permissions role-based <superreader | superuser> yes
  3. **set mgt-config users** *<username>* password
  4. **Enter password:xxxxx**
  5. **Confirm password:xxxxx**
  6. **commit**
- To set up an administrative user with RADIUS authentication:
  1. Create RADIUS server profile.

     set shared server-profile radius *<server_profile_name>* server *<server_name>* ip-address *<ip_address>* port *<port_no>* secret *<shared_password>*
  2. Create authentication-profile.

     set shared authentication-profile *<auth_profile_name>* user-domain *<domain_name_for_authentication>* allow-list <all> method radius server-profile *<server_profile_name>*
  3. Attach the authentication-profile to the user.

     set mgt-config users *<username>* authentication-profile *<auth_profile_name>*
  4. Commit the changes.

     commit
- To view the list of users:.

```
show mgt-config users
users {
admin {
phash fnRL/G5lXVMug;
permissions {
```

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

```
role-based {
superuser yes;
}
}
}
admin_user_2 {
permissions {
role-based {
superreader yes;
}
}
authentication-profile RADIUS;
}
}
```

STEP 7 | Configure the firewalls to access the PAN-DB private cloud.

Configure the Firewalls to Access the PAN-DB Private Cloud

When using the PAN-DB public cloud, each firewall accesses the PAN-DB servers in the AWS cloud to download the list of eligible servers to which it can connect for URL lookups. With the PAN-DB private cloud, you must configure the firewalls with a (static) list of your PAN-DB private cloud servers that will be used for URL lookups. The list can contain up to 20 entries; IPv4 addresses, IPv6 addresses, and FQDNs are supported. Each entry on the list— IP address or FQDN —must be assigned to the management port and/or eth1 of the PAN-DB server.

STEP 1 | From the PAN-OS CLI, add a list of static PAN-DB private cloud servers used for URL lookups.

- Use the following CLI command to add private PAN-DB server IP addresses:

```
> configure
```

```
# set deviceconfig setting pan-url-db cloud-static-list <IP
 addresses>
```

Or, in the web interface for each firewall, select **Device** > **Setup** > **Content-ID**, edit the URL Filtering section and enter the **PAN-DB Server** IP address(es) or FQDN(s). The list must be comma separated.

- To delete the entries for the private PAN-DB servers, use the following command:

```
# delete deviceconfig setting pan-url-db cloud-static-list <IP
 addresses>
```

When you delete the list of private PAN-DB servers, a re-election process is triggered on the firewall. The firewall first checks for the list of PAN-DB private cloud servers and when it cannot find one, the firewall accesses the PAN-DB servers in the AWS cloud to download the list of eligible servers to which it can connect.

STEP 2 | Enter **# commit** to save your changes.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 3 | To verify that the change is effective, use the following CLI command on the firewall:

```
> show url-cloud status
Cloud status:            Up
URL database version:    20150417-220
```

Configure Authentication with Custom Certificates on the PAN-DB Private Cloud

By default, a PAN-DB server uses predefined certificates for mutual authentication to establish the SSL connections used for management access and inter-device communication. However, you can configure authentication using custom certificates instead. Custom certificates allow you to establish a unique chain of trust to ensure mutual authentication between your PAN-DB server and firewalls. In the case of a PAN-DB private cloud, the firewall acts as the client and the PAN-DB server acts as the server.

STEP 1 | Obtain key pairs and certificate authority (CA) certificates for the PAN-DB server and firewall.

STEP 2 | Import the CA certificate to validate the certificate on the firewall.

1. Log in to the CLI on the PAN-DB server and enter configuration mode.

   ```
   admin@M-600> configure
   ```

2. Use TFTP or SCP to import the CA certificate.

   ```
   admin@M-600# {tftp | scp} import certificate from <value>
    file <value> remote-port <1-65535> source-ip <ip/netmask>
    certificate-name <value> passphrase <value> format {pkcs12 |
    pem}
   ```

STEP 3 | Use TFTP or SCP to import the key pair that contains the server certificate and private key for the PAN-DB M-600 appliance.

```
admin@M-600# {tftp | scp} import keypair from <value> file <value>
 remote-port <1-65535> source-ip <ip/netmask> certificate-
name <value> passphrase <value> format {pkcs12 | pem}
```

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 4 | Configure a certificate profile that includes the root CA and intermediate CA. This certificate profile defines the device authentication between the PAN-DB server and the firewall.

1. In the CLI of the PAN-DB server, enter configuration mode.

```
admin@M-600> configure
```

2. Name the certificate profile.

```
admin@M-600# set shared certificate-profile <name>
```

3. (Optional) Set the user domain.

```
admin@M-600# set shared certificate-profile <name>
 domain <value>
```

4. Configure the CA.

   Default-ocsp-url *and* ***ocsp-verify-cert*** *are optional parameters.*

   admin@M-600# **set shared certificate-profile <*name*> CA <*name*>**

```
admin@M-600# set shared certificate-profile <name> CA <name>
 [default-ocsp-url <value>]
```

```
admin@M-600# set shared certificate-profile <name> CA <name>
 [ocsp-verify-cert <value>]
```


©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 5 | Configure an SSL/TLS profile for the PAN-DB M-600 appliance. This profile defines the certificate and protocol range that PAN-DB and client devices use for SSL/TLS services.

1. Identify the SSL/TLS profile.

```
admin@M-600# set shared ssl-tls-service-profile <name>
```

2. Select the certificate.

```
admin@M-600# set shared ssl-tls-service-profile <name>
 certificate <value>
```

3. Define the SSL/TLS range.

   *PAN-OS 8.0 and later releases support TLS 1.2 and later TLS versions only. You must set the max version to **TLS 1.2** or **max**.*

```
admin@M-600# set shared ssl-tls-service-profile <name>
 protocol-settings min-version {tls1-0 | tls1-1 | tls1-2
```

```
admin@M-600# set shared ssl-tls-service-profile <name>
 protocol-settings max-version {tls1-0 | tls1-1 | tls1-2 | max
```

STEP 6 | Configure secure server communication on PAN-DB.

1. Set the SSL/TLS profile. This SSL/TLS service profile applies to all SSL connections between PAN-DB and firewalls.

```
admin@M-600# set deviceconfig setting management secure-conn-
server ssl-tls-service-profile <ssltls-profile>
```

2. Set the certificate profile.

```
admin@M-600# set deviceconfig setting management secure-conn-
server certificate-profile <certificate-profile>
```

3. Set the disconnect wait time in number of minutes that PAN-DB should wait before breaking and reestablishing the connection with its firewall (range is 0 to 44,640).

```
admin@M-600# set deviceconfig setting management secure-conn-
server disconnect-wait-time <0-44640
```

STEP 7 | Import the CA certificate to validate the certificate for the PAN-DB M-600 appliance.

1. Log in to the firewall web interface.
2. Import the CA certificate.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 8 | Configure a local or a SCEP certificate for the firewall.

1. If you are a local certificate, then import the key pair for the firewall.
2. If you are a SCEP certificate for the firewall, configure a SCEP profile.

STEP 9 | Configure the certificate profile for the firewall. You can configure this on each firewall individually or you can push the configuration from Panorama to the firewalls as part of a template.

1. Select **Device** > **Certificate Management** > **Certificate Profile** for firewalls or **Panorama** > **Certificate Management** > **Certificate Profile** for Panorama.
2. Configure a Certificate Profile.

STEP 10 | Deploy custom certificates on each firewall. You can either deploy certificates centrally from Panorama or configure them manually on each firewall.

1. Log in to the firewall web interface.
2. Select **Device** > **Setup** > **Management** for a firewall or **Panorama** > **Setup** > **Management** for Panorama and **Edit** the Secure Communication
3. Select the **Certificate Type**, **Certificate**, and **Certificate Profile** from the respective drop-downs.
4. In the Customize Communication settings, select **PAN-DB Communication**.
5. Click **OK**.
6. **Commit** your changes.

After committing your changes, the firewalls do not terminate their current sessions with the PAN-DB server until after the **Disconnect Wait Time**. The disconnect wait time begins counting down after you enforce the use of custom certificates in the next step.

STEP 11 | After deploying custom certificates on all firewalls, enforce custom certificate authentication.

1. Log in to the CLI on the PAN-DB server and enter configuration mode.

```
admin@M-600> configure
```

2. Enforce the use of custom certificates.

```
admin@M-600# set deviceconfig setting management secure-conn-
server disable-pre-defined-cert yes
```

After committing this change, the disconnect wait time begins counting down (if you configured setting on PAN-DB). When the wait time ends, PAN-DB and its firewall connect using only the configured certificates.

STEP 12 | You have two choices when adding new firewalls or Panorama to your PAN-DB private cloud deployment.

- If you did not enable **Custom Certificates Only** then you can add a new firewall to the PAN-DB private cloud and then deploy the custom certificate as described above.
- If you enabled **Custom Certificates Only** on the PAN-DB private cloud, then you can must deploy the custom certificates on the firewalls before connecting them to the PAN-DB private cloud.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Enable SSL/TLS Handshake Inspection

SSL/TLS handshake inspection closes a gap in threat detection for SSL/TLS web traffic marked for decryption. When enabled, the Content and Threat Detection (CTD) engine of the firewall inspects HTTPS traffic for potential threats during the SSL/TLS handshake. The firewall uses data in the handshake to identify the traffic and enforce applicable Security policy rules. Examining the handshake improves network security and optimizes our URL Filtering solution by preventing threats and enforcing Security policy actions on web traffic as early as possible.

Specifically, the firewall scans the Client Hello message for the *Server Name Indication (SNI)* field, an extension to the SSL/TLS protocol that contains the hostname of a requested website. From the hostname, the firewall can derive the URL category and server destination of the traffic. Then, it evaluates the URL category against the URL Filtering profiles of matching Security policy rules to determine which actions to enforce. If the firewall detects a threat, such as a malicious web server in the SNI field, or policy dictates that the website be blocked, it will terminate the handshake and end the web session immediately. If no threat is detected and the traffic is allowed per policy, the client and server can complete the SSL/TLS handshake and exchange application data through the secure connection.

URL Filtering response pages do not display for sites blocked by the firewall during SSL/TLS handshake inspections. After detecting traffic from blocked categories, the firewall resets the HTTPS connection, ending the handshake and preventing user notification by response page. Instead, the browser displays a standard connection error message.

Details of successful SSL/TLS handshakes and sessions will be in the Traffic and Decryption logs. If the firewall blocks web sessions during the SSL/TLS handshake, it will not generate Decryption logs. You can find details of failed sessions in the URL Filtering logs, however.

The following procedure details the requirements and steps needed to enable SSL/TLS handshake inspection:

STEP 1 | Select **Device > Licenses** to confirm that you have an active Advanced URL Filtering or legacy URL Filtering license.

STEP 2 | Verify that you decrypt SSL/TLS traffic through either SSL Forward Proxy or SSL Inbound Inspection.

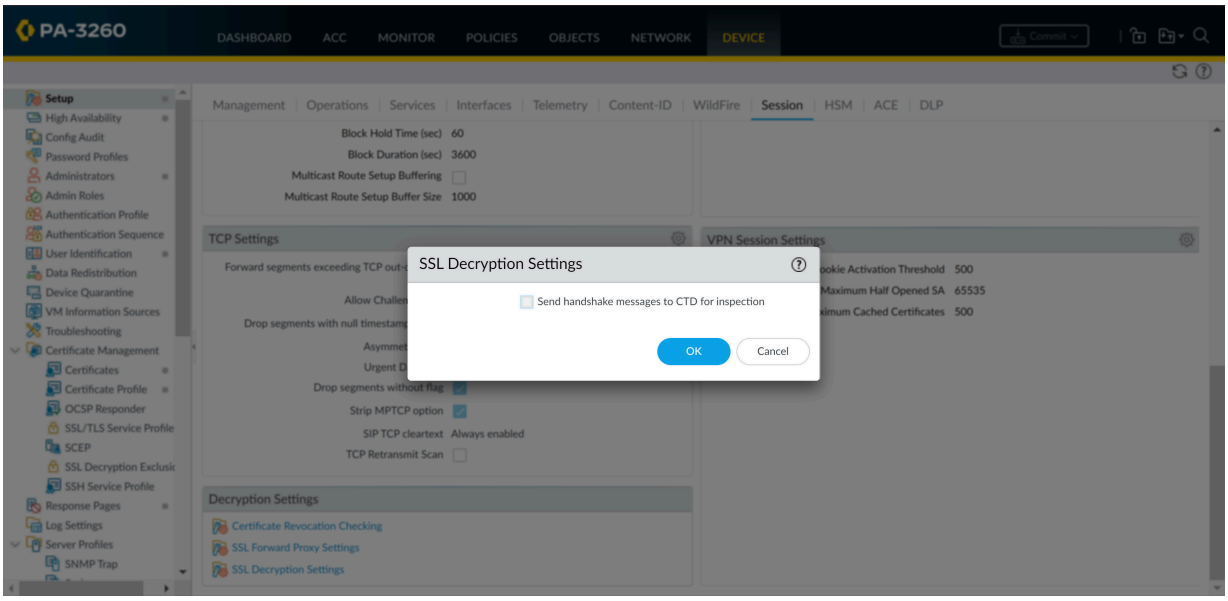
 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 3 | Enable inspection of SSL/TLS handshakes by CTD. By default, the option is disabled.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets



https://t.me/learningnets

1. Select **Device** > **Setup** > **Session** > **Decryption Settings** > **SSL Decryption Settings**.
2. Select **Send handshake messages to CTD for inspection**.

   Alternatively, you can use the **`set deviceconfig setting ssl-decrypt scan-handshake <yes|no>`** CLI command.
3. Click **OK**.

STEP 4 | **Commit** your configuration changes.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Quality of Service

Quality of Service (QoS) is a set of technologies that work on a network to guarantee its ability to dependably run high-priority applications and traffic under limited network capacity. QoS technologies accomplish this by providing differentiated handling and capacity allocation to specific flows in network traffic. This enables the network administrator to assign the order in which traffic is handled, and the amount of bandwidth afforded to traffic.

Palo Alto Networks Application Quality of Service (QoS) provides basic QoS applied to networks and extends it to provide QoS to applications and users.

Use the following topics to learn about and configure Palo Alto Networks application-based QoS:

> QoS Overview
> QoS Concepts
> Configure QoS
> Configure QoS for a Virtual System
> Enforce QoS Based on DSCP Classification
> QoS Use Cases

Use the Palo Alto Networks product comparison tool to view the QoS features supported on your firewall model. Select two or more product models and click **Compare Now** to view QoS feature support for each model (for example, you can check if your firewall model supports QoS on subinterfaces and if so, the maximum number of subinterfaces on which QoS can be enabled).

QoS on Aggregate Ethernet (AE) interfaces is supported on PA-7000 Series, PA-5200 Series, and PA-3200 Series firewalls running PAN-OS 7.0 or later release versions.

https://t.me/learningnets

QoS Overview

Use QoS to prioritize and adjust quality aspects of network traffic. You can assign the order in which packets are handled and allot bandwidth, ensuring preferred treatment and optimal levels of performance are afforded to selected traffic, applications, and users.

Service quality measurements subject to a QoS implementation are bandwidth (maximum rate of transfer), throughput (actual rate of transfer), latency (delay), and jitter (variance in latency). The capability to shape and control these service quality measurements makes QoS of particular importance to high-bandwidth, real-time traffic such as voice over IP (VoIP), video conferencing, and video-on-demand that has a high sensitivity to latency and jitter. Additionally, use QoS to achieve outcomes such as the following:

- Prioritize network and application traffic, guaranteeing high priority to important traffic or limiting non-essential traffic.
- Achieve equal bandwidth sharing among different subnets, classes, or users in a network.
- Allocate bandwidth externally or internally or both, applying QoS to both upload and download traffic or to only upload or download traffic.
- Ensure low latency for customer and revenue-generating traffic in an enterprise environment.
- Perform traffic profiling of applications to ensure bandwidth usage.

QoS implementation on a Palo Alto Networks firewall begins with three primary configuration components that support a full QoS solution: a QoS Profile, a QoS Policy, and setting up the QoS Egress Interface. Each of these options in the QoS configuration task facilitate a broader process that optimizes and prioritizes the traffic flow and allocates and ensures bandwidth according to configurable parameters.

The figure QoS Traffic Flow shows traffic as it flows from the source, is shaped by the firewall with QoS enabled, and is ultimately prioritized and delivered to its destination.

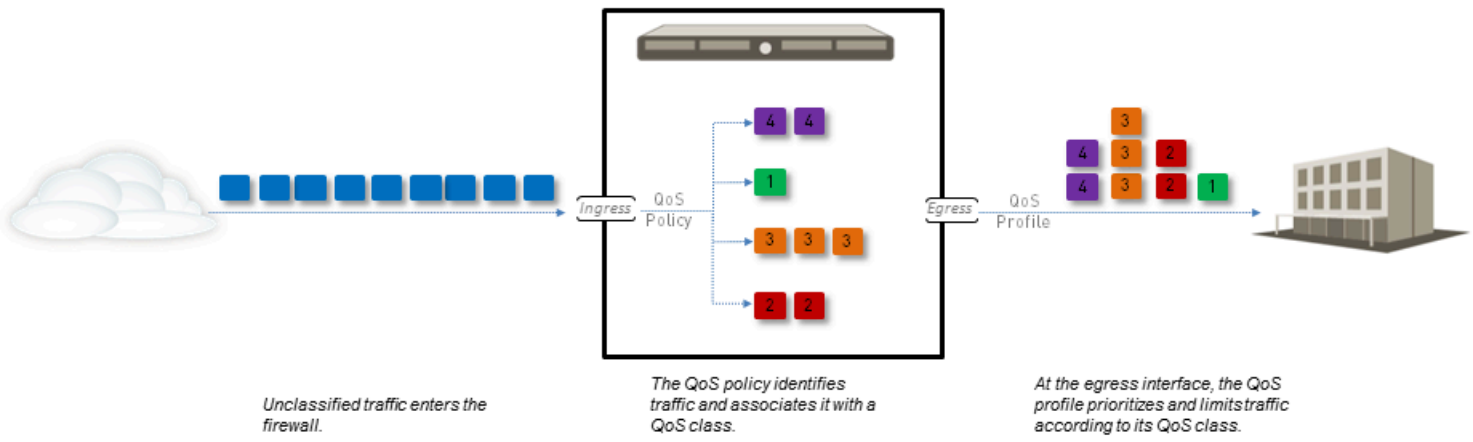


Figure 6: QoS Traffic Flow

The QoS configuration options allow you to control the traffic flow and define it at different points in the flow. The figure QoS Traffic Flow indicates where the configurable options define the traffic flow. A QoS policy rule allows you to define traffic you want to receive QoS treatment and assign

https://t.me/learningnets

that traffic a QoS class. The matching traffic is then shaped based on the QoS profile class settings as it exits the physical interface.

Each of the QoS configuration components influence each other and the QoS configuration options can be used to create a full and granular QoS implementation or can be used sparingly with minimal administrator action.

Each firewall model supports a maximum number of ports that can be configured with QoS. Refer to the spec sheet for your firewall model or use the product comparison tool to view QoS feature support for two or more firewalls on a single page.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

QoS Concepts

Use the following topics to learn about the different components and mechanisms of a QoS configuration on a Palo Alto Networks firewall:

- QoS for Applications and Users
- QoS Policy
- QoS Profile
- QoS Classes
- QoS Priority Queuing
- QoS Bandwidth Management
- QoS Egress Interface
- QoS for Clear Text and Tunneled Traffic

QoS for Applications and Users

A Palo Alto Networks firewall provides basic QoS, controlling traffic leaving the firewall according to network or subnet, and extends the power of QoS to also classify and shape traffic according to application and user. The Palo Alto Networks firewall provides this capability by integrating the features App-ID and User-ID with the QoS configuration. App-ID and User-ID entries that exist to identify specific applications and users in your network are available in the QoS configuration so that you can easily specify applications and users for which you want to manage and/or guarantee bandwidth.

QoS Policy

Use a QoS policy rule to define traffic to receive QoS treatment (either preferential treatment or bandwidth-limiting) and assigns such traffic a QoS class of service.

Define a QoS policy rule to match to traffic based on:

- Applications and application groups.
- Source zones, source addresses, and source users.
- Destination zones and destination addresses.
- Services and service groups limited to specific TCP and/or UDP port numbers.
- URL categories, including custom URL categories.
- Differentiated Services Code Point (DSCP) and Type of Service (ToS) values, which are used to indicate the level of service requested for traffic, such as high priority or best effort delivery.

You cannot apply DSCP code points or QoS to SSL Forward Proxy, SSL Inbound Inspection, and SSH Proxy traffic.

Set up multiple QoS policy rules (**Policies** > **QoS**) to associate different types of traffic with different QoS Classes of service.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Because QoS is enforced on traffic as it egresses the firewall, your QoS policy rule is applied to traffic after the firewall has enforced all other security policy rules, including Network Address Translation (NAT) rules. If you want to apply QoS treatment to traffic based on source, make sure to specify the post-NAT source address in a QoS policy rule (do not use the pre-NAT source address).

QoS Profile

Use a QoS profile rule to define values of up to eight QoS Classes contained within that single profile rule.

With a QoS profile rule, you can define QoS Priority Queuing and QoS Bandwidth Management for QoS classes. Each QoS profile rule allows you to configure individual bandwidth and priority settings for up eight QoS classes, as well as the total bandwidth alloted for the eight classes combined. Attach the QoS profile rule (or multiple QoS profile rules) to a physical interface to apply the defined priority and bandwidth settings to the traffic exiting that interface.

A default QoS profile rule is available on the firewall. The default profile rule and the classes defined in the profile do not have predefined maximum or guaranteed bandwidth limits.

To define priority and bandwidth settings for QoS classes, see Step Add a QoS profile rule.

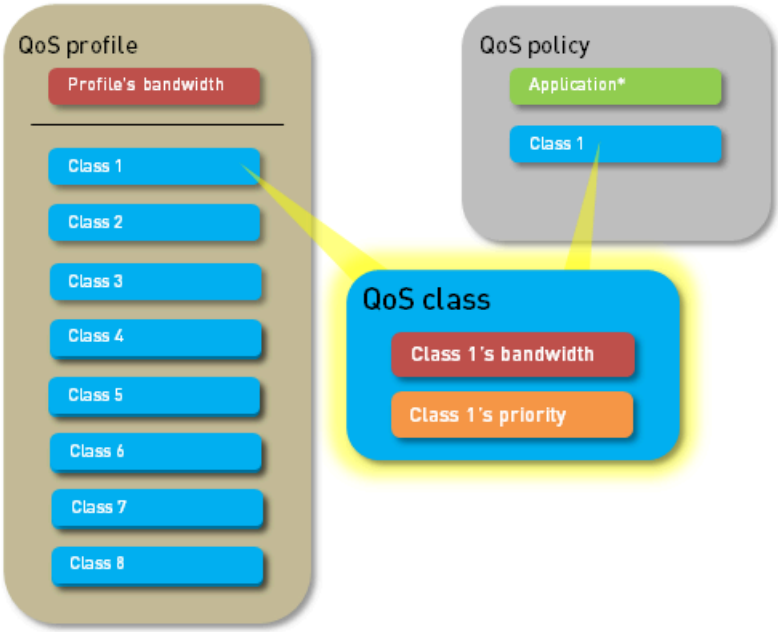
QoS Classes

A QoS class determines the priority and bandwidth for traffic matching a QoS Policy rule. You can use a QoS Profile rule to define QoS classes. There are up to eight definable QoS classes in a single QoS profile. Unless otherwise configured, traffic that does not match a QoS class is assigned a class of 4.

QoS Priority Queuing and QoS Bandwidth Management, the fundamental mechanisms of a QoS configuration, are configured within the QoS class definition (see Step 4). For each QoS class, you can set a priority (real-time, high, medium, and low) and the maximum and guaranteed bandwidth for matching traffic. QoS priority queuing and bandwidth management determine the order of traffic and how traffic is handled upon entering or leaving a network.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets



QoS Priority Queuing

One of four priorities can be enforced for a QoS class: real-time, high, medium, and low. Traffic matching a QoS policy rule is assigned the QoS class associated with that rule, and the firewall treats the matching traffic based on the QoS class priority. Packets in the outgoing traffic flow are queued based on their priority until the network is ready to process the packets. Priority queuing allows you to ensure that important traffic, applications, and users take precedence. Real-time priority is typically used for applications that are particularly sensitive to latency, such as voice and video applications.

QoS Bandwidth Management

QoS bandwidth management allows you to control traffic flows on a network so that traffic does not exceed network capacity (resulting in network congestion) and also allows you to allocate bandwidth for certain types of traffic and for applications and users. With QoS, you can enforce bandwidth for traffic on a narrow or a broad scale. A QoS profile rule allows you to set bandwidth limits for individual QoS classes and the total combined bandwidth for all eight QoS classes. As part of the steps to Configure QoS, you can attach the QoS profile rule to a physical interface to enforce bandwidth settings on the traffic exiting that interface—the individual QoS class settings are enforced for traffic matching that QoS class (QoS classes are assigned to traffic matching QoS Policy rules) and the overall bandwidth limit for the profile can be applied to all clear text traffic, specific clear text traffic originating from source interfaces and source subnets, all tunneled traffic, and individual tunnel interfaces. You can add multiple profile rules to a single QoS interface to apply varying bandwidth settings to the traffic exiting that interface.

The following fields support QoS bandwidth settings:

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

- **Egress Guaranteed**—The amount of bandwidth guaranteed for matching traffic. When the egress guaranteed bandwidth is exceeded, the firewall passes traffic on a best-effort basis. Bandwidth that is guaranteed but is unused continues to remain available for all traffic. Depending on your QoS configuration, you can guarantee bandwidth for a single QoS class, for all or some clear text traffic, and for all or some tunneled traffic.

  Example:

  Class 1 traffic has 5 Gbps of egress guaranteed bandwidth, which means that 5 Gbps is available but is not reserved for class 1 traffic. If Class 1 traffic does not use or only partially uses the guaranteed bandwidth, the remaining bandwidth can be used by other classes of traffic. However, during high traffic periods, 5 Gbps of bandwidth is absolutely available for class 1 traffic. During these periods of congestion, any Class 1 traffic that exceeds 5 Gbps is best effort.
- **Egress Max**—The overall bandwidth allocation for matching traffic. The firewall drops traffic that exceeds the egress max limit that you set. Depending on your QoS configuration, you can set a maximum bandwidth limit for a QoS class, for all or some clear text traffic, for all or some tunneled traffic, and for all traffic exiting the QoS interface.

  The cumulative guaranteed bandwidth for the QoS profile rules attached to the interface must not exceed the total bandwidth allocated to the interface.

To define bandwidth settings for QoS classes, see Step Add a QoS profile rule. To then apply those bandwidth settings to clear text and tunneled traffic, and to set the overall bandwidth limit for a QoS interface, see Step Enable QoS on a physical interface.

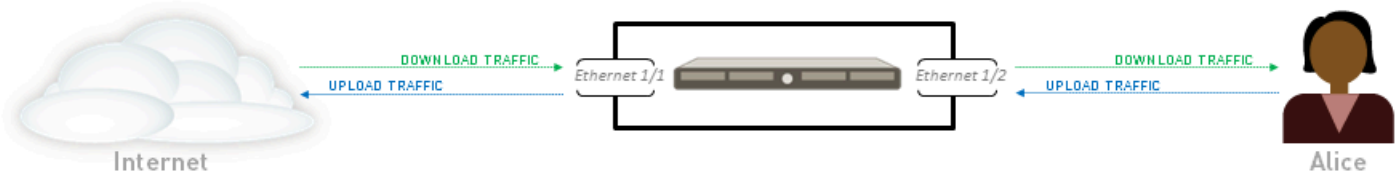
QoS Egress Interface

Enabling a QoS profile rule on the egress interface of the traffic identified for QoS treatment completes a QoS configuration. The ingress interface for QoS traffic is the interface on which the traffic enters the firewall. The egress interface for QoS traffic is the interface that traffic leaves the firewall from. QoS is always enabled and enforced on the egress interface for a traffic flow. The egress interface in a QoS configuration can either be the external- or internal-facing interface of the firewall, depending on the flow of the traffic receiving QoS treatment.

For example, in an enterprise network, if you are limiting employees' download traffic from a specific website, the egress interface in the QoS configuration is the firewall's internal interface, as the traffic flow is from the Internet, through the firewall, and to your company network. Alternatively, when limiting employees' upload traffic to the same website, the egress interface in the QoS configuration is the firewall's external interface, as the traffic you are limiting flows from your company network, through the firewall, and then to the Internet.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets



- The egress interface for Alice's download traffic is Ethernet 1/2. To prioritize or limit her download traffic, Alice enables QoS on Ethernet 1/2.
- The egress interface for Alice's upload traffic is Ethernet 1/1. To prioritize or limit her upload traffic, Alice enables QoS on Ethernet 1/1.

Because QoS is enforced on traffic as it egresses the firewall, your QoS policy rule is applied to traffic after the firewall has enforced all other security policy rules, including Network Address Translation (NAT) rules. If you want to apply QoS treatment to traffic based on source, you must specify the post-NAT source address in a QoS policy rule (do not use the pre-NAT source address).

Learn more about how to Identify the egress interface for applications that you want to receive QoS treatment.

QoS for Clear Text and Tunneled Traffic

At the minimum, enabling a QoS interfaces requires you to select a default QoS profile rule that defines bandwidth and priority settings for clear text traffic egressing the interface. However, when setting up or modifying a QoS interface, you can apply granular QoS settings to outgoing clear text traffic and tunneled traffic. QoS preferential treatment and bandwidth limiting can be enforced for tunneled traffic, for individual tunnel interfaces, and/or for clear text traffic originating from different source interfaces and source subnets. On Palo Alto Networks firewalls, *tunneled traffic* refers to tunnel interface traffic, specifically IPSec traffic in tunnel mode.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Configure QoS

Follow these steps to configure Quality of Service (QoS), which includes creating a QoS profile, creating a QoS policy, and enabling QoS on an interface.

STEP 1 | Identify the traffic you want to manage with QoS.

This example shows how to use QoS to limit web browsing.

Select **ACC** to view the **Application Command Center** page. Use the settings and charts on the **ACC** page to view trends and traffic related to Applications, URL filtering, Threat Prevention, Data Filtering, and HIP Matches.

Click any application name to display detailed application information.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 2 | Identify the egress interface for applications that you want to receive QoS treatment.

The egress interface for traffic depends on the traffic flow. If you are shaping incoming traffic, the egress interface is the internal-facing interface. If you are shaping outgoing traffic, the egress interface is the external-facing interface.

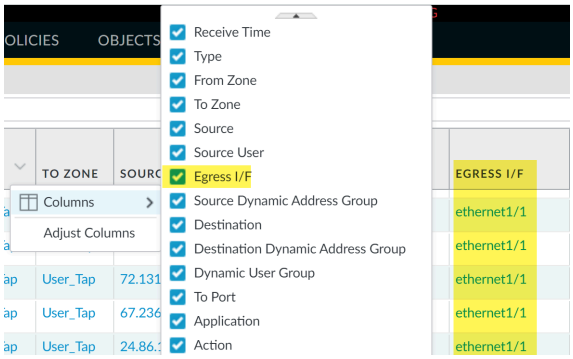
Select **Monitor** > **Logs** > **Traffic** to view the Traffic logs.

To filter and only show logs for a specific application:

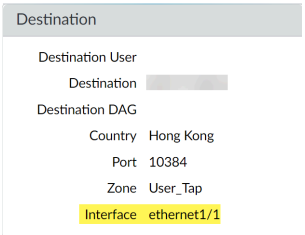
- If an entry is displayed for the application, click the underlined link in the Application column then click the Submit icon.
- If an entry is not displayed for the application, click the Add Log icon and search for the application.

The **Egress I/F** in the traffic logs displays each application's egress interface. To display the **Egress I/F** column if it is not displayed by default:

- Click any column header to add a column to the log:



- Click the spyglass icon to the left of any entry to display a detailed log that includes the application's egress interface listed in the Destination section:



https://t.me/learningnets

STEP 3 | Add a QoS policy rule.

A QoS policy rule defines the traffic to receive QoS treatment. The firewall assigns a QoS class of service to the traffic matched to the policy rule.



Because QoS is enforced on traffic as it egresses the firewall, your QoS policy rule is applied to traffic after the firewall has enforced all other security policy rules, including Network Address Translation (NAT) rules. If you want to apply QoS treatment to traffic based on source, you must specify the post-NAT source address in a QoS policy rule (do not use the pre-NAT source address).

1. Select **Policies** > **QoS** and **Add** a new policy rule.
2. On the **General** tab, give the QoS Policy Rule a descriptive **Name**.
3. Specify traffic to receive QoS treatment based on **Source**, **Destination**, **Application**, **Service/URL Category**, and **DSCP/ToS** values (the **DSCP/ToS** settings allow you to Enforce QoS Based on DSCP Classification).

   For example, select the **Application**, click **Add**, and select **web-browsing** to apply QoS to web browsing traffic.
4. (Optional) Continue to define additional parameters. For example, select **Source** and **Add** a **Source User** to provide QoS for a specific user's web traffic.
5. Select **Other Settings** and assign a **QoS Class** to traffic matching the policy rule. For example, assign Class 2 to the user1's web traffic.
6. Click **OK**.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 4 | Add a QoS profile rule.

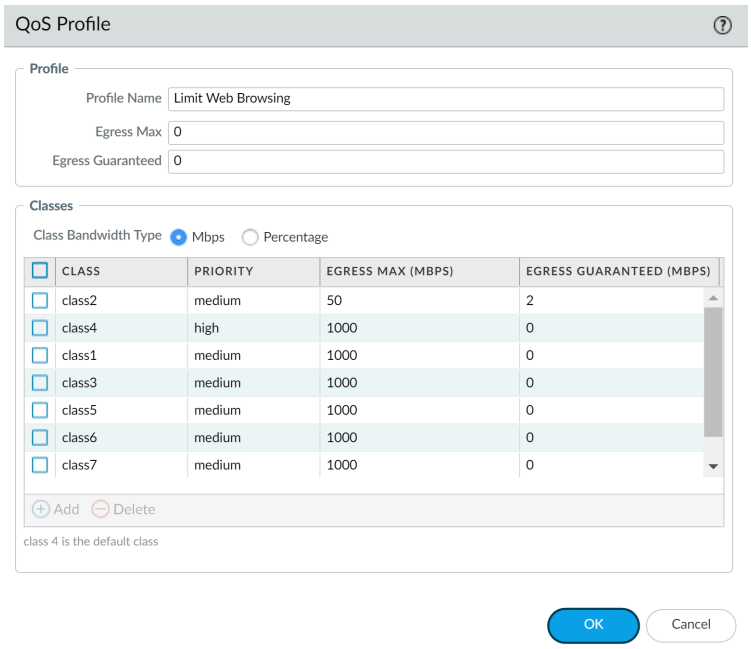
A QoS profile rule allows you to define the eight classes of service that traffic can receive, including priority, and enables QoS Bandwidth Management.

You can edit any existing QoS profile, including the default, by clicking the QoS profile name.

1. Select **Network** > **Network Profiles** > **QoS Profile** and **Add** a new profile.
2. Enter a descriptive **Profile Name**.
3. Set the overall bandwidth limits for the QoS profile rule:
   - Enter an **Egress Max** value to set the overall bandwidth allocation for the QoS profile rule.
   - Enter an **Egress Guaranteed** value to set the guaranteed bandwidth for the QoS Profile.

   Any traffic that exceeds the Egress Guaranteed value is best effort and not guaranteed. Bandwidth that is guaranteed but is unused continues to remain available for all traffic.
4. In the Classes section, specify how to treat up to eight individual QoS classes:
   1. **Add** a class to the QoS Profile.
   2. Select the **Priority** for the class: real-time, high, medium, or low.
   3. Enter the **Egress Max** and **Egress Guaranteed** bandwidth for traffic assigned to each QoS class.
5. Click **OK**.

In the following example, the QoS profile rule Limit Web Browsing limits Class 2 traffic to a maximum bandwidth of 50Mbps and a guaranteed bandwidth of 2Mbps.



https://t.me/learningnets

STEP 5 | Enable QoS on a physical interface.

Part of this step includes the option to select clear text and tunneled traffic for unique QoS treatment.

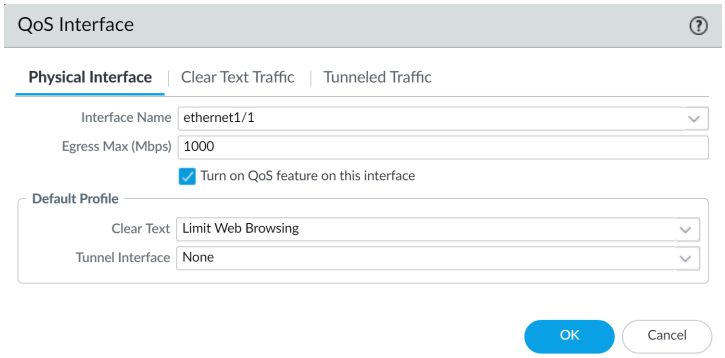
Check if the firewall model you're using supports enabling QoS on a subinterface by reviewing a summary of the Product Specifications.

1. Select **Network** > **QoS** and **Add** a QoS interface.
2. Select **Physical Interface** and choose the **Interface Name** of the interface on which to enable QoS.

   In the example, Ethernet 1/1 is the egress interface for web-browsing traffic (see Step 2).
3. Set the **Egress Max** bandwidth for all traffic exiting this interface.

   It is a best practice to always define the Egress Max value for a QoS interface. Ensure that the cumulative guaranteed bandwidth for the QoS profile rules attached to the interface does not exceed the total bandwidth allocated to the interface.
4. Select **Turn on QoS feature on this interface**.
5. In the Default Profile section, select a QoS profile rule to apply to all **Clear Text** traffic exiting the physical interface.
6. (Optional) Select a default QoS profile rule to apply to all tunneled traffic exiting the interface.

For example, enable QoS on ethernet 1/1 and apply the bandwidth and priority settings you defined for the QoS profile rule Limit Web Browsing (Step 4) to be used as the default settings for clear text egress traffic.



1. (Optional) Continue to define more granular settings to provide QoS for Clear Text and Tunneled Traffic. Settings configured on the **Clear Text Traffic** tab and the **Tunneled**

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Traffic tab automatically override the default profile settings for clear text and tunneled traffic on the Physical Interface tab.

- Select **Clear Text Traffic** and:
  - Set the **Egress Guaranteed** and **Egress Max** bandwidths for clear text traffic.
  - Click **Add** and apply a QoS profile rule to enforce clear text traffic based on source interface and source subnet.

    (PA-3200 Series, PA-5200 Series, PA-5400 Series, PA-7000 Series only) You must also select a destination interface when configuring a QoS policy rule if the rule is applied to a specific subinterface.
- Select **Tunneled Traffic** and:
  - Set the **Egress Guaranteed** and **Egress Max** bandwidths for tunneled traffic.
  - Click **Add** and attach a QoS profile rule to a single tunnel interface.

2. Click **OK**.

STEP 6 | Commit your changes.

Click **Commit**.

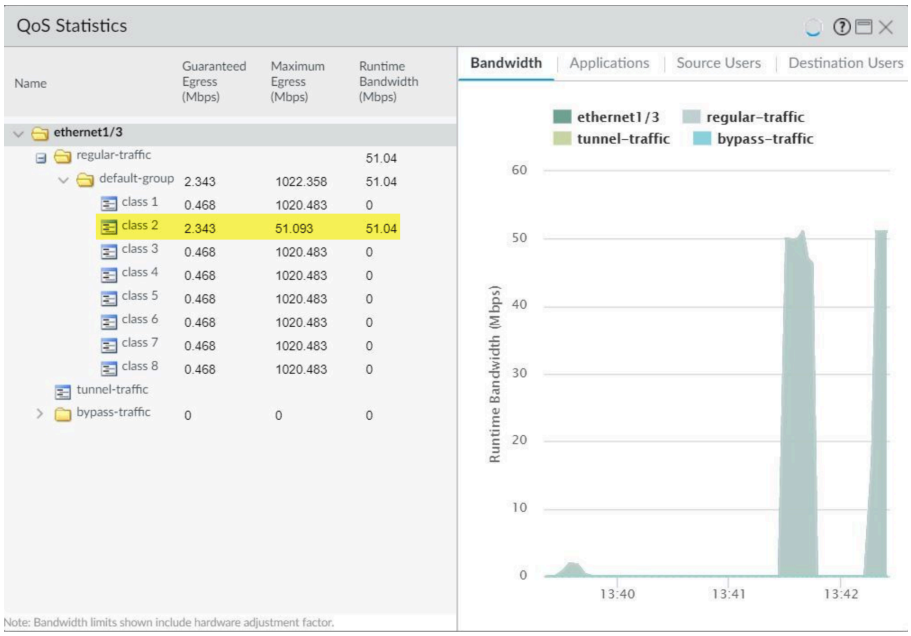
 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 7 | Verify a QoS configuration.

Select **Network** > **QoS** and then **Statistics** to view QoS bandwidth, active sessions of a selected QoS class, and active applications for the selected QoS class.

For example, see the statistics for ethernet 1/3 with QoS enabled:



Class 2 traffic limited to 2.343 Mbps of guaranteed bandwidth and a maximum bandwidth of 51.093 Mbps.

Continue to click the tabs to display further information regarding applications, source users, destination users, security rules and QoS rules.

*Bandwidth limits shown on the **QoS Statistics** window include a hardware adjustment factor.*

 ©2022 Palo Alto Networks, Inc.

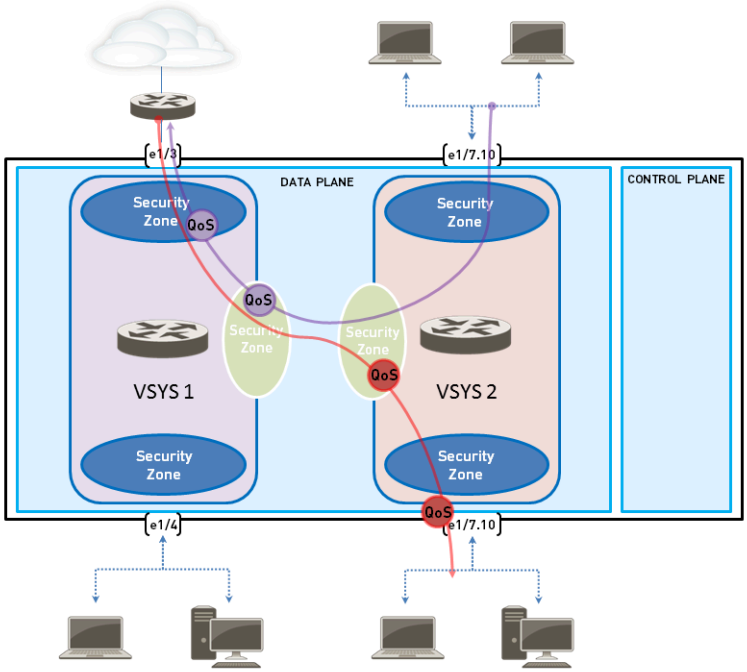
https://t.me/learningnets

Configure QoS for a Virtual System

QoS can be configured for a single or several virtual systems configured on a Palo Alto Networks firewall. Because a virtual system is an independent firewall, QoS must be configured independently for a single virtual system.

Configuring QoS for a virtual system is similar to configuring QoS on a physical firewall, with the exception that configuring QoS for a virtual system requires specifying the source and destination of traffic. Because a virtual system exists without set physical boundaries and because traffic in a virtual environment spans more than one virtual system, specifying source and destination zones and interfaces for traffic is necessary to control and shape traffic for a single virtual system.

The example below shows two virtual systems configured on firewall. VSYS 1 (purple) and VSYS 2 (red) each have QoS configured to prioritize or limit two distinct traffic flows, indicated by their corresponding purple (VSYS 1) and red (VSYS 2) lines. The QoS nodes indicate the points at traffic is matched to a QoS policy and assigned a QoS class of service, and then later indicate the point at which traffic is shaped as it egresses the firewall.



Refer to Virtual Systems for information on virtual systems and how to configure them.

STEP 1 | Confirm that the appropriate interfaces, virtual routers, and security zones are associated with each virtual system.

- To view configured interfaces, select **Network** > **Interface**.
- To view configured zones, select **Network** > **Zones**.
- To view information on defined virtual routers, select **Network** > **Virtual Routers**.

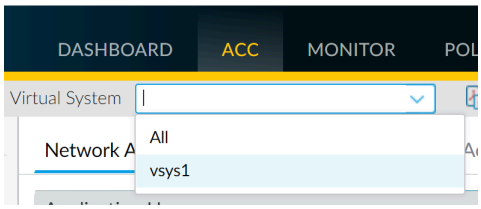
 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 2 | Identify traffic to apply QoS to.

Select **ACC** to view the **Application Command Center** page. Use the settings and charts on the **ACC** page to view trends and traffic related to Applications, URL filtering, Threat Prevention, Data Filtering, and HIP Matches.

To view information for a specific virtual system, select the virtual system from the **Virtual System** drop-down:



Click any application name to display detailed application information.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 3 | Identify the egress interface for applications that you identified as needing QoS treatment.

In a virtual system environment, QoS is applied to traffic on the traffic's egress point on the virtual system. Depending the configuration and QoS policy for a virtual system, the egress point of QoS traffic could be associated with a physical interface or could be a zone.

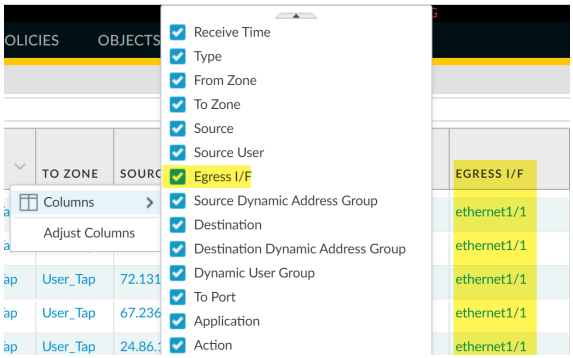
This example shows how to limit web-browsing traffic on vsys 1.

Select **Monitor** > **Logs** > **Traffic** to view traffic logs. Each entry has the option to display columns with information necessary to configure QoS in a virtual system environment:

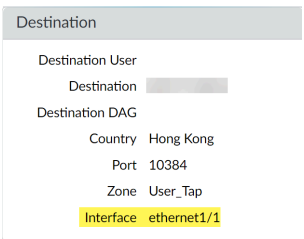
- virtual system
- egress interface
- ingress interface
- source zone
- destination zone

To display a column if it is not displayed by default:

- Click any column header to add a column to the log:



- Click the spyglass icon to the left of any entry to display a detailed log that includes the application's egress interface, as well as source and destination zones, in the **Source** and **Destination** sections:



For example, for web-browsing traffic from VSYS 1, the ingress interface is ethernet 1/2, the egress interface is ethernet 1/1, the source zone is trust and the destination zone is untrust.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 4 | Create a QoS Profile.

You can edit any existing QoS Profile, including the default, by clicking the profile name.

1. Select **Network** > **Network Profiles** > **QoS Profile** and click **Add** to open the QoS Profile dialog.
2. Enter a descriptive **Profile Name**.
3. Enter an **Egress Max** to set the overall bandwidth allocation for the QoS profile.
4. Enter an **Egress Guaranteed** to set the guaranteed bandwidth for the QoS profile.

   Any traffic that exceeds the QoS profile's egress guaranteed limit is best effort but is not guaranteed.
5. In the Classes section of the **QoS Profile**, specify how to treat up to eight individual QoS classes:
   1. Click **Add** to add a class to the QoS Profile.
   2. Select the **Priority** for the class.
   3. Enter an **Egress Max** for a class to set the overall bandwidth limit for that individual class.
   4. Enter an **Egress Guaranteed** for the class to set the guaranteed bandwidth for that individual class.
6. Click **OK** to save the QoS profile.

STEP 5 | Create a QoS policy.

In an environment with multiple virtual systems, traffic spans more than one virtual system. Because of this, when you are enabling QoS for a virtual system, you must define traffic to receive QoS treatment based on source and destination zones. This ensures that the traffic is

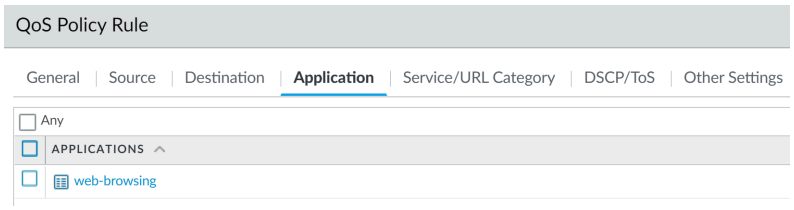
 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

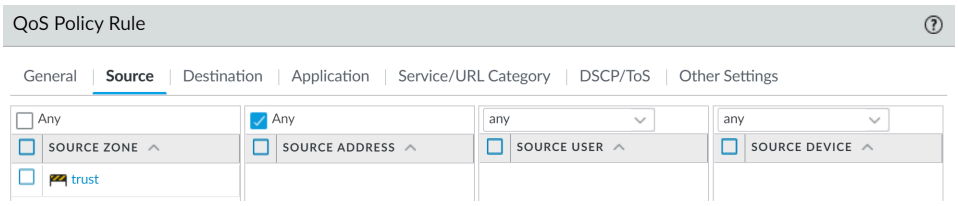
prioritized and shaped only for that virtual system (and not for other virtual systems through which the traffic might flow).

1. Select **Policies** > **QoS** and **Add** a QoS Policy Rule.
2. Select **General** and give the QoS Policy Rule a descriptive **Name**.
3. Specify the traffic to which the QoS policy rule will apply. Use the **Source**, **Destination**, **Application**, and **Service/URL Category** tabs to define matching parameters for identifying traffic.

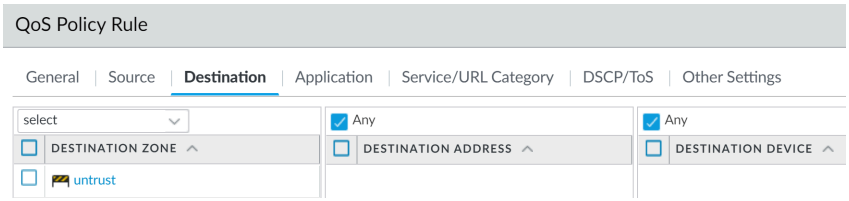
   For example, select **Application** and **Add** web-browsing to apply the QoS policy rule to that application:



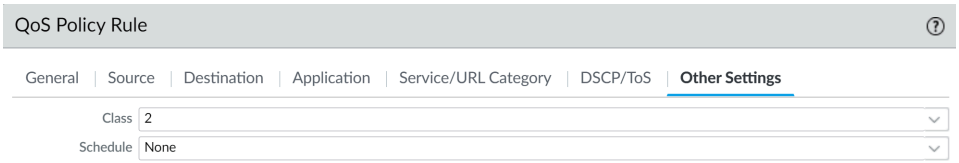
4. Select **Source** and **Add** the source zone of vsys 1 web-browsing traffic.



5. Select **Destination** and **Add** the destination zone of vsys 1 web-browsing traffic.



6. Select **Other Settings** and select a **QoS Class** to assign to the QoS policy rule. For example, assign Class 2 to web-browsing traffic on vsys 1:



7. Click **OK** to save the QoS policy rule.

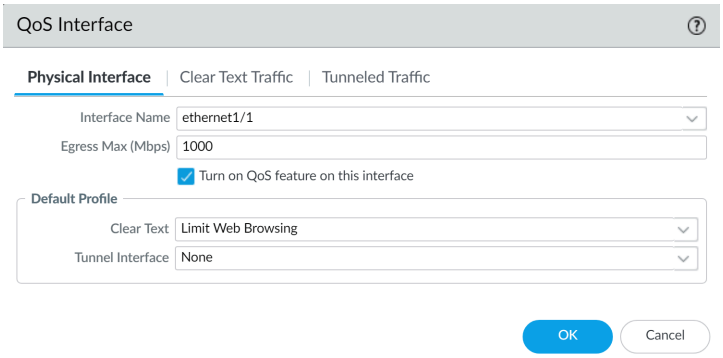
https://t.me/learningnets

STEP 6 | Enable the QoS Profile on a physical interface.

*It is a best practice to always define the **Egress Max** value for a QoS interface.*

1. Select **Network** > **QoS** and click **Add** to open the QoS Interface dialog.
2. Enable QoS on the physical interface:
   1. On the **Physical Interface** tab, select the **Interface Name** of the interface to apply the QoS Profile to.

      In this example, ethernet 1/1 is the egress interface for web-browsing traffic on vsys 1 (see Step 2).



   2. Select **Turn on QoS feature on this interface**.
3. On the **Physical Interface** tab, select the default QoS profile to apply to all **Clear Text** traffic.

   (Optional) Use the **Tunnel Interface** field to apply a default QoS profile to all tunneled traffic.
4. (Optional) On the **Clear Text Traffic** tab, configure additional QoS settings for clear text traffic:
   - Set the **Egress Guaranteed** and **Egress Max** bandwidths for clear text traffic.
   - Click **Add** to apply a QoS Profile to selected clear text traffic, further selecting the traffic for QoS treatment according to source interface and source subnet (creating a QoS node).
5. (Optional) On the **Tunneled Traffic** tab, configure additional QoS settings for tunnel interfaces:
   - Set the **Egress Guaranteed** and **Egress Max** bandwidths for tunneled traffic.
   - Click **Add** to associate a selected tunnel interface with a QoS Profile.
6. Click **OK** to save changes.
7. **Commit** the changes.

STEP 7 | Verify QoS configuration.

- Select **Network** > **QoS** to view the QoS Policies page. The **QoS Policies** page verifies that QoS is enabled and includes a **Statistics** link. Click the Statistics link to view QoS bandwidth,

https://t.me/learningnets

active sessions of a selected QoS node or class, and active applications for the selected QoS node or class.

- In a multi-vsys environment, sessions cannot span multiple systems. Multiple sessions are created for one traffic flow if the traffic passes through more than one virtual system. To browse sessions running on the firewall and view applied QoS Rules and QoS Classes, select **Monitor** > **Session Browser**.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Enforce QoS Based on DSCP Classification

A Differentiated Services Code Point (DSCP) is a packet header value that can be used to request (for example) high priority or best effort delivery for traffic. Session-Based DSCP Classification allows you to both honor DSCP values for incoming traffic and to mark a session with a DSCP value as session traffic exits the firewall. This enables all inbound and outbound traffic for a session can receive continuous QoS treatment as it flows through your network. For example, inbound return traffic from an external server can now be treated with the same QoS priority that the firewall initially enforced for the outbound flow based on the DSCP value the firewall detected at the beginning of the session. Network devices between the firewall and end user will also then enforce the same priority for the return traffic (and any other outbound or inbound traffic for the session).

You cannot apply DSCP code points or QoS to SSL Forward Proxy, SSL Inbound Inspection, and SSH Proxy traffic.

Different types of DSCP markings indicate different levels of service:

Completing this step enables the firewall to mark traffic with the same DSCP value that was detected at the beginning of a session (in this example, the firewall would mark return traffic with the DSCP AF11 value). While configuring QoS allows you to shape traffic as it egresses the firewall, enabling this option in a security rule allows the other network devices intermediate to the firewall and the client to continue to enforce priority for DSCP marked traffic.

- **Expedited Forwarding (EF)**: Can be used to request low loss, low latency and guaranteed bandwidth for traffic. Packets with EF codepoint values are typically guaranteed highest priority delivery.
- **Assured Forwarding (AF)**: Can be used to provide reliable delivery for applications. Packets with AF codepoint indicate a request for the traffic to receive higher priority treatment than best effort service provides (though packets with an EF codepoint will continue to take precedence over those with an AF codepoint).
- **Class Selector (CS)**: Can be used to provide backward compatibility with network devices that use the IP precedence field to mark priority traffic.
- **IP Precedence (ToS)**: Can be used by legacy network devices to mark priority traffic (the IP Precedence header field was used to indicate the priority for a packet before the introduction of the DSCP classification).
- **Custom Codepoint**: Create a custom codepoint to match to traffic by entering a **Codepoint Name** and **Binary Value**.

For example, select the **Assured Forwarding (AF)** to ensure traffic marked with an AF codepoint value has higher priority for reliable delivery over applications marked to receive lower priority.Use the following steps to enable Session-Based DSCP Classification. Start by configuring QoS based on DSCP marking detected at the beginning of a session. You can then continue to enable the firewall to mark the return flow for a session with the same DSCP value used to enforce QoS for the initial outbound flow.

STEP 1 | Perform the preliminary steps to Configure QoS.

https://t.me/learningnets

STEP 2 | Define the traffic to receive QoS treatment based on DSCP value.

1. Select **Policies** > **QoS** and **Add** or modify an existing QoS rule and populate required fields.
2. Select **DSCP/ToS** and select **Codepoints**.
3. **Add** DSCP/ToS codepoints for which you want to enforce QoS.
4. Select the **Type** of DSCP/ToS marking for the QoS rule to match to traffic:

   It is a best practice to use a single DSCP type to manage and prioritize your network traffic.
5. Match the QoS policy to traffic on a more granular scale by specifying the **Codepoint** value. For example, with Assured Forwarding (AF) selected as the **Type** of DSCP value for the policy to match, further specify an AF **Codepoint** value such as AF11.

   *When Expedited Forwarding (EF) is selected as the **Type** of DSCP marking, a granular **Codepoint** value cannot be specified. The QoS policy rule matches to traffic marked with any EF codepoint value.*
6. Select **Other Settings** and assign a **QoS Class** to traffic matched to the QoS rule. In this example, assign Class 1 to sessions where a DSCP marking of AF11 is detected for the first packet in the session.
7. Click **OK** to save the QoS rule.

STEP 3 | Define the QoS priority for traffic to receive when it is matched to a QoS rule based the DSCP marking detected at the beginning of a session.

1. Select **Network** > **Network Profiles** > **QoS Profile** and **Add** or modify an existing QoS profile. For details on profile options to set priority and bandwidth for traffic, see QoS Concepts and Configure QoS.
2. **Add** or modify a profile class. For example, because Step 2 showed steps to classify AF11 traffic as Class 1 traffic, you could add or modify a **class1** entry.
3. Select a **Priority** for the class of traffic, such as **high**.
4. Click **OK** to save the QoS Profile.

STEP 4 | Enable QoS on an interface.

Select **Network** > **QoS** and **Add** or modify an existing interface and **Turn on QoS feature on this interface**.

In this example, traffic with an AF11 DSCP marking is matched to the QoS rule and assigned Class 1. The QoS profile enabled on the interface enforces high priority treatment for Class 1 traffic as it egresses the firewall (the session *outbound* traffic).

https://t.me/learningnets

STEP 5 | Enable DSCP Marking.

Mark return traffic with a DSCP value, enabling the inbound flow for a session to be marked with the same DSCP value detected for the outbound flow.

1. Select **Policies** > **Security** and **Add** or modify a security policy.
2. Select **Actions** and in the **QoS Marking** drop-down, choose **Follow Client-to-Server Flow**.
3. Click **OK** to save your changes.

Completing this step enables the firewall to mark traffic with the same DSCP value that was detected at the beginning of a session (in this example, the firewall would mark return traffic with the DSCP AF11 value). While configuring QoS allows you to shape traffic as it egresses the firewall, enabling this option in a security rule allows the other network devices intermediate to the firewall and the client to continue to enforce priority for DSCP marked traffic.

STEP 6 | Commit the configuration.

Commit your changes.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

QoS Use Cases

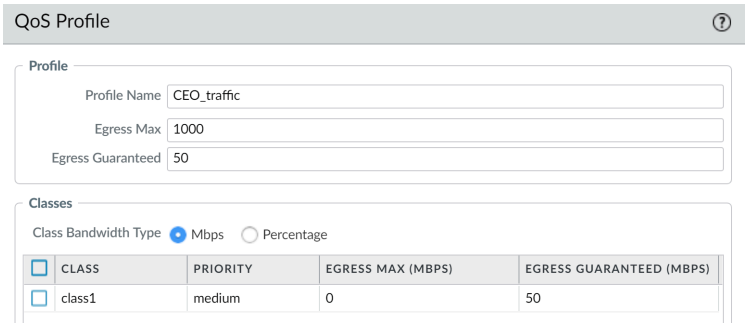
The following use cases demonstrate how to use QoS in common scenarios:

- Use Case: QoS for a Single User
- Use Case: QoS for Voice and Video Applications

Use Case: QoS for a Single User

A CEO finds that during periods of high network usage, she is unable to access enterprise applications to respond effectively to critical business communications. The IT admin wants to ensure that all traffic to and from the CEO receives preferential treatment over other employee traffic so that she is guaranteed not only access to, but high performance of, critical network resources.

STEP 1 | The admin creates the QoS profile *CEO_traffic* to define how traffic originating from the CEO will be treated and shaped as it flows out of the company network:



The admin assigns a guaranteed bandwidth (**Egress Guaranteed**) of 50 Mbps to ensure that the CEO will have that amount that bandwidth guaranteed to her at all times (more than she would need to use), regardless of network congestion.

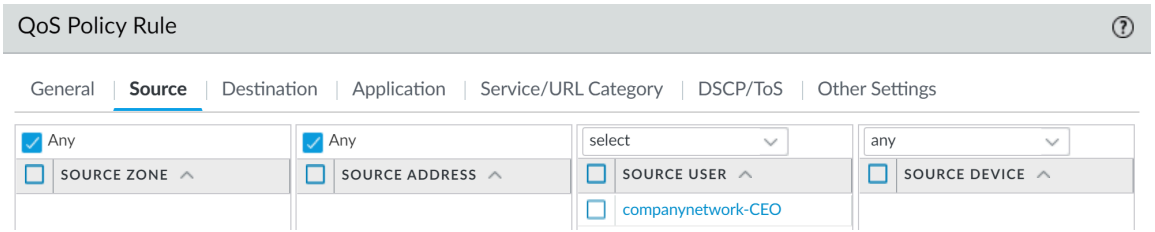
The admin continues by designating Class 1 traffic as high priority and sets the profile's maximum bandwidth usage (**Egress Max**) to 1000 Mbps, the same maximum bandwidth for the interface that the admin will enable QoS on. The admin is choosing to not restrict the CEO's bandwidth usage in any way.

*It is a best practice to populate the **Egress Max** field for a QoS profile, even if the max bandwidth of the profile matches the max bandwidth of the interface. The QoS profile's max bandwidth should never exceed the max bandwidth of the interface you are planning to enable QoS on.*

STEP 2 | The admin creates a QoS policy to identify the CEO's traffic (**Policies** > **QoS**) and assigns it the class that he defined in the QoS profile (see prior step). Because User-ID is configured, the admin uses the **Source** tab in the QoS policy to singularly identify the CEO's traffic by her

https://t.me/learningnets

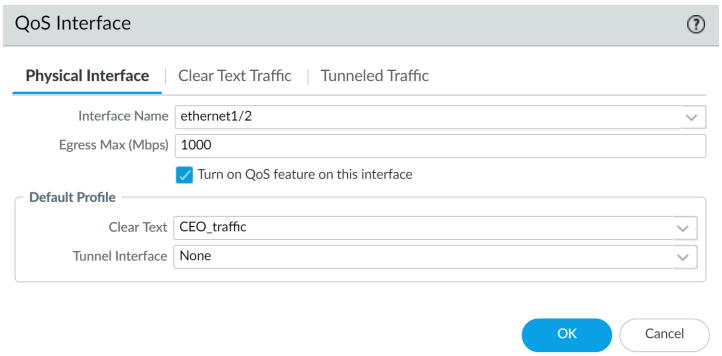
company network username. (If User-ID is not configured, the administrator could **Add** the CEO's IP address under **Source Address**. See User-ID.):



The admin associates the CEO's traffic with Class 1 (**Other Settings** tab) and then continues to populate the remaining required policy fields; the admin gives the policy a descriptive **Name** (**General** tab) and selects **Any** for the **Source Zone** (**Source** tab) and **Destination Zone** (**Destination** tab):

	NAME	TAGS	Source ZONE	Source ADDRESS	Source USER	Source DEVICE	Destination ZONE	Destination ADDRESS	Destination DEVICE	APPLICATION	SERVICE	DSCP/TOS	CLASS
1	HTTPS	none	trust	any	any	any	untrust	any	any	web-browsing	any	any	2
2	Voice-Video	none	any	any	any	any	any	any	any	voip-video-l...	any	any	1
3	Guarantee CEO bandwidth	none	any	any	companynet...	any	any	any	any	any	any	any	1

STEP 3 | Now that Class 1 is associated with the CEO's traffic, the admin enables QoS by checking **Turn on QoS feature on interface** and selecting the traffic flow's egress interface. The egress interface for the CEO's traffic flow is the external-facing interface, in this case, ethernet 1/2:



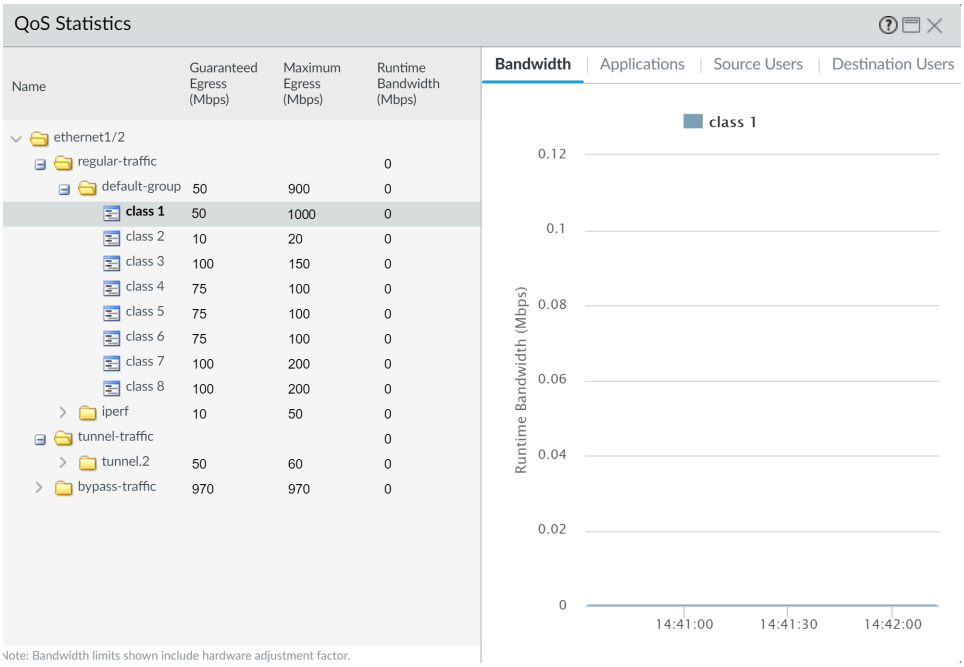
Because the admin wants to ensure that all traffic originating from the CEO is guaranteed by the QoS profile and associated QoS policy he created, he selects the *CEO_traffic* to apply to **Clear Text** traffic flowing from ethernet 1/2.

STEP 4 | After committing the QoS configuration, the admin navigates to the **Network** > **QoS** page to confirm that the QoS profile CEO_traffic is enabled on the external-facing interface, ethernet 1/2:

NAME	GUARANTEED EGRESS (MBPS)	MAXIMUM EGRESS (MBPS)	PROFILE	ENABLED	
ethernet1/2		1,000.000		✓	Statistics
Tunneled Traffic					
Clear Text Traffic	50.000		CEO_traffic		

https://t.me/learningnets

STEP 5 | He clicks **Statistics** to view how traffic originating with the CEO (Class 1) is being shaped as it flows from ethernet 1/2:



*This case demonstrates how to apply QoS to traffic originating from a single source user. However, if you also wanted to guarantee or shape traffic to a destination user, you could configure a similar QoS setup. Instead of, or in addition to this work flow, create a QoS policy that specifies the user's IP address as the **Destination Address** on the **Policies** > **QoS** page (instead of specifying the user's source information) and then enable QoS on the network's internal-facing interface on the **Network** > **QoS** page (instead of the external-facing interface).*

Use Case: QoS for Voice and Video Applications

Voice and video traffic is particularly sensitive to measurements that the QoS feature shapes and controls, especially latency and jitter. For voice and video transmissions to be audible and clear, voice and video packets cannot be dropped, delayed, or delivered inconsistently. A best practice for voice and video applications, in addition to guaranteeing bandwidth, is to guarantee priority to voice and video traffic.

In this example, employees at a company branch office are experiencing difficulties and unreliability in using video conferencing and Voice over IP (VoIP) technologies to conduct business communications with other branch offices, with partners, and with customers. An IT admin intends to implement QoS in order to address these issues and ensure effective and reliable business communication for the branch employees. Because the admin wants to guarantee QoS to both incoming and outgoing network traffic, he will enable QoS on both the firewall's internal- and external-facing interfaces.

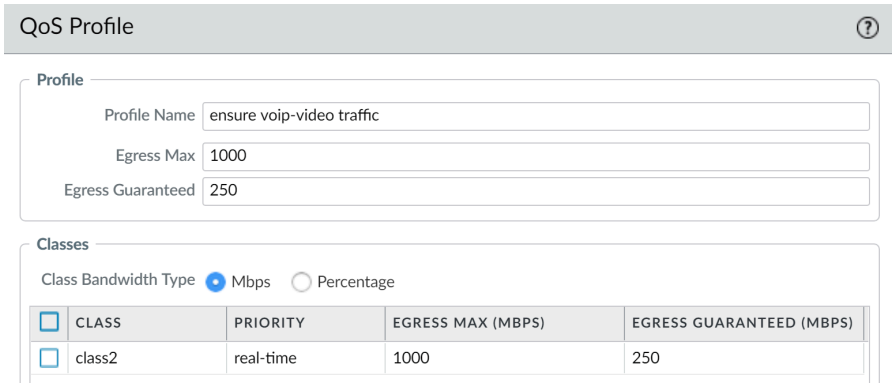
 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 1 | The admin creates a QoS profile, defining Class 2 so that Class 2 traffic receives real-time priority and on an interface with a maximum bandwidth of 1000 Mbps, is guaranteed a bandwidth of 250 Mbps at all times, including peak periods of network usage.

Real-time priority is typically recommended for applications affected by latency, and is particularly useful in guaranteeing performance and quality of voice and video applications.

On the firewall web interface, the admin selects **Network** > **Network Profiles** > **Qos Profile** page, clicks **Add**, enters the **Profile Name** ensure voip-video traffic and defines Class 2 traffic.

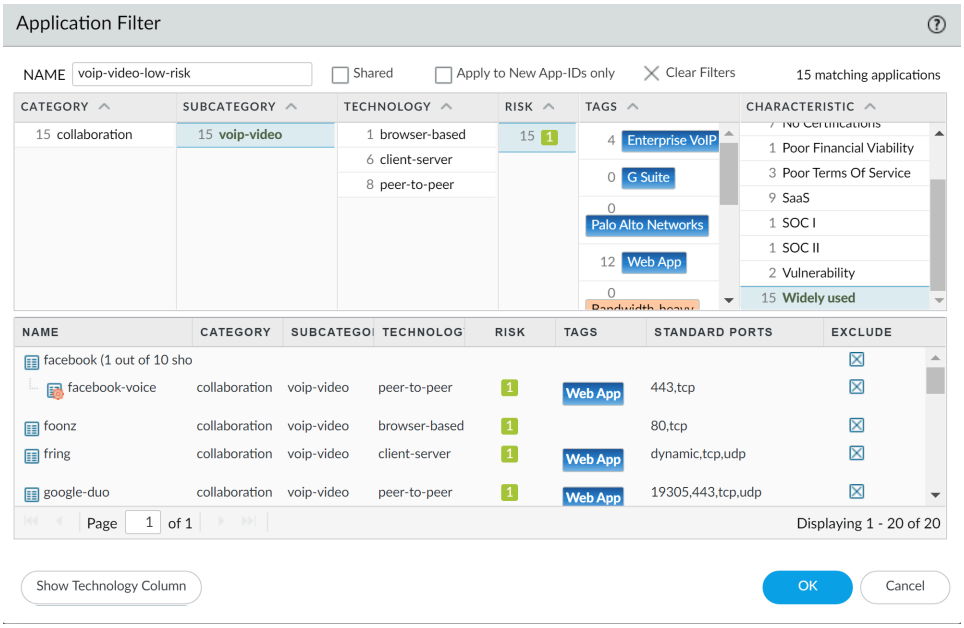


STEP 2 | The admin creates a QoS policy to identify voice and video traffic. Because the company does not have one standard voice and video application, the admin wants to ensure QoS is applied to a few applications that are widely and regularly used by employees to communicate with other offices, with partners, and with customers. On the **Policies** > **QoS** > **QoS Policy Rule** > **Applications** tab, the admin clicks **Add** and opens the **Application Filter** window. The admin continues by selecting criteria to filter the applications he wants to apply

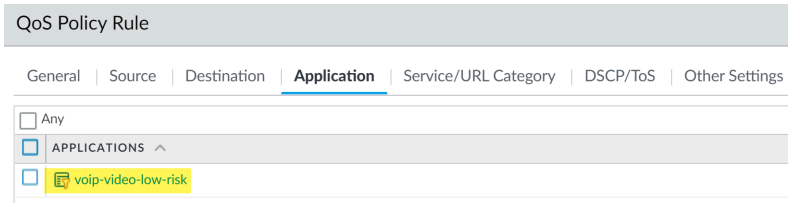
https://t.me/learningnets

QoS to, choosing the Subcategory voip-video, and narrowing that down by specifying only voip-video applications that are both low-risk and widely-used.

The application filter is a dynamic tool that, when used to filter applications in the QoS policy, allows QoS to be applied to all applications that meet the criteria of voip-video, low risk, and widely used at any given time.



The admin names the **Application Filter** voip-video-low-risk and includes it in the QoS policy:



The admin names the QoS policy Voice-Video and selects Other Settings to assign all traffic matched to the policy Class 2. He is going to use the Voice-Video QoS policy for both incoming and outgoing QoS traffic, so he sets **Source** and **Destination** information to **Any**:

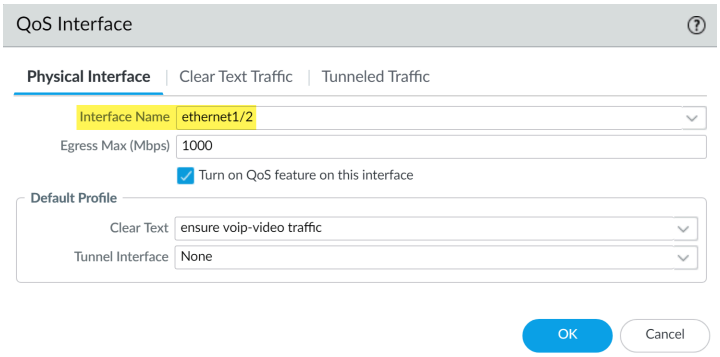
			Source				Destination						
	NAME	TAGS	ZONE	ADDRESS	USER	DEVICE	ZONE	ADDRESS	DEVICE	APPLICATION	SERVICE	DSCP/TOS	CLASS
1	HTTPS	none	trust	any	any	any	untrust	any	any	web-browsing	any	any	2
2	Voice-Video	none	any	any	any	any	any	any	any	voip-video-l...	any	any	1

STEP 3 | Because the admin wants to ensure QoS for both incoming and outgoing voice and video communications, he enables QoS on the network's external-facing interface (to apply QoS

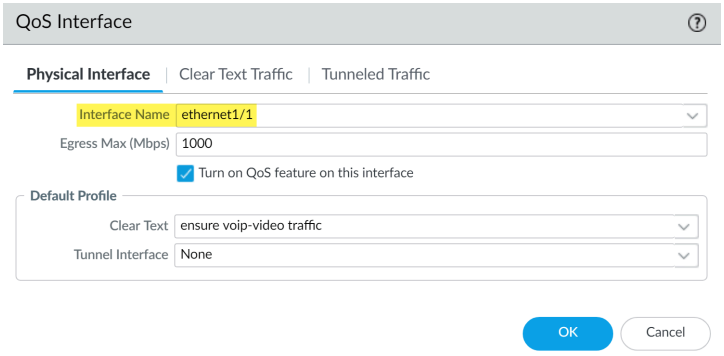
https://t.me/learningnets

to outgoing communications) and to the internal-facing interface (to apply QoS to incoming communications).

The admin begins by enabling the QoS profile he created, ensure voice-video traffic (Class 2 in this profile is associated with policy, Voice-Video) on the external-facing interface, in this case, ethernet 1/2.



He then enables the same QoS profile ensure voip-video traffic on a second interface, the internal-facing interface (in this case, ethernet 1/1).



STEP 4 | The admin selects **Network** > **QoS** to confirm that QoS is enabled for both incoming and outgoing voice and video traffic:

NAME	GUARANTEED EGRESS (MBPS)	MAXIMUM EGRESS (MBPS)	PROFILE	ENABLED	
ethernet1/1		1,000.000		☑	Statistics
Tunneled Traffic					
Clear Text Traffic	250.000		ensure voip-video traffic		
ethernet1/2		1,000.000		☑	Statistics
Tunneled Traffic					
Clear Text Traffic	250.000		ensure voip-video traffic		

The admin has successfully enabled QoS on both the network's internal- and external-facing interfaces. Real-time priority is now ensured for voice and video application traffic as it flows both into and out of the network, ensuring that these communications, which are particularly sensitive to latency and jitter, can be used reliably and effectively to perform both internal and external business communications.

https://t.me/learningnets

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

VPNs

Virtual private networks (VPNs) create tunnels that allow users/systems to connect securely over a public network, as if they were connecting over a local area network (LAN). To set up a VPN tunnel, you need a pair of devices that can authenticate each other and encrypt the flow of information between them. The devices can be a pair of Palo Alto Networks firewalls, or a Palo Alto Networks firewall along with a VPN-capable device from another vendor.

> VPN Deployments
> Site-to-Site VPN Overview
> Site-to-Site VPN Concepts
> Set Up Site-to-Site VPN
> Site-to-Site VPN Quick Configs

https://t.me/learningnets

VPN Deployments

The Palo Alto Networks firewall supports the following VPN deployments:

- **Site-to-Site VPN**— A simple VPN that connects a central site and a remote site, or a hub and spoke VPN that connects a central site with multiple remote sites. The firewall uses the IP Security (IPSec) set of protocols to set up a secure tunnel for the traffic between the two sites. See Site-to-Site VPN Overview.
- **Remote User-to-Site VPN**—A solution that uses the GlobalProtect agent to allow a remote user to establish a secure connection through the firewall. This solution uses SSL and IPSec to establish a secure connection between the user and the site. Refer to the GlobalProtect Administrator's Guide.
- **Large Scale VPN**— The Palo Alto Networks GlobalProtect Large Scale VPN (LSVPN) provides a simplified mechanism to roll out a scalable hub and spoke VPN with up to 1,024 satellite offices. The solution requires Palo Alto Networks firewalls to be deployed at the hub and at every spoke. It uses certificates for device authentication, SSL for securing communication between all components, and IPSec to secure data. See Large Scale VPN (LSVPN).

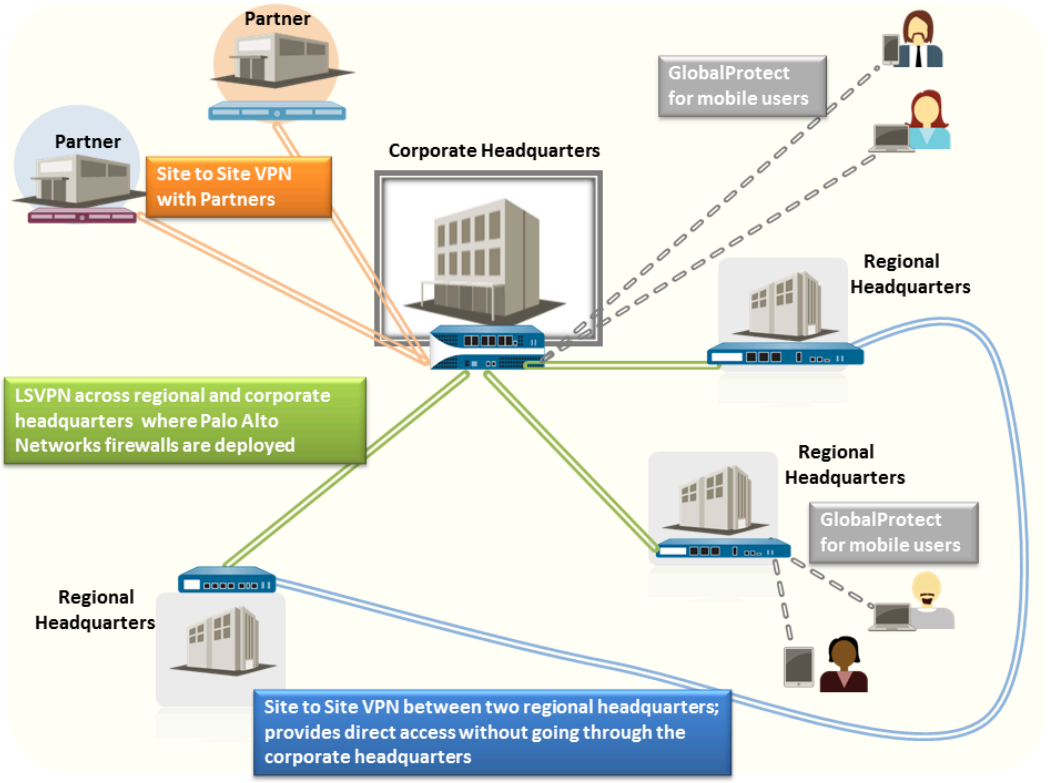


Figure 7: VPN Deployments

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Site-to-Site VPN Overview

A VPN connection that allows you to connect two Local Area Networks (LANs) is called a site-to-site VPN. You can configure route-based VPNs to connect Palo Alto Networks firewalls located at two sites or to connect a Palo Alto Networks firewall with a third-party security device at another location. The firewall can also interoperate with third-party policy-based VPN devices; the Palo Alto Networks firewall supports route-based VPN.

The Palo Alto Networks firewall sets up a route-based VPN, where the firewall makes a routing decision based on the destination IP address. If traffic is routed to a specific destination through a VPN tunnel, then it is handled as VPN traffic.

The IP Security (IPSec) set of protocols is used to set up a secure tunnel for the VPN traffic, and the information in the TCP/IP packet is secured (and encrypted if the tunnel type is ESP). The IP packet (header and payload) is embedded in another IP payload, and a new header is applied and then sent through the IPSec tunnel. The source IP address in the new header is that of the local VPN peer and the destination IP address is that of the VPN peer on the far end of the tunnel. When the packet reaches the remote VPN peer (the firewall at the far end of the tunnel), the outer header is removed and the original packet is sent to its destination.

In order to set up the VPN tunnel, first the peers need to be authenticated. After successful authentication, the peers negotiate the encryption mechanism and algorithms to secure the communication. The Internet Key Exchange (IKE) process is used to authenticate the VPN peers, and IPSec Security Associations (SAs) are defined at each end of the tunnel to secure the VPN communication. IKE uses digital certificates or preshared keys, and the Diffie Hellman keys to set up the SAs for the IPSec tunnel. The SAs specify all of the parameters that are required for secure transmission— including the security parameter index (SPI), security protocol, cryptographic keys, and the destination IP address—encryption, data authentication, data integrity, and endpoint authentication.

The following figure shows a VPN tunnel between two sites. When a client that is secured by VPN Peer A needs content from a server located at the other site, VPN Peer A initiates a connection request to VPN Peer B. If the security policy permits the connection, VPN Peer A uses the IKE Crypto profile parameters (IKE phase 1) to establish a secure connection and authenticate VPN Peer B. Then, VPN Peer A establishes the VPN tunnel using the IPSec Crypto profile, which defines the IKE phase 2 parameters to allow the secure transfer of data between the two sites.



Figure 8: Site-to-Site VPN

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Site-to-Site VPN Concepts

A VPN connection provides secure access to information between two or more sites. In order to provide secure access to resources and reliable connectivity, a VPN connection needs the following components:

- IKE Gateway
- Tunnel Interface
- Tunnel Monitoring
- Internet Key Exchange (IKE) for VPN
- IKEv2

IKE Gateway

The Palo Alto Networks firewalls or a firewall and another security device that initiate and terminate VPN connections across the two networks are called the IKE Gateways. To set up the VPN tunnel and send traffic between the IKE Gateways, each peer must have an IP address—static or dynamic—or FQDN. The VPN peers use preshared keys or certificates to mutually authenticate each other.

The peers must also negotiate the mode—main or aggressive—for setting up the VPN tunnel and the SA lifetime in IKE Phase 1. Main mode protects the identity of the peers and is more secure because more packets are exchanged when setting up the tunnel. Main mode is the recommended mode for IKE negotiation if both peers support it. Aggressive mode uses fewer packets to set up the VPN tunnel and is hence faster but a less secure option for setting up the VPN tunnel.

See Set Up an IKE Gateway for configuration details.

Tunnel Interface

To set up a VPN tunnel, the Layer 3 interface at each end must have a logical *tunnel* interface for the firewall to connect to and establish a VPN tunnel. A tunnel interface is a logical (virtual) interface that is used to deliver traffic between two endpoints. If you configure any proxy IDs, the proxy ID is counted toward any IPSec tunnel capacity.

The tunnel interface must belong to a security zone to apply policy and it must be assigned to a virtual router in order to use the existing routing infrastructure. Ensure that the tunnel interface and the physical interface are assigned to the same virtual router so that the firewall can perform a route lookup and determine the appropriate tunnel to use.

Typically, the Layer 3 interface that the tunnel interface is attached to belongs to an external zone, for example the untrust zone. While the tunnel interface can be in the same security zone as the physical interface, for added security and better visibility, you can create a separate zone for the tunnel interface. If you create a separate zone for the tunnel interface, say a VPN zone, you will need to create security policies to enable traffic to flow between the VPN zone and the trust zone.

To route traffic between the sites, a tunnel interface does not require an IP address. An IP address is only required if you want to enable tunnel monitoring or if you are using a dynamic routing

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

protocol to route traffic across the tunnel. With dynamic routing, the tunnel IP address serves as the next hop IP address for routing traffic to the VPN tunnel.

If you are configuring the Palo Alto Networks firewall with a VPN peer that performs policy-based VPN, you must configure a local and remote Proxy ID when setting up the IPSec tunnel. Each peer compares the Proxy-IDs configured on it with what is actually received in the packet in order to allow a successful IKE phase 2 negotiation. If multiple tunnels are required, configure unique Proxy IDs for each tunnel interface; a tunnel interface can have a maximum of 250 Proxy IDs. Each Proxy ID counts towards the IPSec VPN tunnel capacity of the firewall, and the tunnel capacity varies by the firewall model.

See Set Up an IPSec Tunnel for configuration details.

Tunnel Monitoring

For a VPN tunnel, you can check connectivity to a destination IP address across the tunnel. The network monitoring profile on the firewall allows you to verify connectivity (using ICMP) to a destination IP address or a next hop at a specified polling interval, and to specify an action on failure to access the monitored IP address.

If the destination IP is unreachable, you either configure the firewall to wait for the tunnel to recover or configure automatic failover to another tunnel. In either case, the firewall generates a system log that alerts you to a tunnel failure and renegotiates the IPSec keys to accelerate recovery.

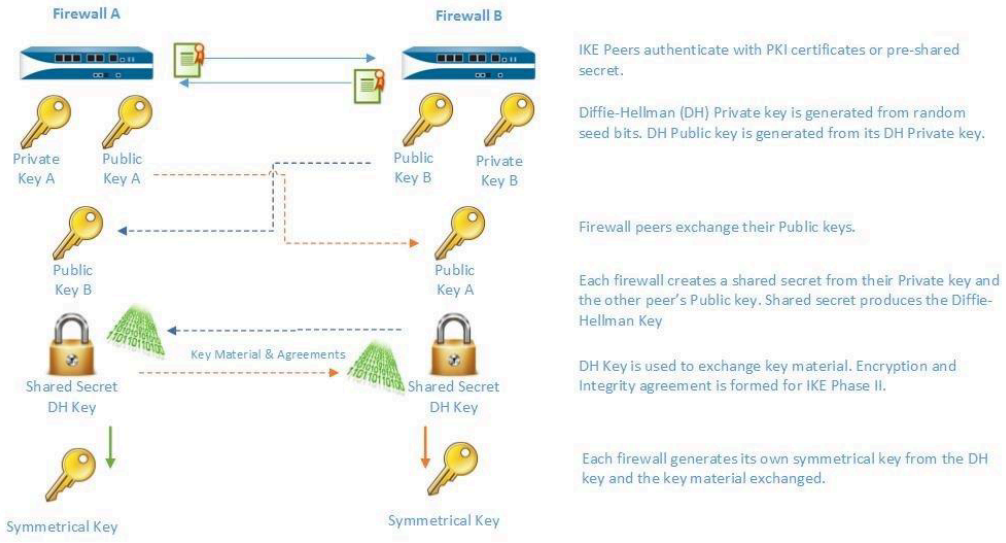
See Set Up Tunnel Monitoring for configuration details.

Internet Key Exchange (IKE) for VPN

The IKE process allows the VPN peers at both ends of the tunnel to encrypt and decrypt packets using mutually agreed-upon keys or certificate and method of encryption. The IKE process occurs in two phases: IKE Phase 1 and IKE Phase 2. Each of these phases use keys and encryption algorithms that are defined using cryptographic profiles— IKE crypto profile and IPSec crypto profile—and the result of the IKE negotiation is a Security Association (SA). An SA is a set of mutually agreed-upon keys and algorithms that are used by both VPN peers to allow the flow of data across the VPN tunnel. The following illustration depicts the key exchange process for setting up the VPN tunnel:

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets



IKE Phase 1

In this phase, the firewalls use the parameters defined in the IKE Gateway configuration and the IKE Crypto profile to authenticate each other and set up a secure control channel. IKE Phase supports the use of preshared keys or digital certificates (which use public key infrastructure, PKI) for mutual authentication of the VPN peers. Preshared keys are a simple solution for securing smaller networks because they do not require the support of a PKI infrastructure. Digital certificates can be more convenient for larger networks or implementations that require stronger authentication security.

When using certificates, make sure that the CA issuing the certificate is trusted by both gateway peers and that the maximum length of certificates in the certificate chain is 5 or less. With IKE fragmentation enabled, the firewall can reassemble IKE messages with up to 5 certificates in the certificate chain and successfully establish a VPN tunnel.

The IKE Crypto profile defines the following options that are used in the IKE SA negotiation:

- Diffie-Hellman (DH) group for generating symmetrical keys for IKE.

  The Diffie-Hellman algorithm uses the private key of one party and the public key of the other to create a shared secret, which is an encrypted key that both VPN tunnel peers share. The DH groups supported on the firewall are: Group 1—768 bits, Group 2—1024 bits (default), Group 5—1536 bits, Group 14—2048 bits, Group 19—256-bit elliptic curve group, and Group 20—384-bit elliptic curve group.
- Authentication algorithms—sha1, sha 256, sha 384, sha 512, or md5
- Encryption algorithms—aes-256-gcm, aes-128-gcm, 3des, aes-128-cbc, aes-192-cbc, aes-256-cbc, or des

IKE Phase 2

After the tunnel is secured and authenticated, in Phase 2 the channel is further secured for the transfer of data between the networks. IKE Phase 2 uses the keys that were established in Phase 1 of the process and the IPSec Crypto profile, which defines the IPSec protocols and keys used for the SA in IKE Phase 2.

The IPSEC uses the following protocols to enable secure communication:

https://t.me/learningnets

- Encapsulating Security Payload (ESP)—Allows you to encrypt the entire IP packet, and authenticate the source and verify integrity of the data. While ESP requires that you encrypt and authenticate the packet, you can choose to only encrypt or only authenticate by setting the encryption option to Null; using encryption without authentication is discouraged.
- Authentication Header (AH)—Authenticates the source of the packet and verifies data integrity. AH does not encrypt the data payload and is unsuited for deployments where data privacy is important. AH is commonly used when the main concern is to verify the legitimacy of the peer, and data privacy is not required.

Table 5: Algorithms Supported for IPSEC Authentication and Encryption

ESP	AH
Diffie Hellman (DH) exchange options supported	
• Group 1—768 bits • Group 2—1024 bits (the default) • Group 5—1536 bits • Group 14—2048 bits. • Group 19— 256-bit elliptic curve group • Group 20—384-bit elliptic curve group • no-pfs—By default, perfect forward secrecy (PFS) is enabled, which means a new DH key is generated in IKE phase 2 using one of the groups listed above. This key is independent of the keys exchanged in IKE phase1 and provides better data transfer security. If you select no-pfs, the DH key created at phase 1 is not renewed and a single key is used for the IPSec SA negotiations. Both VPN peers must be enabled or disabled for PFS.	
Encryption algorithms supported	
• 3des	Triple Data Encryption Standard (3DES) with a security strength of 112 bits
• aes-128-cbc	Advanced Encryption Standard (AES) using cipher block chaining (CBC) with a security strength of 128 bits
• aes-192-cbc	AES using CBC with a security strength of 192 bits
• aes-256-cbc	AES using CBC with a security strength of 256 bits
• aes-128-ccm	AES using Counter with CBC-MAC (CCM) with a security strength of 128 bits
• aes-128-gcm	AES using Galois/Counter Mode (GCM) with a security strength of 128 bits
• aes-256-gcm	AES using GCM with a security strength of 256 bits

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

ESP	AH
• des	Data Encryption Standard (DES) with a security strength of 56 bits
Authentication algorithms supported	
• md5	• md5
• sha 1	• sha 1
• sha 256	• sha 256
• sha 384	• sha 384
• sha512	• sha 512

Methods of Securing IPSec VPN Tunnels (IKE Phase 2)

IPSec VPN tunnels can be secured using manual keys or auto keys. In addition, IPSec configuration options include Diffie-Hellman Group for key agreement, and/or an encryption algorithm and a hash for message authentication.

- **Manual Key**—Manual key is typically used if the Palo Alto Networks firewall is establishing a VPN tunnel with a legacy device, or if you want to reduce the overhead of generating session keys. If using manual keys, the same key must be configured on both peers.

  Manual keys are not recommended for establishing a VPN tunnel because the session keys can be compromised when relaying the key information between the peers; if the keys are compromised, the data transfer is no longer secure.
- **Auto Key**— Auto Key allows you to automatically generate keys for setting up and maintaining the IPSec tunnel based on the algorithms defined in the IPSec Crypto profile.

IKEv2

An IPSec VPN gateway uses IKEv1 or IKEv2 to negotiate the IKE security association (SA) and IPSec tunnel. IKEv2 is defined in RFC 5996.

Unlike IKEv1, which uses Phase 1 SA and Phase 2 SA, IKEv2 uses a child SA for Encapsulating Security Payload (ESP) or Authentication Header (AH), which is set up with an IKE SA.

NAT traversal (NAT-T) must be enabled on both gateways if you have NAT occurring on a device that sits between the two gateways. A gateway can see only the public (globally routable) IP address of the NAT device.

IKEv2 provides the following benefits over IKEv1:

- Tunnel endpoints exchange fewer messages to establish a tunnel. IKEv2 uses four messages; IKEv1 uses either nine messages (in main mode) or six messages (in aggressive mode).
- Built-in NAT-T functionality improves compatibility between vendors.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

- Built-in health check automatically re-establishes a tunnel if it goes down. The liveness check replaces the Dead Peer Detection used in IKEv1.
- Supports traffic selectors (one per exchange). The traffic selectors are used in IKE negotiations to control what traffic can access the tunnel.
- Supports Hash and URL certificate exchange to reduce fragmentation.
- Resiliency against DoS attacks with improved peer validation. An excessive number of half-open SAs can trigger cookie validation.

Before configuring IKEv2, you should be familiar with the following concepts:

- Liveness Check
- Cookie Activation Threshold and Strict Cookie Validation
- Traffic Selectors
- Hash and URL Certificate Exchange
- SA Key Lifetime and Re-Authentication Interval

After you Set Up an IKE Gateway, if you chose IKEv2, perform the following optional tasks related to IKEv2 as required by your environment:

- Export a Certificate for a Peer to Access Using Hash and URL
- Import a Certificate for IKEv2 Gateway Authentication
- Change the Key Lifetime or Authentication Interval for IKEv2
- Change the Cookie Activation Threshold for IKEv2
- Configure IKEv2 Traffic Selectors

Liveness Check

The liveness check for IKEv2 is similar to Dead Peer Detection (DPD), which IKEv1 uses as the way to determine whether a peer is still available.

In IKEv2, the liveness check is achieved by any IKEv2 packet transmission or an empty informational message that the gateway sends to the peer at a configurable interval, five seconds by default. If necessary, the sender attempts the retransmission up to ten times. If it doesn't get a response, the sender closes and deletes the IKE_SA and corresponding CHILD_SAs. The sender will start over by sending out another IKE_SA_INIT message.

Cookie Activation Threshold and Strict Cookie Validation

Cookie validation is always enabled for IKEv2; it helps protect against half-SA DoS attacks. You can configure the global threshold number of half-open SAs that will trigger cookie validation. You can also configure individual IKE gateways to enforce cookie validation for every new IKEv2 SA.

- The **Cookie Activation Threshold** is a global VPN session setting that limits the number of simultaneous half-opened IKE SAs (default is 500). When the number of half-opened IKE SAs exceeds the **Cookie Activation Threshold**, the Responder will request a cookie, and the Initiator must respond with an IKE_SA_INIT containing a cookie to validate the connection. If the cookie

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

validation is successful, another SA can be initiated. A value of 0 means that cookie validation is always on.

The Responder does not maintain a state of the Initiator, nor does it perform a Diffie-Hellman key exchange, until the Initiator returns the cookie. IKEv2 cookie validation mitigates a DoS attack that would try to leave numerous connections half open.

The **Cookie Activation Threshold** must be lower than the **Maximum Half Opened SA** setting. If you Change the Cookie Activation Threshold for IKEv2 to a very high number (for example, 65534) and the **Maximum Half Opened SA** setting remained at the default value of 65535, cookie validation is essentially disabled.

- You can enable **Strict Cookie Validation** if you want cookie validation performed for every new IKEv2 SA a gateway receives, regardless of the global threshold. **Strict Cookie Validation** affects only the IKE gateway being configured and is disabled by default. With **Strict Cookie Validation** disabled, the system uses the **Cookie Activation Threshold** to determine whether a cookie is needed or not.

Traffic Selectors

In IKEv1, a firewall that has a route-based VPN needs to use a local and remote Proxy ID in order to set up an IPSec tunnel. Each peer compares its Proxy IDs with what it received in the packet in order to successfully negotiate IKE Phase 2. IKE Phase 2 is about negotiating the SAs to set up an IPSec tunnel. (For more information on Proxy IDs, see Tunnel Interface.)

In IKEv2, you can Configure IKEv2 Traffic Selectors, which are components of network traffic that are used during IKE negotiation. Traffic selectors are used during the CHILD_SA (tunnel creation) Phase 2 to set up the tunnel and to determine what traffic is allowed through the tunnel. The two IKE gateway peers must negotiate and agree on their traffic selectors; otherwise, one side narrows its address range to reach agreement. One IKE connection can have multiple tunnels; for example, you can assign different tunnels to each department to isolate their traffic. Separation of traffic also allows features such as QoS to be implemented.

The IPv4 and IPv6 traffic selectors are:

- **Source IP address**—A network prefix, address range, specific host, or wildcard.
- **Destination IP address**—A network prefix, address range, specific host, or wildcard.
- **Protocol**—A transport protocol, such as TCP or UDP.
- **Source port**—The port where the packet originated.
- **Destination port**—The port the packet is destined for.

During IKE negotiation, there can be multiple traffic selectors for different networks and protocols. For example, the Initiator might indicate that it wants to send TCP packets from 172.168.0.0/16 through the tunnel to its peer, destined for 198.5.0.0/16. It also wants to send UDP packets from 172.17.0.0/16 through the same tunnel to the same gateway, destined for 0.0.0.0 (any network). The peer gateway must agree to these traffic selectors so that it knows what to expect.

It is possible that one gateway will start negotiation using a traffic selector that is a more specific IP address than the IP address of the other gateway.

- For example, gateway A offers a source IP address of 172.16.0.0/16 and a destination IP address of 192.16.0.0/16. But gateway B is configured with 0.0.0.0 (any source) as the source IP address and 0.0.0.0 (any destination) as the destination IP address. Therefore,

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

gateway B narrows down its source IP address to 192.16.0.0/16 and its destination address to 172.16.0.0/16. Thus, the narrowing down accommodates the addresses of gateway A and the traffic selectors of the two gateways are in agreement.

- If gateway B (configured with source IP address 0.0.0.0) is the Initiator instead of the Responder, gateway A will respond with its more specific IP addresses, and gateway B will narrow down its addresses to reach agreement.

Hash and URL Certificate Exchange

IKEv2 supports Hash and URL Certificate Exchange, which is used during an IKEv2 negotiation of an SA. You store the certificate on an HTTP server, which is specified by a URL. The peer fetches the certificate from the server based on receiving the URL to the server. The hash is used to check whether the content of the certificate is valid or not. Thus, the two peers exchange certificates with the HTTP CA rather than with each other.

The hash part of Hash and URL reduces the message size and thus Hash and URL is a way to reduce the likelihood of packet fragmentation during IKE negotiation. The peer receives the certificate and hash that it expects, and thus IKE Phase 1 has validated the peer. Reducing fragmentation occurrences helps protect against DoS attacks.

You can enable the Hash and URL certificate exchange when configuring an IKE gateway by selecting **HTTP Certificate Exchange** and entering the **Certificate URL**. The peer must also use Hash and URL certificate exchange in order for the exchange to be successful. If the peer cannot use Hash and URL, X.509 certificates are exchanged similarly to how they are exchanged in IKEv1.

If you enable the Hash and URL certificate exchange, you must export your certificate to the certificate server if it is not already there. When you export the certificate, the file format should be **Binary Encoded Certificate (DER)**. See Export a Certificate for a Peer to Access Using Hash and URL.

SA Key Lifetime and Re-Authentication Interval

In IKEv2, two IKE crypto profile values, **Key Lifetime** and **IKEv2 Authentication Multiple**, control the establishment of IKEv2 IKE SAs. The key lifetime is the length of time that a negotiated IKE SA key is effective. Before the key lifetime expires, the SA must be re-keyed; otherwise, upon expiration, the SA must begin a new IKEv2 IKE SA re-key. The default value is 8 hours.

The re-authentication interval is derived by multiplying the **Key Lifetime** by the **IKEv2 Authentication Multiple**. The authentication multiple defaults to 0, which disables the re-authentication feature.

The range of the authentication multiple is 0-50. So, if you were to configure an authentication multiple of 20, for example, the system would perform re-authentication every 20 re-keys, which is every 160 hours. That means the gateway could perform Child SA creation for 160 hours before the gateway must re-authenticate with IKE to recreate the IKE SA from scratch.

In IKEv2, the Initiator and Responder gateways have their own key lifetime value, and the gateway with the shorter key lifetime is the one that will request that the SA be re-keyed.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Set Up Site-to-Site VPN

To set up site-to-site VPN:

- Make sure that your Ethernet interfaces, virtual routers, and zones are configured properly. For more information, see Configure Interfaces and Zones.
- Create your tunnel interfaces. Ideally, put the tunnel interfaces in a separate zone, so that tunneled traffic can use different policies.
- Set up static routes or assign routing protocols to redirect traffic to the VPN tunnels. To support dynamic routing (OSPF, BGP, RIP are supported), you must assign an IP address to the tunnel interface.
- Define IKE gateways for establishing communication between the peers across each end of the VPN tunnel; also define the cryptographic profile that specifies the protocols and algorithms for identification, authentication, and encryption to be used for setting up VPN tunnels in IKEv1 Phase 1. See Set Up an IKE Gateway and Define IKE Crypto Profiles.
- Configure the parameters that are needed to establish the IPSec connection for transfer of data across the VPN tunnel; See Set Up an IPSec Tunnel. For IKEv1 Phase-2, see Define IPSec Crypto Profiles.
- (Optional) Specify how the firewall will monitor the IPSec tunnels. See Set Up Tunnel Monitoring.
- Define security policies to filter and inspect the traffic.

If there is a deny rule at the end of the security rulebase, intra-zone traffic is blocked unless otherwise allowed. Rules to allow IKE and IPSec applications must be explicitly included above the deny rule.

If your VPN traffic is passing through (not originating or terminating on) a PA-7000 Series or PA-5200 Series firewall, configure bi-directional Security policy rules to allow the ESP or AH traffic in both directions.

When these tasks are complete, the tunnel is ready for use. Traffic destined for the zones/addresses defined in policy is automatically routed properly based on the destination route in the routing table, and handled as VPN traffic. For a few examples on site-to-site VPN, see Site-to-Site VPN Quick Configs.

For troubleshooting purposes, you can Enable/Disable, Refresh or Restart an IKE Gateway or IPSec Tunnel.

Set Up an IKE Gateway

To set up a VPN tunnel, the VPN peers or gateways must authenticate each other—using pre-shared keys or digital certificates—and establish a secure channel in which to negotiate the IPSec security association (SA) that will be used to secure traffic between the hosts on each side.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 1 | Define the IKE Gateway.

1. Select **Network** > **Network Profiles** > **IKE Gateways**, **Add** a gateway, and enter the gateway **Name** (**General** tab).
2. Set the **Version** to **IKEv1 only mode**, **IKEv2 only mode**, or **IKEv2 preferred mode**. The IKE gateway begins its negotiation with its peer in the mode you specify here. If you select **IKEv2 preferred mode**, the two peers will use IKEv2 if the remote peer supports it; otherwise they will use IKEv1.

   The **Version** you select also determines which options are available for you to configure on the **Advanced Options** tab.

STEP 2 | Establish the local endpoint of the tunnel (gateway).

1. Select the **Address Type**: **IPv4** or **IPv6**.
2. Select the physical, outgoing **Interface** on the firewall where the local gateway resides.
3. From the **Local IP Address** list, select the IP address that the VPN connection will use as the endpoint; this is the external-facing interface with a publicly routable IP address on the firewall.

STEP 3 | Establish the peer at the far end of the tunnel (gateway).

For **Peer IP Address Type**, select one of the following and enter the corresponding information for the peer:

- **IP**—Enter a **Peer Address** that is either an IPv4 or IPv6 address or enter an address object that is an IPv4 or IPv6 address.
- **FQDN**—Enter a **Peer Address** that is an FQDN string or an address object that uses an FQDN string. If the FQDN or FQDN address object resolves to more than one IP address, the firewall selects the preferred address from the set of addresses that match the Address Type (IPv4 or IPv6) of the IKE gateway as follows:
  - If no IKE security association (SA) is negotiated, the preferred address is the IP address with the smallest value.
  - If the IKE gateway uses an address that is in the set of returned addresses, the firewall selects that address (whether or not it is the smallest address in the set).
  - If the IKE gateway uses an address that isn't in the set of returned addresses, the firewall selects a new address, and it is the smallest address in the set.
- **Dynamic**—Select **Dynamic** if the peer IP address or FQDN value is unknown so that the peer will initiate the negotiation.



Using an FQDN or FQDN address object reduces issues in environments where the peer is subject to dynamic IP address changes (and would otherwise require you to reconfigure this IKE gateway peer address).

STEP 4 | Specify how to authenticate the peer.

Select the **Authentication** method: **Pre-Shared Key** or **Certificate**. If you choose a pre-shared key, proceed to the next step. If you choose a certificate, skip ahead to Step 6, Configure certificate-based authentication.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 5 | Configure a pre-shared key.

1. Enter a **Pre-shared Key**, which is the security key for authentication across the tunnel. Re-enter the value to **Confirm Pre-shared Key**. Use a maximum of 255 ASCII or non-ASCII characters.

   Generate a key that is difficult to crack with dictionary attacks; use a pre-shared key generator, if necessary.

2. For **Local Identification**, choose from the following types and enter a value that you determine: **FQDN (hostname)**, **IP address**, **KEYID (binary format ID string in HEX)**, and **User FQDN (email address)**. Local identification defines the format and identification of the local gateway. If you do not specify a value, the local IP address is used as the local identification value.
3. For **Peer Identification**, choose from the following types and enter a value that you determine: **FQDN (hostname)**, **IP address**, **KEYID (binary format ID string in HEX)**, and **User FQDN (email address)**. Peer identification defines the format and identification of the peer gateway. If you do not specify a value, the peer IP address is used as the peer identification value.
4. Proceed to Step 7 (Configure advanced options for the gateway).

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 6 | Configure certificate-based authentication.

Perform the remaining steps in this procedure if you selected **Certificate** as the method of authenticating the peer gateway at the opposite end of the tunnel.

1. Select a **Local Certificate**—one that is already on the firewall, **Import** a certificate, or **Generate** a new certificate.
   - If you need to **Import** a certificate, then first Import a Certificate for IKEv2 Gateway Authentication and then return to this task.
   - If you want to **Generate** a new certificate, then first generate a certificate on the firewall and then return to this task.
2. (Optional) Enable (select) the **HTTP Certificate Exchange** to configure Hash and URL (IKEv2 only). For an HTTP certificate exchange, enter the **Certificate URL**. For more information, see Hash and URL Certificate Exchange.
3. Select the **Local Identification** type—**Distinguished Name (Subject), FQDN (hostname)**, **IP address**, or **User FQDN (email address)**—and then enter the value. Local identification defines the format and identification of the local gateway.
4. Select the **Peer Identification** type—**Distinguished Name (Subject), FQDN (hostname)**, **IP address**, or **User FQDN (email address)**—and then enter the value. Peer identification defines the format and identification of the peer gateway.
5. Specify the type of **Peer ID Check**:
   - **Exact**—Ensures that the local setting and peer IKE ID payload match exactly.
   - **Wildcard**—Allows the peer identification to match as long as every character before the wildcard (*) matches. The characters after the wildcard need not match.
6. (Optional) **Permit peer identification and certificate payload identification mismatch** to allow a successful IKE SA even when the peer identification does not match the peer identification in the certificate.
7. Choose a **Certificate Profile**. A certificate profile contains information about how to authenticate the peer gateway.
8. (Optional) **Enable strict validation of peer's extended key use** to strictly control how the key can be used.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 7 | Configure advanced options for the gateway.

1. (Optional) **Enable Passive Mode** in the Common Options (**Advanced Options**) to specify that the firewall only respond to IKE connection requests and never initiate them.
2. If you have a device performing NAT between the gateways, **Enable NAT Traversal** to use UDP encapsulation on IKE and UDP protocols, which enables them to pass through intermediate NAT devices.
3. If you configured **IKEv1 only mode** in Step 1, then, on the IKEv1 tab:
   - Select the **Exchange Mode**: **auto**, **aggressive**, or **main**. When you set a firewall to use **auto** exchange mode, it can accept both **main** mode and **aggressive** mode negotiation requests; however, when possible, it will initiate exchanges in **main** mode.

     *If you do not set the exchange mode to **auto**, then you must configure both peers with the same exchange mode to allow each peer to accept negotiation requests.*
   - Select an existing profile or keep the default profile from the **IKE Crypto Profile** list. If needed, you can Define IKE Crypto Profiles.
   - (Only when using certificate-based authentication and when exchange mode is not set to aggressive mode) Click **Enable Fragmentation** to enable the firewall to operate with IKE Fragmentation.
   - Click **Dead Peer Detection** and enter an **Interval** (range is 2 to 100 seconds). For **Retry**, define the time to delay (range is 2 to 100 seconds) before attempting to re-check availability. Dead peer detection identifies inactive or unavailable IKE peers by sending an IKE phase 1 notification payload to the peer and waiting for an acknowledgment.
4. If you configured **IKEv2 only mode** or **IKEv2 preferred mode** in Step 1, then on the IKEv2 tab:
   - Select an **IKE Crypto Profile**, which configures IKE Phase 1 options such, as the DH group, hash algorithm, and ESP authentication. For information about IKE crypto profiles, see IKE Phase 1.
   - (Optional) Enable **Strict Cookie Validation** Cookie Activation Threshold and Strict Cookie Validation.
   - (Optional) **Enable Liveness Check** and enter an **Interval (sec)** (default is 5) if you want to have the gateway send a message request to its gateway peer, requesting a response. If necessary, the Initiator attempts the liveness check as many as 10 times. If it doesn't get a response, the Initiator closes and deletes the IKE_SA and CHILD_SA. The Initiator will start over by sending out another IKE_SA_INIT.

STEP 8 | Click **OK** and **Commit** your changes.

Export a Certificate for a Peer to Access Using Hash and URL

IKEv2 supports Hash and URL Certificate Exchange as a method of having the peer at the remote end of the tunnel fetch the certificate from a server where you have exported the certificate. Perform this task to export your certificate to that server. You must have already created a certificate using **Device** > **Certificate Management**.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 1 | Select **Device** > **Certificates**, and if your platform supports multiple virtual systems, for **Location**, select the appropriate virtual system.

STEP 2 | On the **Device Certificates** tab, select the certificate to **Export** to the server.

The status of the certificate should be valid, not expired. The firewall will not stop you from exporting an invalid certificate.

STEP 3 | For **File Format**, select **Binary Encoded Certificate (DER)**.

STEP 4 | Leave **Export private key** clear. Exporting the private key is unnecessary for Hash and URL.

STEP 5 | Click **OK**.

Import a Certificate for IKEv2 Gateway Authentication

Perform this task if you are authenticating a peer for an IKEv2 gateway and you did not use a local certificate already on the firewall; you want to import a certificate from elsewhere.

This task presumes that you selected **Network** > **IKE Gateways**, added a gateway, and for **Local Certificate**, you clicked **Import**.

STEP 1 | Import a certificate.

1. Select **Network** > **IKE Gateways**, **Add** a gateway, and on the **General** tab, for **Authentication**, select **Certificate**. For **Local Certificate**, click **Import**.
2. In the Import Certificate window, enter a **Certificate Name** for the certificate you are importing.
3. Select **Shared** if this certificate is to be shared among multiple virtual systems.
4. For **Certificate File**, **Browse** to the certificate file. Click on the file name and click **Open**, which populates the **Certificate File** field.
5. For **File Format**, select one of the following:
   - **Base64 Encoded Certificate (PEM)**—Contains the certificate, but not the key. It is cleartext.
   - **Encrypted Private Key and Certificate (PKCS12)**—Contains both the certificate and the key.
6. Select **Import private key** if the key is in a different file from the certificate file. The key is optional, with the following exception:
   - You must import a key if you set the **File Format** to **PEM**. Enter a **Key file** by clicking **Browse** and navigating to the key file to import.
   - Enter a **Passphrase** and **Confirm Passphrase**.
7. Click **OK**.

STEP 2 | Continue to the next task.

Step Configure certificate-based authentication.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Change the Key Lifetime or Authentication Interval for IKEv2

This task is optional; the default setting of the IKEv2 IKE SA re-key lifetime is 8 hours. The default setting of the IKEv2 Authentication Multiple is 0, meaning the re-authentication feature is disabled. For more information, see SA Key Lifetime and Re-Authentication Interval.

To change the default values, perform the following task. A prerequisite is that an IKE crypto profile already exists.

STEP 1 | Change the SA key lifetime or authentication interval for an IKE Crypto profile.

1. Select **Network** > **Network Profiles** > **IKE Crypto** and select the IKE Crypto profile that applies to the local gateway.
2. For the **Key Lifetime**, select a unit (**Seconds**, **Minutes**, **Hours**, or **Days**) and enter a value. The minimum is three minutes.
3. For **IKE Authentication Multiple**, enter a value, which is multiplied by the lifetime to determine the re-authentication interval.

STEP 2 | Commit your changes.

Click **OK** and **Commit**.

Change the Cookie Activation Threshold for IKEv2

Perform the following task if you want a firewall to have a threshold different from the default setting of 500 half-opened SA sessions before cookie validation is required. For more information about cookie validation, see Cookie Activation Threshold and Strict Cookie Validation.

STEP 1 | Change the Cookie Activation Threshold.

1. Select **Device** > **Setup** > **Session** and edit the VPN Session Settings. For **Cookie Activation Threshold**, enter the maximum number of half-opened SAs that are allowed before the responder requests a cookie from the initiator (range is 0-65,535; default is 500).
2. Click **OK**.

STEP 2 | Commit your changes.

Click **OK** and **Commit**.

Configure IKEv2 Traffic Selectors

In IKEv2, you can configure Traffic Selectors, which are components of network traffic that are used during IKE negotiation. Traffic selectors are used during the CHILD_SA (tunnel creation) Phase 2 to set up the tunnel and to determine what traffic is allowed through the tunnel. The two IKE gateway peers must negotiate and agree on their traffic selectors; otherwise, one side narrows its address range to reach agreement. One IKE connection can have multiple tunnels; for example, you can assign different tunnels to each department to isolate their traffic. Separation of traffic also allows features such as QoS to be implemented. Use the following workflow to configure traffic selectors.

STEP 1 | Select **Network** > **IPSec Tunnels** > **Proxy IDs**.

STEP 2 | Select the **IPv4** or **IPv6** tab.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 3 | Click **Add** and enter the **Name** in the **Proxy ID** field.

STEP 4 | In the **Local** field, enter the **Source IP Address**.

STEP 5 | In the **Remote** field, enter the **Destination IP Address**.

STEP 6 | In the **Protocol** field, select the transport protocol (**TCP** or **UDP**).

STEP 7 | Click **OK**.

Define Cryptographic Profiles

A cryptographic profile specifies the ciphers used for authentication and/or encryption between two IKE peers, and the lifetime of the key. The time period between each renegotiation is known as the lifetime; when the specified time expires, the firewall renegotiates a new set of keys.

For securing communication across the VPN tunnel, the firewall requires IKE and IPSec cryptographic profiles for completing IKE phase 1 and phase 2 negotiations, respectively. The firewall includes a default IKE crypto profile and a default IPSec crypto profile that are ready for use.

- Define IKE Crypto Profiles
- Define IPSec Crypto Profiles

Define IKE Crypto Profiles

The IKE crypto profile is used to set up the encryption and authentication algorithms used for the key exchange process in IKE Phase 1, and lifetime of the keys, which specifies how long the keys are valid. To invoke the profile, you must attach it to the IKE Gateway configuration.

*All IKE gateways configured on the same interface or local IP address must use the same crypto profile when the IKE gateway's **Peer IP Address Type** is configured as **Dynamic** and IKEv1 main mode or IKEv2 is applied.*

STEP 1 | Create a new IKE profile.

1. Select **Network** > **Network Profiles** > **IKE Crypto** and select **Add**.
2. Enter a **Name** for the new profile.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 2 | Specify the DH (Diffie–Hellman) Group for key exchange and the Authentication and Encryption algorithms.

Click **Add** in the corresponding sections (DH Group, Authentication, and Encryption) and select from the menus.

If you are not certain what the VPN peers support, add multiple groups or algorithms in the order of most-to-least secure; the peers negotiate the strongest supported group or algorithm to establish the tunnel.

- DH Group—
  - **group20**
  - **group19**
  - **group14**
  - **group5**
  - **group2**
  - **group1**
- Authentication—
  - **sha512**
  - **sha384**
  - **sha256**
  - **sha1**
  - **md5**
  - (PAN-OS 10.0.3 and later 10.1 releases) **none**

  *If you select an AES-GCM algorithm for encryption, you must select the Authentication setting **none** or the commit will fail. The hash is automatically selected based on the DH Group selected. DH Group 19 and below uses **sha256**; DH Group 20 uses **sha384**.*
- Encryption—
  - (PAN-OS 10.0.3 and later 10.1 releases) **aes-256-gcm** (requires IKEv2; DH Group should be set to **group20**)
  - (PAN-OS 10.0.3 and later 10.1 releases) **aes-128-gcm** (requires IKEv2 and DH Group set to **group19**)
  - **aes-256-cbc**
  - **aes-192-cbc**
  - **aes-128-cbc**
  - **3des**
  - **des**

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

*Choose the strongest authentication and encryption algorithms the peer can support. For the authentication algorithm, use SHA-256 or higher (SHA-384 or higher preferred for long-lived transactions). Do not use SHA-1 or MD5. For the encryption algorithm, use AES; DES and 3DES are weak and vulnerable. AES with Galois/Counter Mode (AES-GCM) provides the strongest security and has built-in authentication, so you must set Authentication to **none** if you select **aes-256-gcm** or **aes-128-gcm** encryption.*

STEP 3 | Specify the duration for which the key is valid and the re-authentication interval.

For details, see SA Key Lifetime and Re-Authentication Interval.

1. In the **Key Lifetime** fields, specify the period (in seconds, minutes, hours, or days) for which the key is valid (range is 3 minutes to 365 days; default is 8 hours). When the key expires, the firewall renegotiates a new key. A lifetime is the period between each renegotiation.
2. For the **IKEv2 Authentication Multiple**, specify a value (range is 0-50; default is 0) that is multiplied by the **Key Lifetime** to determine the authentication count. The default value of 0 disables the re-authentication feature.

STEP 4 | Commit your IKE Crypto profile.

Click **OK** and click **Commit**.

STEP 5 | Attach the IKE Crypto profile to the IKE Gateway configuration.

See Configure advanced options for the gateway.

Define IPSec Crypto Profiles

The IPSec crypto profile is invoked in IKE Phase 2. It specifies how the data is secured within the tunnel when Auto Key IKE is used to automatically generate keys for the IKE SAs.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 1 | Create a new IPSec profile.

1. Select **Network** > **Network Profiles** > **IPSec Crypto** and select **Add**.
2. Enter a **Name** for the new profile.
3. Select the **IPSec Protocol**—ESP or AH—that you want to apply to secure the data as it traverses across the tunnel.

   As a best practice, select ESP (Encapsulating Security payload) over AH (Authentication Header) because ESP offers both confidentiality and authentication for the connection whereas AH offers only authentication.
4. Click **Add** and select the **Authentication** and **Encryption** algorithms for ESP, and **Authentication** algorithms for AH, so that the IKE peers can negotiate the keys for the secure transfer of data across the tunnel.

   If you are not certain of what the IKE peers support, add multiple algorithms in the order of most-to-least secure as follows; the peers negotiate the strongest supported algorithm to establish the tunnel:

   - Encryption—**aes-256-gcm**, **aes-256-cbc**, **aes-192-cbc**, **aes-128-gcm**, **aes-128-ccm** (the VM-Series firewall doesn't support this option), **aes-128-cbc**, **3des**, **des**.

     As a best practice, choose the strongest authentication and encryption algorithms the peer can support. For the authentication algorithm, use SHA-256 or higher (SHA-384 or higher preferred for long-lived transactions). Do not use SHA-1, MD5 or none. For the encryption algorithm, use AES; DES and 3DES are weak and vulnerable.
   - Authentication—**sha512**, **sha384**, **sha256**, **sha1**, **md5**.

STEP 2 | Select the DH Group to use for the IPSec SA negotiations in IKE phase 2.

From **DH Group**, select the key strength you want to use: **group1**, **group2**, **group5**, **group14**, **group19**, or **group20**. For highest security, choose the group with the highest number.

If you don't want to renew the key that the firewall creates during IKE phase 1, select **no-pfs** (no perfect forward secrecy); the firewall reuses the current key for the IPSec security association (SA) negotiations.

STEP 3 | Specify the duration of the key—time and volume of traffic.

Using a combination of time and traffic volume allows you to ensure safety of data.

Select the **Lifetime** or time period for which the key is valid in seconds, minutes, hours, or days (range is 3 minutes to 365 days). When the specified time expires, the firewall will renegotiate a new set of keys.

Select the **Lifesize** or volume of data after which the keys must be renegotiated.

STEP 4 | Commit your IPSec profile.

Click **OK** and click **Commit**.

STEP 5 | Attach the IPSec Profile to an IPSec tunnel configuration.

See Set up key exchange.

https://t.me/learningnets

Set Up an IPSec Tunnel

The IPSec tunnel configuration allows you to authenticate and/or encrypt the data (IP packet) as it traverses the tunnel.

If you are setting up the firewall to work with a peer that supports policy-based VPN, you must define Proxy IDs. Devices that support policy-based VPN use specific security rules/policies or access-lists (source addresses, destination addresses and ports) for permitting interesting traffic through an IPSec tunnel. These rules are referenced during quick mode/IKE phase 2 negotiation, and are exchanged as Proxy-IDs in the first or the second message of the process. So, if you are configuring the firewall to work with a policy-based VPN peer, for a successful phase 2 negotiation you must define the Proxy-ID so that the setting on both peers is identical. If the Proxy-ID is not configured, because the firewall supports route-based VPN, the default values used as Proxy-ID are source ip: 0.0.0.0/0, destination ip: 0.0.0.0/0 and application: any; and when these values are exchanged with the peer, it results in a failure to set up the VPN connection.

STEP 1 | Select **Network** > **IPSec Tunnels** and then **Add** a new tunnel configuration.

STEP 2 | On the **General** tab, enter a **Name** for the tunnel.

STEP 3 | Select the **Tunnel interface** on which to set up the IPSec tunnel.

To create a new tunnel interface:

1. Select **Tunnel Interface** > **New Tunnel Interface**. (You can also select **Network** > **Interfaces** > **Tunnel** and click **Add**.)
2. In the **Interface Name** field, specify a numeric suffix, such as **.2**.
3. On the **Config** tab, select the **Security Zone** list to define the zone as follows:

Use your trust zone as the termination point for the tunnel—Select the zone. Associating the tunnel interface with the same zone (and virtual router) as the external-facing interface on which the packets enter the firewall mitigates the need to create inter-zone routing.

Or:

Create a separate zone for VPN tunnel termination (Recommended)—Select **New Zone**, define a **Name** for the new zone (for example vpn-corp), and click **OK**.

1. For **Virtual Router**, select **default**.
2. (Optional) If you want to assign an IPv4 address to the tunnel interface, select the **IPv4** tab, and **Add** the IP address and network mask, for example 10.31.32.1/32.
3. Click **OK**.

STEP 4 | (Optional) Enable IPv6 on the tunnel interface.

1. Select the IPv6 tab on **Network** > **Interfaces** > **Tunnel** > **IPv6**.
2. Select **Enable IPv6 on the interface**.

   This option allows you to route IPv6 traffic over an IPv4 IPSec tunnel and will provide confidentiality between IPv6 networks. The IPv6 traffic is encapsulated by IPv4 and then

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

ESP. To route IPv6 traffic to the tunnel, you can use a static route to the tunnel, or use OSPFv3, or use a Policy-Based Forwarding (PBF) rule.

3. Enter the 64-bit extended unique **Interface ID** in hexadecimal format, for example, 00:26:08:FF:FE:DE:4E:29. By default, the firewall will use the EUI-64 generated from the physical interface's MAC address.
4. To assign an IPv6 **Address** to the tunnel interface, **Add** the IPv6 address and prefix length, for example 2001:400:f00::1/64. If Prefix is not selected, the IPv6 address assigned to the interface will be wholly specified in the address text box.
   1. Select **Use interface ID as host portion** to assign an IPv6 address to the interface that will use the interface ID as the host portion of the address.
   2. Select **Anycast** to include routing through the nearest node.

STEP 5 | Set up key exchange.

On the **General** tab, configure one of the following types of key exchange:

Set up Auto Key exchange

1. Select the IKE Gateway. To set up an IKE gateway, see Set Up an IKE Gateway.
2. (Optional) Select the default IPSec Crypto Profile. To create a new IPSec Profile, see Define IPSec Crypto Profiles.

Set up Manual Key exchange

1. Specify the **Local SPI** for the local firewall. SPI is a 32-bit hexadecimal index that is added to the header for IPSec tunneling to assist in differentiating between IPSec traffic flows; it is used to create the SA required for establishing a VPN tunnel.
2. Select the **Interface** that will be the tunnel endpoint, and optionally select the IP address for the local interface that is the endpoint of the tunnel.
3. Select the protocol to be used—**AH** or **ESP**.
4. For AH, select the **Authentication** method and enter a **Key** and then **Confirm Key**.
5. For ESP, select the **Authentication** method and enter a **Key** and then **Confirm Key**. Then, select the **Encryption** method and enter a **Key** and then **Confirm Key**, if needed.
6. Specify the **Remote SPI** for the remote peer.
7. Enter the **Remote Address**, the IP address of the remote peer.

STEP 6 | Protect against a replay attack.

Anti-replay is a sub-protocol of IPSec and is part of the Internet Engineering Task Force (IETF) Request for Comments (RFC) 6479. The anti-replay protocol is used to prevent hackers from injecting or making changes in packets that travel from a source to a destination and uses a unidirectional security association in order to establish a secure connection between two nodes in the network.

After a secure connection is established, the anti-replay protocol uses packet sequence numbers to defeat replay attacks. When the source sends a message, it adds a sequence number to its packet; the sequence number starts at 0 and is incremented by 1 for each subsequent packet. The destination maintains the sequence of numbers in a *sliding window* format, maintains a record of the sequence numbers of validated received packets, and rejects all packets that have a sequence number that is lower than the lowest in the sliding window

https://t.me/learningnets

(packets that are too old) or packets that already appear in the sliding window (duplicate or replayed packets). Accepted packets, after they are validated, update the sliding window, displacing the lowest sequence number out of the window if it was already full.

1. On the General tab, select **Show Advanced Options** and select **Enable Replay Protection** to detect and neutralize against replay attacks.
2. Select the **Anti Replay Window** to use. You can select a anti-replay window size of 64, 128, 256, 512, 1024, 2048, or 4096. The default is 1024.

STEP 7 | (Optional) Preserve the Type of Service header for the priority or treatment of IP packets.

In the Show Advanced Options section, select **Copy TOS Header**. This copies the Type of Service (TOS) header from the inner IP header to the outer IP header of the encapsulated packets in order to preserve the original TOS information.

If there are multiple sessions inside the tunnel (each with a different TOS value), copying the TOS header can cause the IPSec packets to arrive out of order.

STEP 8 | (Optional) Select **Add GRE Encapsulation** to enable GRE over IPSec.

Add GRE encapsulation in cases where the remote endpoint requires traffic to be encapsulated within a GRE tunnel before IPSec encrypts the traffic. For example, some implementations require multicast traffic to be encapsulated before IPSec encrypts it. Add GRE Encapsulation when the GRE packet encapsulated in IPSec has the same source IP address and destination IP address as the encapsulating IPSec tunnel.

STEP 9 | Enable Tunnel Monitoring.

You must assign an IP address to the tunnel interface for monitoring.

To alert the device administrator to tunnel failures and to provide automatic failover to another tunnel interface:

1. Select **Tunnel Monitor**.
2. Specify a **Destination IP** address on the other side of the tunnel to determine if the tunnel is working properly.
3. Select a **Profile** to determine the action upon tunnel failure. To create a new profile, see Define a Tunnel Monitoring Profile.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 10 | Create a Proxy ID to identify the VPN peers.

This step is required only if the VPN peer uses policy-based VPN.

1. Select **Network** > **IPSec Tunnels** and click **Add**.
2. Select the **Proxy IDs** tab.
3. Select the **IPv4** or **IPv6** tab.
4. Click **Add** and enter the **Proxy ID** name.
5. Enter the **Local** IP address or subnet for the VPN gateway.
6. Enter the **Remote** address for the VPN gateway.
7. Select the **Protocol**:
   - **Number**—Specify the protocol number (used for interoperability with third-party devices).
   - **Any**—Allows TCP and/or UDP traffic.
   - **TCP**—Specify the Local Port and Remote Port numbers.
   - **UDP**—Specify the Local Port and Remote Port numbers.
8. Click **OK**.

STEP 11 | Commit your changes.

Click **OK** and **Commit**.

Set Up Tunnel Monitoring

To provide uninterrupted VPN service, you can use the Dead Peer Detection capability along with the tunnel monitoring capability on the firewall. You can also monitor the status of the tunnel. These monitoring tasks are described in the following sections:

- Define a Tunnel Monitoring Profile
- View the Status of the Tunnels

Define a Tunnel Monitoring Profile

A tunnel monitoring profile allows you to verify connectivity between the VPN peers; you can configure the tunnel interface to ping a destination IP address at a specified interval and specify the action if the communication across the tunnel is broken.

STEP 1 | Select **Network** > **Network Profiles** > **Monitor**. A default tunnel monitoring profile is available for use.

STEP 2 | Click **Add**, and enter a **Name** for the profile.

STEP 3 | Select the **Action** to take if the destination IP address is unreachable.

- **Wait Recover**—the firewall waits for the tunnel to recover. It continues to use the tunnel interface in routing decisions as if the tunnel were still active.
- **Fail Over**—forces traffic to a back-up path if one is available. The firewall disables the tunnel interface, and thereby disables any routes in the routing table that use the interface.

In either case, the firewall attempts to accelerate the recovery by negotiating new IPSec keys.

https://t.me/learningnets

STEP 4 | Specify the **Interval (sec)** and **Threshold** to trigger the specified action.

- **Threshold** specifies the number of heartbeats to wait before taking the specified action (range is 2-100; default is 5).
- **Interval (sec)** specifies the time (in seconds) between heartbeats (range is 2-10; default is 3).

STEP 5 | Attach the monitoring profile to the IPsec Tunnel configuration. See Enable Tunnel Monitoring.

View the Status of the Tunnels

The status of the tunnel informs you about whether or not valid IKE phase-1 and phase-2 SAs have been established, and whether the tunnel interface is up and available for passing traffic.

Because the tunnel interface is a logical interface, it cannot indicate a physical link status. Therefore, you must enable tunnel monitoring so that the tunnel interface can verify connectivity to an IP address and determine if the path is still usable. If the IP address is unreachable, the firewall will either wait for the tunnel to recover or failover. When a failover occurs, the existing tunnel is torn down and routing changes are triggered to set up a new tunnel and redirect traffic.

STEP 1 | Select **Network** > **IPSec Tunnels**.

STEP 2 | View the **Tunnel Status**.

- Green indicates a valid IPSec SA tunnel.
- Red indicates that IPSec SA is not available or has expired.

STEP 3 | View the **IKE Gateway Status**.

- Green indicates a valid IKE phase-1 SA.
- Red indicates that IKE phase-1 SA is not available or has expired.

STEP 4 | View the **Tunnel Interface Status**.

- Green indicates that the tunnel interface is up.
- Red indicates that the tunnel interface is down, because tunnel monitoring is enabled and the status is down.

To troubleshoot a VPN tunnel that is not yet up, see Interpret VPN Error Messages.

Enable/Disable, Refresh or Restart an IKE Gateway or IPSec Tunnel

You can enable, disable, refresh or restart an IKE gateway or VPN tunnel to make troubleshooting easier.

- Enable or Disable an IKE Gateway or IPSec Tunnel
- Refresh and Restart Behaviors
- Refresh or Restart an IKE Gateway or IPSec Tunnel

Enable or Disable an IKE Gateway or IPSec Tunnel

Enable or disable an IKE gateway or IPSec tunnel to make troubleshooting easier.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

- Enable or disable an IKE gateway.
  1. Select **Network** > **Network Profiles** > **IKE Gateways** and select the gateway you want to enable or disable.
  2. At the bottom of the screen, click **Enable** or **Disable**.

- Enable or disable an IPSec tunnel.
  1. Select **Network** > **IPSec Tunnels** and select the tunnel you want to enable or disable.
  2. At the bottom of the screen, click **Enable** or **Disable**.

Refresh and Restart Behaviors

You can Refresh or Restart an IKE Gateway or IPSec Tunnel. The refresh and restart behaviors for an IKE gateway and IPSec tunnel are as follows:

Phase	Refresh	Restart
IKE Gateway (IKE Phase 1)	Updates the onscreen statistics for the selected IKE gateway. Equivalent to issuing a second `show` command in the CLI (after an initial `show` command).	Restarts the selected IKE gateway. **IKEv2**: Also restarts any associated child IPSec security associations (SAs). **IKEv1**: Does not restart the associated IPSec SAs. A restart is disruptive to all existing sessions. Equivalent to issuing a **`clear`**, **`test`**, **`show`** command sequence in the CLI.
IPSec Tunnel (IKE Phase 2)	Updates the onscreen statistics for the selected IPSec tunnel. Equivalent to issuing a second `show` command in the CLI (after an initial `show` command).	Restarts the IPSec tunnel. A restart is disruptive to all existing sessions. Equivalent to issuing a **`clear`**, **`test`**, **`show`** command sequence in the CLI.

Refresh or Restart an IKE Gateway or IPSec Tunnel

Keep in mind that the result of restarting an IKE gateway depends on whether it is IKEv1 or IKEv2. See Refresh and Restart Behaviors for an IKE gateway (IKEv1 and IKEv2) and for an IPSec tunnel.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

- Refresh or restart an IKE gateway.
  1. Select **Network** > **IPSec Tunnels** and select the tunnel for the gateway you want to refresh or restart.
  2. In the row for that tunnel, under the Status column, click **IKE Info**.
  3. At the bottom of the IKE Info screen, click the action you want:
     - **Refresh**—Updates the statistics on the screen.
     - **Restart**—Clears the SAs, so traffic is dropped until the IKE negotiation starts over and the tunnel is recreated.

- Refresh or restart an IPSec tunnel.

  You might determine that the tunnel needs to be refreshed or restarted because you use the tunnel monitor to monitor the tunnel status, or you use an external network monitor to monitor network connectivity through the IPSec tunnel.
  1. Select **Network** > **IPSec Tunnels** and select the tunnel you want to refresh or restart.
  2. In the row for that tunnel, under the Status column, click **Tunnel Info**.
  3. At the bottom of the Tunnel Info screen, click the action you want:
     - **Refresh**—Updates the onscreen statistics.
     - **Restart**—Clears the SAs, so traffic is dropped until the IKE negotiation starts over and the tunnel is recreated.

Test VPN Connectivity

Perform this task to test VPN connectivity.

STEP 1 | Initiate IKE phase 1 by either pinging a host across the tunnel or using the following CLI command:

```
test vpn ike-sa gateway <gateway_name>
```

STEP 2 | Enter the following command to test if IKE phase 1 is set up:

```
show vpn ike-sa gateway <gateway_name>
```

In the output, check whether the Security Association displays. If it doesn't, review the system log messages to interpret the reason for failure.

STEP 3 | Initiate IKE phase 2 by either pinging a host from across the tunnel or using the following CLI command:

```
test vpn ipsec-sa tunnel <tunnel_name>
```

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 4 | Enter the following command to test if IKE phase 2 is set up:

```
show vpn ipsec-sa tunnel <tunnel_name>
```

In the output, check whether the Security Association displays. If it doesn't, review the system log messages to interpret the reason for failure.

STEP 5 | To view the VPN traffic flow information, use the following command:

```
show vpn flow
total tunnels configured:               1
filter - type IPSec, state any

total IPSec tunnel configured:        1
total IPSec tunnel shown:                1

name                    id      state      local-ip       peer-ip
       tunnel-i/f
------------------------------------------------------------------------------
vpn-to-siteB       5       active
    100.1.1.1     200.1.1.1     tunnel.41
```

Interpret VPN Error Messages

The following table lists some of the common VPN error messages that are logged in the system log.

Table 6: Syslog Error Messages for VPN Issues

If error is this:	Try this:

`IKE phase-1 negotiation is failed as initiator, main mode. Failed SA: x.x.x.x[500]-y.y.y.y[500] cookie:84222f276c2fa2e9:0000000000000000 due to timeout.`
or
`IKE phase 1 negotiation is failed. Couldn't find configuration for IKE phase-1 request for peer IP x.x.x.x[1929]`
	• Verify that the public IP address for each VPN peer is accurate in the IKE Gateway configuration.
	• Verify that the IP addresses can be pinged and that routing issues are not causing the connection failure.

`Received unencrypted notify payload (no proposal chosen) from IP x.x.x.x[500] to y.y.y.y[500], ignored...` or	Check the IKE Crypto profile configuration to verify that the proposals on both sides have a common encryption, authentication, and DH Group proposal.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

If error is this:	Try this:
`IKE phase-1 negotiation is failed. Unable to process peer's SA payload.`	

`pfs group mismatched:my: 2peer: 0`
or
`IKE phase-2 negotiation failed when processing SA payload. No suitable proposal found in peer's SA payload.`
	Check the IPSec Crypto profile configuration to verify that:
	• pfs is either enabled or disabled on both VPN peers
	• the DH Groups proposed by each peer has at least one DH Group in common

`IKE phase-2 negotiation failed when processing Proxy ID. Received local id x.x.x.x/x type IPv4 address protocol 0 port 0, received remote id y.y.y.y/y type IPv4 address protocol 0 port 0.`	The VPN peer on one end is using policy-based VPN. You must configure a Proxy ID on the Palo Alto Networks firewall. See Create a Proxy ID to identify the VPN peers..

©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

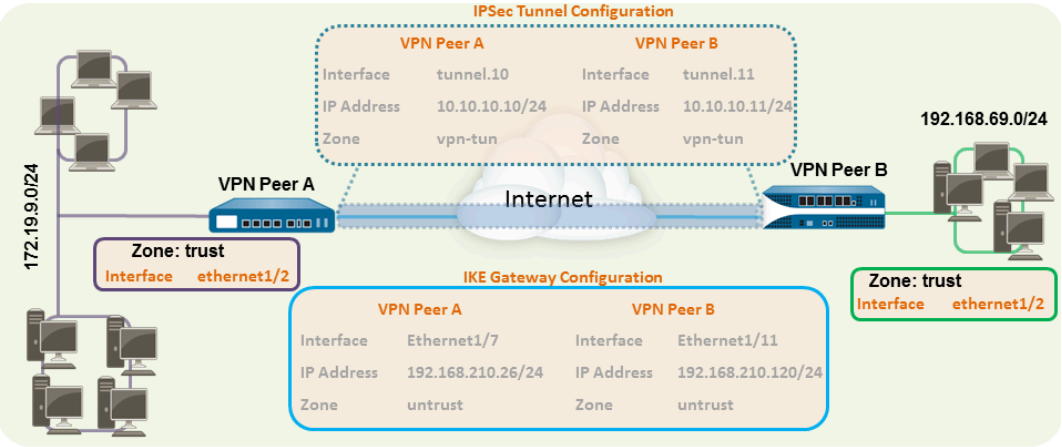
Site-to-Site VPN Quick Configs

The following sections provide instructions for configuring some common VPN deployments:

- Site-to-Site VPN with Static Routing
- Site-to-Site VPN with OSPF
- Site-to-Site VPN with Static and Dynamic Routing

Site-to-Site VPN with Static Routing

The following example shows a VPN connection between two sites that use static routes. Without dynamic routing, the tunnel interfaces on VPN Peer A and VPN Peer B do not require an IP address because the firewall automatically uses the tunnel interface as the next hop for routing traffic across the sites. However, to enable tunnel monitoring, a static IP address has been assigned to each tunnel interface.



 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 1 | Configure a Layer 3 interface.

This interface is used for the IKE phase-1 tunnel.

1. Select **Network** > **Interfaces** > **Ethernet** and then select the interface you want to configure for VPN.
2. Select **Layer3** from the **Interface Type**.
3. On the **Config** tab, select the **Security Zone** to which the interface belongs:
   - The interface must be accessible from a zone outside of your trust network. Consider creating a dedicated VPN zone for visibility and control over your VPN traffic.
   - If you have not yet created the zone, select **New Zone** from the **Security Zone**, define a **Name** for the new zone and then click **OK**.
4. Select the **Virtual Router** to use.
5. To assign an IP address to the interface, select the **IPv4** tab, click **Add** in the IP section, and enter the IP address and network mask to assign to the interface, for example 192.168.210.26/24.
6. To save the interface configuration, click **OK**.

   In this example, the configuration for VPN Peer A is:

   - **Interface**—ethernet1/7
   - **Security Zone**—untrust
   - **Virtual Router**—default
   - **IPv4**—192.168.210.26/24

   The configuration for VPN Peer B is:

   - **Interface**—ethernet1/11
   - **Security Zone**—untrust
   - **Virtual Router**—default
   - **IPv4**—192.168.210.120/24

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 2 | Create a tunnel interface and attach it to a virtual router and security zone.

1. Select **Network** > **Interfaces** > **Tunnel** and click **Add**.
2. In the **Interface Name** field, specify a numeric suffix, such as **.1**.
3. On the **Config** tab, expand the **Security Zone** to define the zone as follows:
   - To use your trust zone as the termination point for the tunnel, select the zone.
   - (Recommended) To create a separate zone for VPN tunnel termination, click **New Zone**. In the Zone dialog, define a **Name** for new zone (for example *vpn-tun*), and then click **OK**.
4. Select the **Virtual Router**.
5. (Optional) Assign an IP address to the tunnel interface, select the **IPv4** or **IPv6** tab, click **Add** in the IP section, and enter the IP address and network mask to assign to the interface.

   With static routes, the tunnel interface does not require an IP address. For traffic that is destined to a specified subnet/IP address, the tunnel interface will automatically become the next hop. Consider adding an IP address if you want to enable tunnel monitoring.
6. To save the interface configuration, click **OK**.

   In this example, the configuration for VPN Peer A is:

   - **Interface**—tunnel.10
   - **Security Zone**—vpn_tun
   - **Virtual Router**—default
   - **IPv4**—172.19.9.2/24

   The configuration for VPN Peer B is:

   - **Interface**—tunnel.11
   - **Security Zone**—vpn_tun
   - **Virtual Router**—default
   - **IPv4**—192.168.69.2/24

STEP 3 | Configure a static route, on the virtual router, to the destination subnet.

1. Select **Network** > **Virtual Router** and click the router you defined in the prior step.
2. Select **Static Route**, click **Add**, and enter a new route to access the subnet that is at the other end of the tunnel.

   In this example, the configuration for VPN Peer A is:

   - **Destination**—192.168.69.0/24
   - **Interface**—tunnel.10

   The configuration for VPN Peer B is:

   - **Destination**—172.19.9.0/24
   - **Interface**—tunnel.11

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 4 | Set up the Crypto profiles (IKE Crypto profile for phase 1 and IPSec Crypto profile for phase 2).

Complete this task on both peers and make sure to set identical values.

1. Select **Network** > **Network Profiles** > **IKE Crypto**. In this example, we use the default profile.
2. Select **Network** > **Network Profiles** > **IPSec Crypto**. In this example, we use the default profile.

STEP 5 | Set up the IKE Gateway.

1. Select **Network** > **Network Profiles** > **IKE Gateway**.
2. Click **Add** and configure the options in the **General** tab.

   In this example, the configuration for VPN Peer A is:

   - **Interface**—ethernet1/7
   - **Local IP address**—192.168.210.26/24
   - **Peer IP type/address**—static/192.168.210.120
   - **Preshared keys**—enter a value
   - **Local identification**—None; this means that the local IP address will be used as the local identification value.
   - The configuration for VPN Peer B is:
   - **Interface**—ethernet1/11
   - **Local IP address**—192.168.210.120/24
   - **Peer IP type/address**—static/192.168.210.26
   - **Preshared keys**—enter same value as on Peer A
   - **Local identification**—None
3. Select **Advanced Phase 1 Options** and select the IKE Crypto profile you created earlier to use for IKE phase 1.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 6 | Set up the IPSec Tunnel.

1. Select **Network** > **IPSec Tunnels**.
2. Click **Add** and configure the options in the **General** tab.

   In this example, the configuration for VPN Peer A is:

   - **Tunnel Interface**—tunnel.10
   - **Type**—Auto Key
   - **IKE Gateway**—Select the IKE Gateway defined above.
   - **IPSec Crypto Profile**—Select the IPSec Crypto profile defined in Step 4.

   The configuration for VPN Peer B is:

   - **Tunnel Interface**—tunnel.11
   - **Type**—Auto Key
   - **IKE Gateway**—Select the IKE Gateway defined above.
   - **IPSec Crypto Profile**—Select the IPSec Crypto defined in Step 4.
3. (Optional) Select **Show Advanced Options**, select **Tunnel Monitor**, and specify a Destination IP address to ping for verifying connectivity. Typically, the tunnel interface IP address for the VPN Peer is used.
4. (Optional) To define the action on failure to establish connectivity, see Define a Tunnel Monitoring Profile.

STEP 7 | Create policies to allow traffic between the sites (subnets).

1. Select **Policies** > **Security**.
2. Create rules to allow traffic between the untrust and the vpn-tun zone and the vpn-tun and the untrust zone for traffic originating from specified source and destination IP addresses.

STEP 8 | Commit any pending configuration changes.

Click **Commit**.

STEP 9 | Test VPN Connectivity.

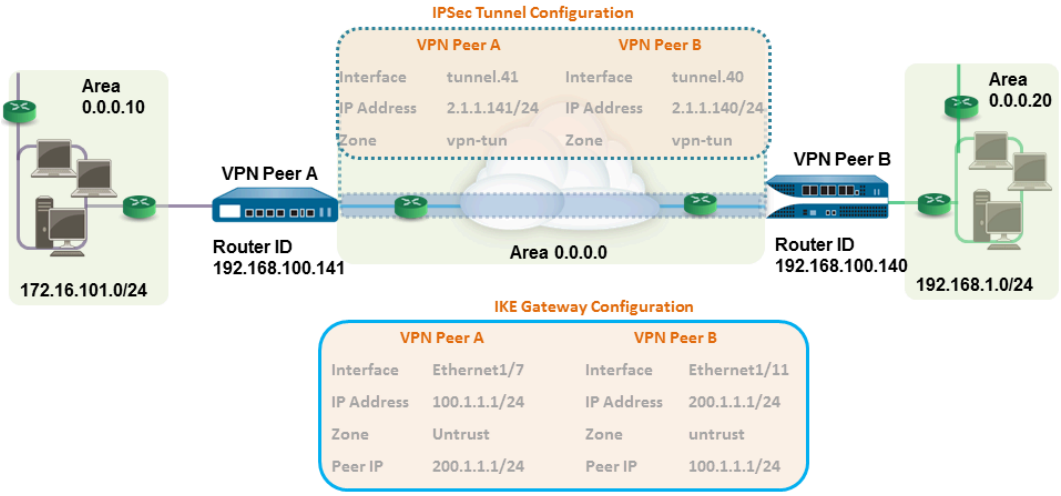
See also View the Status of the Tunnels.

Site-to-Site VPN with OSPF

In this example, each site uses OSPF for dynamic routing of traffic. The tunnel IP address on each VPN peer is statically assigned and serves as the next hop for routing traffic between the two sites.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets



STEP 1 | Configure the Layer 3 interfaces on each firewall.

1. Select **Network** > **Interfaces** > **Ethernet** and then select the interface you want to configure for VPN.
2. Select **Layer3** from the **Interface Type** list.
3. On the **Config** tab, select the **Security Zone** to which the interface belongs:
   - The interface must be accessible from a zone outside of your trust network. Consider creating a dedicated VPN zone for visibility and control over your VPN traffic.
   - If you have not yet created the zone, select **New Zone** from the **Security Zone** list, define a **Name** for the new zone and then click **OK**.
4. Select the **Virtual Router** to use.
5. To assign an IP address to the interface, select the **IPv4** tab, click **Add** in the IP section, and enter the IP address and network mask to assign to the interface, for example 192.168.210.26/24.
6. To save the interface configuration, click **OK**.

   In this example, the configuration for VPN Peer A is:

   - **Interface**—ethernet1/7
   - **Security Zone**—untrust
   - **Virtual Router**—default
   - **IPv4**—100.1.1.1/24

   The configuration for VPN Peer B is:

   - **Interface**—ethernet1/11
   - **Security Zone**—untrust
   - **Virtual Router**—default
   - **IPv4**—200.1.1.1/24

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 2 | Create a tunnel interface and attach it to a virtual router and security zone.

1. Select **Network** > **Interfaces** > **Tunnel** and click **Add**.
2. In the **Interface Name** field, specify a numeric suffix, such as, **.11**.
3. On the **Config** tab, expand **Security Zone** to define the zone as follows:
   - To use your trust zone as the termination point for the tunnel, select the zone.
   - (Recommended) To create a separate zone for VPN tunnel termination, click **New Zone**. In the Zone dialog, define a **Name** for new zone (for example, vpn-tun), and then click **OK**.
4. Select the **Virtual Router**.
5. Assign an IP address to the tunnel interface, select the **IPv4** or **IPv6** tab, click **Add** in the IP section, and enter the IP address and network mask/prefix to assign to the interface, for example, 172.19.9.2/24.

   This IP address will be used as the next hop IP address to route traffic to the tunnel and can also be used to monitor the status of the tunnel.
6. To save the interface configuration, click **OK**.

   In this example, the configuration for VPN Peer A is:

   - **Interface**—tunnel.41
   - **Security Zone**—vpn_tun
   - **Virtual Router**—default
   - **IPv4**—2.1.1.141/24

   The configuration for VPN Peer B is:

   - **Interface**—tunnel.40
   - **Security Zone**—vpn_tun
   - **Virtual Router**—default
   - **IPv4**—2.1.1.140/24

STEP 3 | Set up the Crypto profiles (IKE Crypto profile for phase 1 and IPSec Crypto profile for phase 2).

Complete this task on both peers and make sure to set identical values.

1. Select **Network** > **Network Profiles** > **IKE Crypto**. In this example, we use the default profile.
2. Select **Network** > **Network Profiles** > **IPSec Crypto**. In this example, we use the default profile.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 4 | Set up the OSPF configuration on the virtual router and attach the OSPF areas with the appropriate interfaces on the firewall.

For more information on the OSPF options that are available on the firewall, see Configure OSPF.

Use Broadcast as the link type when there are more than two OSPF routers that need to exchange routing information.

1. Select **Network** > **Virtual Routers**, and select the default router or add a new router.
2. Select **OSPF** (for IPv4) or **OSPFv3** (for IPv6) and select **Enable**.
3. In this example, the OSPF configuration for VPN Peer A is:
   - **Router ID**: 192.168.100.141
   - **Area ID**: 0.0.0.0 that is assigned to the tunnel.1 interface with Link type: p2p
   - **Area ID**: 0.0.0.10 that is assigned to the interface Ethernet1/1 and Link Type: Broadcast

   The OSPF configuration for VPN Peer B is:
   - **Router ID**: 192.168.100.140
   - **Area ID**: 0.0.0.0 that is assigned to the tunnel.1 interface with Link type: p2p
   - **Area ID**: 0.0.0.20 that is assigned to the interface Ethernet1/15 and Link Type: Broadcast

STEP 5 | Set up the IKE Gateway.

This examples uses static IP addresses for both VPN peers. Typically, the corporate office uses a statically configured IP address, and the branch side can be a dynamic IP address; dynamic IP addresses are not best suited for configuring stable services such as VPN.

1. Select **Network** > **Network Profiles** > **IKE Gateway**.
2. Click **Add** and configure the options in the **General** tab.

   In this example, the configuration for VPN Peer A is:
   - **Interface**—ethernet1/7
   - **Local IP address**—100.1.1.1/24
   - **Peer IP address**—200.1.1.1/24
   - **Preshared keys**—enter a value

   The configuration for VPN Peer B is:
   - **Interface**—ethernet1/11
   - **Local IP address**—200.1.1.1/24
   - **Peer IP address**—100.1.1.1/24
   - **Preshared keys**—enter same value as on Peer A
3. Select the IKE Crypto profile you created earlier to use for IKE phase 1.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 6 | Set up the IPSec Tunnel.

1. Select **Network** > **IPSec Tunnels**.
2. Click **Add** and configure the options in the **General** tab.

   In this example, the configuration for VPN Peer A is:

   - **Tunnel Interface**—tunnel.41
   - **Type**—Auto Key
   - **IKE Gateway**—Select the IKE Gateway defined above.
   - **IPSec Crypto Profile**—Select the IKE Gateway defined above.

   The configuration for VPN Peer B is:

   - **Tunnel Interface**—tunnel.40
   - **Type**—Auto Key
   - **IKE Gateway**—Select the IKE Gateway defined above.
   - **IPSec Crypto Profile**—Select the IKE Gateway defined above.
3. Select **Show Advanced Options**, select **Tunnel Monitor**, and specify a Destination IP address to ping for verifying connectivity.
4. To define the action on failure to establish connectivity, see Define a Tunnel Monitoring Profile.

STEP 7 | Create policies to allow traffic between the sites (subnets).

1. Select **Policies** > **Security**.
2. Create rules to allow traffic between the untrust and the vpn-tun zone and the vpn-tun and the untrust zone for traffic originating from specified source and destination IP addresses.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 8 | Verify OSPF adjacencies and routes from the CLI.

Verify that both the firewalls can see each other as neighbors with full status. Also confirm that the IP address of the VPN peer's tunnel interface and the OSPF Router ID. Use the following CLI commands on each VPN peer.

- **show routing protocol ospf neighbor**

```
admin@FW-A> show routing protocol ospf neighbor

  Options: 0x80:reserved, O:Opaq-LSA capability, DC:demand circuits, EA:Ext-Attr LSA capability,
           N/P:NSSA option, MC:multicase, E:AS external LSA capability, T:TOS capability
  ==========
  virtual router:               vr1
  neighbor address:             2.1.1.140
  local address binding:        0.0.0.0
  type:                         dynamic
  status:                       full
  neighbor router ID:           192.168.100.140
  area id:                      0.0.0.0
  neighbor priority:            1
  lifetime remain:              39
  messages pending:             0
  LSA request pending:          0
  options:                      0x42: O E
  hello suppressed:             no

admin@FW-B> show routing protocol ospf neighbor

  Options: 0x80:reserved, O:Opaq-LSA capability, DC:demand circuits, EA:Ext-Attr LSA capability,
           N/P:NSSA option, MC:multicase, E:AS external LSA capability, T:TOS capability
  ==========
  virtual router:               vr1
  neighbor address:             2.1.1.141
  local address binding:        0.0.0.0
  type:                         dynamic
  status:                       full
  neighbor router ID:           192.168.100.141
  area id:                      0.0.0.0
  neighbor priority:            1
  lifetime remain:              39
  messages pending:             0
  LSA request pending:          0
  options:                      0x42: O E
  hello suppressed:             no
```

- **show routing route type ospf**

```
admin@FW-A> show routing route type ospf

flags: A:active, ?:loose, C:connect, H:host, S:static, ~:internal, R:rip, O:ospf, B:bgp,
       Oi:ospf intra-area, Oo:ospf inter-area, O1:ospf ext-type-1, O2:ospf ext-type-2


VIRTUAL ROUTER: vr1 (id 1)
  ==========
destination             nexthop             metric flags      age   interface          next-AS
2.1.1.0/24              0.0.0.0                 10   Oi       6760  tunnel.41
172.16.101.0/24         0.0.0.0                 10   Oi       6854  ethernet1/1
192.168.1.0/24          2.1.1.140               20 A Oo       6754  tunnel.40
total routes shown: 3
    admin@FW-B> show routing route type ospf

    flags: A:active, C:connect, H:host, S:static, R:rip, O:ospf,
           Oi:ospf intra-area, Oo:ospf inter-area, O1:ospf ext-type-1, O2:ospf ext-type-2


    VIRTUAL ROUTER: vr1 (id 1)
      ==========
    destination             nexthop             metric flags      age   interface
    2.1.1.0/24              0.0.0.0                 10   Oi       20033 tunnel.40
    172.16.101.0/24         2.1.1.141               20   AOo      6896  tunnel.40
    192.168.1.0/24          0.0.0.0                 10   Oi       8058  ethernet1/15
    total routes shown: 3
```


https://t.me/learningnets

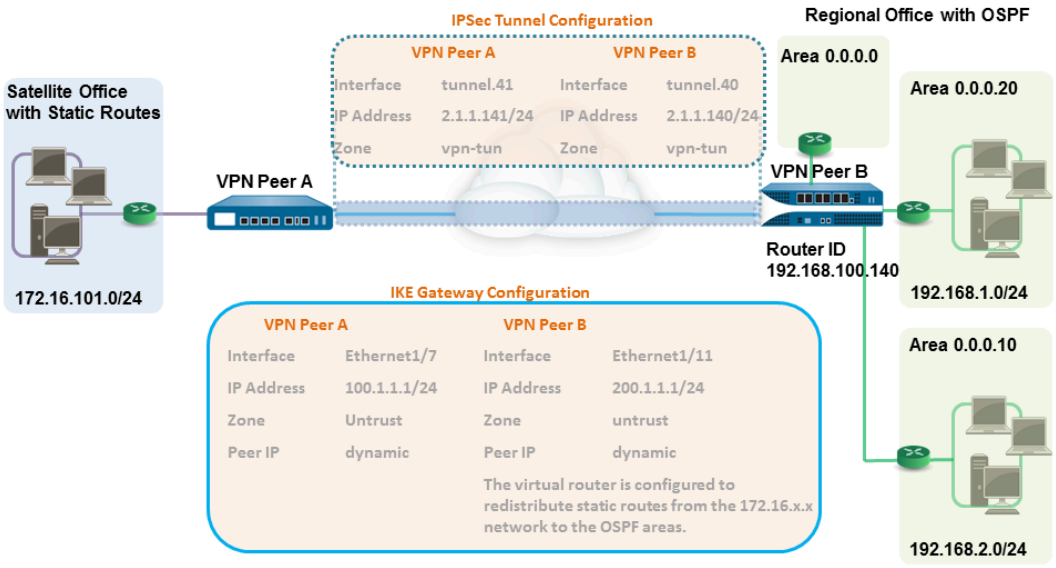
STEP 9 | Test VPN Connectivity.

See Set Up Tunnel Monitoring and View the Status of the Tunnels.

Site-to-Site VPN with Static and Dynamic Routing

In this example, one site uses static routes and the other site uses OSPF. When the routing protocol is not the same between the locations, the tunnel interface on each firewall must be configured with a static IP address. Then, to allow the exchange of routing information, the firewall that participates in both the static and dynamic routing process must be configured with a Redistribution profile. Configuring the redistribution profile enables the virtual router to redistribute and filter routes between protocols—static routes, connected routes, and hosts—from the static autonomous system to the OSPF autonomous system. Without this redistribution profile, each protocol functions on its own and does not exchange any route information with other protocols running on the same virtual router.

In this example, the satellite office has static routes and all traffic destined to the 192.168.x.x network is routed to tunnel.41. The virtual router on VPN Peer B participates in both the static and the dynamic routing process and is configured with a redistribution profile in order to propagate (export) the static routes to the OSPF autonomous system.



 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 1 | Configure the Layer 3 interfaces on each firewall.

1. Select **Network** > **Interfaces** > **Ethernet** and then select the interface you want to configure for VPN.
2. Select **Layer3** from the **Interface Type**.
3. On the **Config** tab, select the **Security Zone** to which the interface belongs:
   - The interface must be accessible from a zone outside of your trust network. Consider creating a dedicated VPN zone for visibility and control over your VPN traffic.
   - If you have not yet created the zone, select **New Zone** from the **Security Zone**, define a **Name** for the new zone and then click **OK**.
4. Select the **Virtual Router** to use.
5. To assign an IP address to the interface, select the **IPv4** tab, click **Add** in the IP section, and enter the IP address and network mask to assign to the interface, for example 192.168.210.26/24.
6. To save the interface configuration, click **OK**.

   In this example, the configuration for VPN Peer A is:

   - **Interface**—ethernet1/7
   - **Security Zone**—untrust
   - **Virtual Router**—default
   - **IPv4**—100.1.1.1/24

   The configuration for VPN Peer B is:

   - **Interface**—ethernet1/11
   - **Security Zone**—untrust
   - **Virtual Router**—default
   - **IPv4**—200.1.1.1/24

STEP 2 | Set up the Crypto profiles (IKE Crypto profile for phase 1 and IPSec Crypto profile for phase 2).

Complete this task on both peers and make sure to set identical values.

1. Select **Network** > **Network Profiles** > **IKE Crypto**. In this example, we use the default profile.
2. Select **Network** > **Network Profiles** > **IPSec Crypto**. In this example, we use the default profile.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 3 | Set up the IKE Gateway.

With pre-shared keys, to add authentication scrutiny when setting up the IKE phase-1 tunnel, you can set up Local and Peer Identification attributes and a corresponding value that is matched in the IKE negotiation process.

1. Select **Network** > **Network Profiles** > **IKE Gateway**.
2. Click **Add** and configure the options in the **General** tab.

   In this example, the configuration for VPN Peer A is:

   - **Interface**—ethernet1/7
   - **Local IP address**—100.1.1.1/24
   - **Peer IP type**—dynamic
   - **Preshared keys**—enter a value
   - **Local identification**—select **FQDN(hostname)** and enter the value for VPN Peer A.
   - **Peer identification**—select **FQDN(hostname)** and enter the value for VPN Peer B

   The configuration for VPN Peer B is:

   - **Interface**—ethernet1/11
   - **Local IP address**—200.1.1.1/24
   - **Peer IP address**—dynamic
   - **Preshared keys**—enter same value as on Peer A
   - **Local identification**—select **FQDN(hostname)** and enter the value for VPN Peer B
   - **Peer identification**—select **FQDN(hostname)** and enter the value for VPN Peer A
3. Select the IKE Crypto profile you created earlier to use for IKE phase 1.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 4 | Create a tunnel interface and attach it to a virtual router and security zone.

1. Select **Network** > **Interfaces** > **Tunnel** and click **Add**.
2. In the **Interface Name** field, specify a numeric suffix, say, **.41**.
3. On the **Config** tab, expand the **Security Zone** to define the zone as follows:
   - To use your trust zone as the termination point for the tunnel, select the zone.
   - (Recommended) To create a separate zone for VPN tunnel termination, click **New Zone**. In the Zone dialog, define a **Name** for new zone (for example *vpn-tun*), and then click **OK**.
4. Select the **Virtual Router**.
5. Assign an IP address to the tunnel interface, select the **IPv4** or **IPv6** tab, click **Add** in the IP section, and enter the IP address and network mask/prefix to assign to the interface, for example, 172.19.9.2/24.

   This IP address will be used to route traffic to the tunnel and to monitor the status of the tunnel.
6. To save the interface configuration, click **OK**.

   In this example, the configuration for VPN Peer A is:
   - **Interface**—tunnel.41
   - **Security Zone**—vpn_tun
   - **Virtual Router**—default
   - **IPv4**—2.1.1.141/24

   The configuration for VPN Peer B is:
   - **Interface**—tunnel.42
   - **Security Zone**—vpn_tun
   - **Virtual Router**—default
   - **IPv4**—2.1.1.140/24

STEP 5 | Specify the interface to route traffic to a destination on the 192.168.x.x network.

1. On VPN Peer A, select the virtual router.
2. Select **Static Routes**, and **Add** tunnel.41 as the **Interface** for routing traffic with a **Destination** in the 192.168.x.x network.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 6 | Set up the static route and the OSPF configuration on the virtual router and attach the OSPF areas with the appropriate interfaces on the firewall.

1. On VPN Peer B, select **Network** > **Virtual Routers**, and select the default router or add a new router.
2. Select **Static Routes** and **Add** the tunnel IP address as the next hop for traffic in the 172.168.x.x. network.

   Assign the desired route metric; using a lower the value makes the a higher priority for route selection in the forwarding table.
3. Select **OSPF** (for IPv4) or **OSPFv3** (for IPv6) and select **Enable**.
4. In this example, the OSPF configuration for VPN Peer B is:
   - Router ID: 192.168.100.140
   - Area ID: 0.0.0.0 is assigned to the interface Ethernet 1/12 Link type: Broadcast
   - Area ID: 0.0.0.10 that is assigned to the interface Ethernet1/1 and Link Type: Broadcast
   - Area ID: 0.0.0.20 is assigned to the interface Ethernet1/15 and Link Type: Broadcast

STEP 7 | Create a redistribution profile to inject the static routes into the OSPF autonomous system.

1. Create a redistribution profile on VPN Peer B.
   1. Select **Network** > **Virtual Routers**, and select the router you used above.
   2. Select **Redistribution Profiles**, and click **Add**.
   3. Enter a Name for the profile and select **Redist** and assign a **Priority** value. If you have configured multiple profiles, the profile with the lowest priority value is matched first.
   4. Set **Source Type** as **static**, and click **OK**. The static route you defined in Step 6 will be used for the redistribution.
2. Inject the static routes in to the OSPF system.
   1. Select **OSPF** > **Export Rules** (for IPv4) or **OSPFv3** > **Export Rules** (for IPv6).
   2. Click **Add**, and select the redistribution profile that you just created.
   3. Select how the external routes are brought into the OSPF system. The default option, **Ext2** calculates the total cost of the route using only the external metrics. To use both internal and external OSPF metrics, use **Ext1**.
   4. Assign a **Metric** (cost value) for the routes injected into the OSPF system. This option allows you to change the metric for the injected route as it comes into the OSPF system.
   5. Click **OK**.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 8 | Set up the IPSec Tunnel.

1. Select **Network** > **IPSec Tunnels**.
2. Click **Add** and configure the options in the **General** tab.

   In this example, the configuration for VPN Peer A is:

   - **Tunnel Interface**—tunnel.41
   - **Type**—Auto Key
   - **IKE Gateway**—Select the IKE Gateway defined above.
   - **IPSec Crypto Profile**—Select the IKE Gateway defined above.

   The configuration for VPN Peer B is:

   - **Tunnel Interface**—tunnel.40
   - **Type**—Auto Key
   - **IKE Gateway**—Select the IKE Gateway defined above.
   - **IPSec Crypto Profile**—Select the IKE Gateway defined above.
3. Select **Show Advanced Options**, select **Tunnel Monitor**, and specify a Destination IP address to ping for verifying connectivity.
4. To define the action on failure to establish connectivity, see Define a Tunnel Monitoring Profile.

STEP 9 | Create policies to allow traffic between the sites (subnets).

1. Select **Policies** > **Security**.
2. Create rules to allow traffic between the untrust and the vpn-tun zone and the vpn-tun and the untrust zone for traffic originating from specified source and destination IP addresses.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 10 | Verify OSPF adjacencies and routes from the CLI.

Verify that both the firewalls can see each other as neighbors with full status. Also confirm that the IP address of the VPN peer's tunnel interface and the OSPF Router ID. Use the following CLI commands on each VPN peer.

- **`show routing protocol ospf neighbor`**

```
admin@FW-A> show routing protocol ospf neighbor

  Options: 0x80:reserved, O:Opaq-LSA capability, DC:demand circuits, EA:Ext-Attr LSA capability,
           N/P:NSSA option, MC:multicase, E:AS external LSA capability, T:TOS capability
  ==========
  virtual router:               vr1
  neighbor address:             2.1.1.140
  local address binding:        0.0.0.0
  type:                         dynamic
  status:                       full
  neighbor router ID:           192.168.100.140
  area id:                      0.0.0.0
  neighbor priority:            1
  lifetime remain:              39
  messages pending:             0
  LSA request pending:          0
  options:                      0x42: O E
  hello suppressed:             no

admin@FW-B> show routing protocol ospf neighbor

  Options: 0x80:reserved, O:Opaq-LSA capability, DC:demand circuits, EA:Ext-Attr LSA capability,
           N/P:NSSA option, MC:multicase, E:AS external LSA capability, T:TOS capability
  ==========
  virtual router:               vr1
  neighbor address:             2.1.1.141
  local address binding:        0.0.0.0
  type:                         dynamic
  status:                       full
  neighbor router ID:           192.168.100.141
  area id:                      0.0.0.0
  neighbor priority:            1
  lifetime remain:              39
  messages pending:             0
  LSA request pending:          0
  options:                      0x42: O E
  hello suppressed:             no
```

- **`show routing route`**

The following is an example of the output on each VPN peer.

```
VPN PeerA

destination             next hop        metric  flags   age     interface       next-AS
192.168.1.0/24          2.1.1.141       20      A S             tunnel.41
192.168.2.0/24          2.1.1.141       20      A S             tunnel.41
172.16.101.0/24         0.0.0.0         1       A H             ethernet1/1
2.1.1.140/24            2.1.1.141       20      A S             tunnel.41

VPN PeerB

destination             next hop        metric  flags   age     interface       next-AS
192.168.1.0/24          0.0.0.0         10      A Oo            ethernet1/1
192.168.2.0/24          0.0.0.0         10      A Oo            ethernet1/15
172.16.101.0/24         2.1.1.140       20      A H             tunnel.40
2.1.1.141/24            2.1.1.140       10      A C             tunnel.40
```

STEP 11 | Test VPN Connectivity.

See Set Up Tunnel Monitoring and View the Status of the Tunnels.

https://t.me/learningnets

Large Scale VPN (LSVPN)

The GlobalProtect Large Scale VPN (LSVPN) feature on the Palo Alto Networks next-generation firewall simplifies the deployment of traditional hub and spoke VPNs, enabling you to quickly deploy enterprise networks with several branch offices with a minimum amount of configuration required on the remote *satellites*. This solution uses certificates for firewall authentication and IPSec to secure data.

LSVPN enables site-to-site VPNs between Palo Alto Networks firewalls. To set up a site-to-site VPN between a Palo Alto Networks firewall and another device, see VPNs.

The following topics describe the LSVPN components and how to set them up to enable site-to-site VPN services between Palo Alto Networks firewalls:

- LSVPN Overview
- Create Interfaces and Zones for the LSVPN
- Enable SSL Between GlobalProtect LSVPN Components
- Configure the Portal to Authenticate Satellites
- Configure GlobalProtect Gateways for LSVPN
- Configure the GlobalProtect Portal for LSVPN
- Prepare the Satellite to Join the LSVPN
- Verify the LSVPN Configuration
- LSVPN Quick Configs

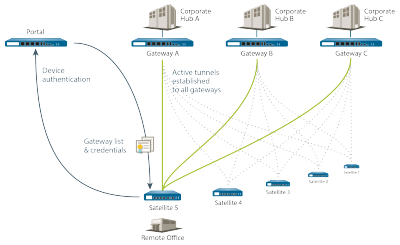
https://t.me/learningnets

LSVPN Overview

GlobalProtect provides a complete infrastructure for managing secure access to corporate resources from your remote sites. This infrastructure includes the following components:

- **GlobalProtect Portal**—Provides the management functions for your GlobalProtect LSVPN infrastructure. Every satellite that participates in the GlobalProtect LSVPN receives configuration information from the portal, including configuration information to enable the satellites (the spokes) to connect to the gateways (the hubs). You configure the portal on an interface on any Palo Alto Networks next-generation firewall.
- **GlobalProtect Gateways**—A Palo Alto Networks firewall that provides the tunnel end point for satellite connections. The resources that the satellites access is protected by security policy on the gateway. It is not required to have a separate portal and gateway; a single firewall can function both as portal and gateway.
- **GlobalProtect Satellite**—A Palo Alto Networks firewall at a remote site that establishes IPSec tunnels with the gateway(s) at your corporate office(s) for secure access to centralized resources. Configuration on the satellite firewall is minimal, enabling you to quickly and easily scale your VPN as you add new sites.

The following diagram illustrates how the GlobalProtect LSVPN components work together.



https://t.me/learningnets

Create Interfaces and Zones for the LSVPN

You must configure the following interfaces and zones for your LSVPN infrastructure:

- **GlobalProtect portal**—Requires a Layer 3 interface for GlobalProtect satellites to connect to. If the portal and gateway are on the same firewall, they can use the same interface. The portal must be in a zone that is accessible from your branch offices.
- **GlobalProtect gateways**—Requires three interfaces: a Layer 3 interface in the zone that is reachable by the remote satellites, an internal interface in the trust zone that connects to the protected resources, and a logical tunnel interface for terminating the VPN tunnels from the satellites. Unlike other site-to-site VPN solutions, the GlobalProtect gateway only requires a single tunnel interface, which it will use for tunnel connections with all of your remote satellites (point-to-multi-point). If you plan to use dynamic routing, you must assign an IP address to the tunnel interface. GlobalProtect supports both IPv6 and IPv4 addressing for the tunnel interface.
- **GlobalProtect satellites**—Requires a single tunnel interface for establishing a VPN with the remote gateways (up to a maximum of 25 gateways). If you plan to use dynamic routing, you must assign an IP address to the tunnel interface. GlobalProtect supports both IPv6 and IPv4 addressing for the tunnel interface.

For more information about portals, gateways, and satellites see LSVPN Overview.

STEP 1 | Configure a Layer 3 interface.

The portal and each gateway and satellite all require a Layer 3 interface to enable traffic to be routed between sites.

If the gateway and portal are on the same firewall, you can use a single interface for both components.

1. Select **Network** > **Interfaces** > **Ethernet** and then select the interface you want to configure for GlobalProtect LSVPN.
2. Select **Layer3** from the **Interface Type** drop-down.
3. On the **Config** tab, select the **Security Zone** to which the interface belongs:
   - The interface must be accessible from a zone outside of your trust network. Consider creating a dedicated VPN zone for visibility and control over your VPN traffic.
   - If you have not yet created the zone, select **New Zone** from the **Security Zone** drop-down, define a **Name** for the new zone and then click **OK**.
4. Select the **Virtual Router** to use.
5. Assign an IP address to the interface:
   - For an IPv4 address, select **IPv4** and **Add** the IP address and network mask to assign to the interface, for example 203.0.11.100/24.
   - For an IPv6 address, select **IPv6**, **Enable IPv6 on the interface**, and **Add** the IP address and network mask to assign to the interface, for example 2001:1890:12f2:11::10.1.8.160/80.
6. To save the interface configuration, click **OK**.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 2 | On the firewall(s) hosting GlobalProtect gateway(s), configure the logical tunnel interface that will terminate VPN tunnels established by the GlobalProtect satellites.

IP addresses are not required on the tunnel interface unless you plan to use dynamic routing. However, assigning an IP address to the tunnel interface can be useful for troubleshooting connectivity issues.

Make sure to enable User-ID in the zone where the VPN tunnels terminate.

1. Select **Network** > **Interfaces** > **Tunnel** and click **Add**.
2. In the **Interface Name** field, specify a numeric suffix, such as **.2**.
3. On the **Config** tab, expand the **Security Zone** drop-down to define the zone as follows:
   - To use your trust zone as the termination point for the tunnel, select the zone from the drop-down.
   - (Recommended) To create a separate zone for VPN tunnel termination, click **New Zone**. In the Zone dialog, define a **Name** for new zone (for example *lsvpn-tun*), select the **Enable User Identification** check box, and then click **OK**.
4. Select the **Virtual Router**.
5. (Optional) To assign an IP address to the tunnel interface:
   - For an IPv4 address, select **IPv4** and **Add** the IP address and network mask to assign to the interface, for example 203.0.11.100/24.
   - For an IPv6 address, select **IPv6**, **Enable IPv6 on the interface**, and **Add** the IP address and network mask to assign to the interface, for example 2001:1890:12f2:11::10.1.8.160/80.
6. To save the interface configuration, click **OK**.

STEP 3 | If you created a separate zone for tunnel termination of VPN connections, create a security policy to enable traffic flow between the VPN zone and your trust zone.

For example, a policy rule enables traffic between the *lsvpn-tun* zone and the *L3-Trust* zone.

STEP 4 | Commit your changes.

Click **Commit**.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Enable SSL Between GlobalProtect LSVPN Components

All interaction between the GlobalProtect components occurs over an SSL/TLS connection. Therefore, you must generate and/or install the required certificates before configuring each component so that you can reference the appropriate certificate(s) and/or certificate profiles in the configurations for each component. The following sections describe the supported methods of certificate deployment, descriptions and best practice guidelines for the various GlobalProtect certificates, and provide instructions for generating and deploying the required certificates:

- About Certificate Deployment
- Deploy Server Certificates to the GlobalProtect LSVPN Components
- Deploy Client Certificates to the GlobalProtect Satellites Using SCEP

About Certificate Deployment

There are two basic approaches to deploying certificates for GlobalProtect LSVPN:

- **Enterprise Certificate Authority**—If you already have your own enterprise certificate authority, you can use this internal CA to issue an intermediate CA certificate for the GlobalProtect portal to enable it to issue certificates to the GlobalProtect gateways and satellites. You can also configure the GlobalProtect portal to act as a Simple Certificate Enrollment Protocol (SCEP) client to issue client certificates to GlobalProtect satellites.
- **Self-Signed Certificates**—You can generate a self-signed root CA certificate on the firewall and use it to issue server certificates for the portal, gateway(s), and satellite(s). When using self-signed root CA certificates, as a best practice, create a self-signed root CA certificate on the portal and use it to issue server certificates for the gateways and satellites. This way, the private key used for certificate signing stays on the portal.

Deploy Server Certificates to the GlobalProtect LSVPN Components

The GlobalProtect LSVPN components use SSL/TLS to mutually authenticate. Before deploying the LSVPN, you must assign an SSL/TLS service profile to each portal and gateway. The profile specifies the server certificate and allowed TLS versions for communication with satellites. You don't need to create SSL/TLS service profiles for the satellites because the portal will issue a server certificate for each satellite during the first connection as part of the satellite registration process.

In addition, you must import the root certificate authority (CA) certificate used to issue the server certificates onto each firewall that you plan to host as a gateway or satellite. Finally, on each gateway and satellite participating in the LSVPN, you must configure a certificate profile that will enable them to establish an SSL/TLS connection using mutual authentication.

The following workflow shows the best practice steps for deploying SSL certificates to the GlobalProtect LSVPN components:

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 1 | On the firewall hosting the GlobalProtect portal, create the root CA certificate for signing the certificates of the GlobalProtect components.

Create a Self-Signed Root CA Certificate:

1. Select **Device** > **Certificate Management** > **Certificates** > **Device Certificates** and click **Generate**.
2. Enter a **Certificate Name**, such as **LSVPN_CA**.
3. Do not select a value in the **Signed By** field (this is what indicates that it is self-signed).
4. Select the **Certificate Authority** check box and then click **OK** to generate the certificate.

STEP 2 | Create SSL/TLS service profiles for the GlobalProtect portal and gateways.

For the portal and each gateway, you must assign an SSL/TLS service profile that references a unique self-signed server certificate.

The best practice is to issue all of the required certificates on the portal, so that the signing certificate (with the private key) doesn't have to be exported.

If the GlobalProtect portal and gateway are on the same firewall interface, you can use the same server certificate for both components.

1. Use the root CA on the portal to Generate a Certificate for each gateway you will deploy:
    1. Select **Device** > **Certificate Management** > **Certificates** > **Device Certificates** and click **Generate**.
    2. Enter a **Certificate Name**.
    3. Enter the FQDN (recommended) or IP address of the interface where you plan to configure the gateway in the **Common Name** field.
    4. In the **Signed By** field, select the **LSVPN_CA** certificate you just created.
    5. In the Certificate Attributes section, click **Add** and define the attributes to uniquely identify the gateway. If you add a **Host Name** attribute (which populates the SAN field of the certificate), it must exactly match the value you defined for the **Common Name**.
    6. **Generate** the certificate.
2. Configure an SSL/TLS Service Profile for the portal and each gateway:
    1. Select **Device** > **Certificate Management** > **SSL/TLS Service Profile** and click **Add**.
    2. Enter a **Name** to identify the profile and select the server **Certificate** you just created for the portal or gateway.
    3. Define the range of TLS versions (**Min Version** to **Max Version**) allowed for communicating with satellites and click **OK**.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 3 | Deploy the self-signed server certificates to the gateways.

Best Practices:

- Export the self-signed server certificates issued by the root CA from the portal and import them onto the gateways.
- Be sure to issue a unique server certificate for each gateway.
- The Common Name (CN) and, if applicable, the Subject Alternative Name (SAN) fields of the certificate must match the IP address or fully qualified domain name (FQDN) of the interface where you configure the gateway.

1. On the portal, select **Device** > **Certificate Management** > **Certificates** > **Device Certificates**, select the gateway certificate you want to deploy, and click **Export**.
2. Select **Encrypted Private Key and Certificate (PKCS12)** from the **File Format** drop-down.
3. Enter (and re-enter) a **Passphrase** to encrypt the private key associated with the certificate and then click **OK** to download the PKCS12 file to your computer.
4. On the gateway, select **Device** > **Certificate Management** > **Certificates** > **Device Certificates** and click **Import**.
5. Enter a **Certificate Name**.
6. Enter the path and name to the **Certificate File** you just downloaded from the portal, or **Browse** to find the file.
7. Select **Encrypted Private Key and Certificate (PKCS12)** as the **File Format**.
8. Enter the path and name to the PKCS12 file in the **Key File** field or **Browse** to find it.
9. Enter and re-enter the **Passphrase** you used to encrypt the private key when you exported it from the portal and then click **OK** to import the certificate and key.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 4 | Import the root CA certificate used to issue server certificates for the LSVPN components.

You must import the root CA certificate onto all gateways and satellites. For security reasons, make sure you export the certificate only, and not the associated private key.

1. Download the root CA certificate from the portal.
   1. Select **Device** > **Certificate Management** > **Certificates** > **Device Certificates**.
   2. Select the root CA certificate used to issue certificates for the LSVPN components and click **Export**.
   3. Select **Base64 Encoded Certificate (PEM)** from the **File Format** drop-down and click **OK** to download the certificate. (Do not export the private key.)
2. On the firewalls hosting the gateways and satellites, import the root CA certificate.
   1. Select **Device** > **Certificate Management** > **Certificates** > **Device Certificates** and click **Import**.
   2. Enter a **Certificate Name** that identifies the certificate as your client CA certificate.
   3. **Browse** to the **Certificate File** you downloaded from the CA.
   4. Select **Base64 Encoded Certificate (PEM)** as the **File Format** and then click **OK**.
   5. Select the certificate you just imported on the **Device Certificates** tab to open it.
   6. Select **Trusted Root CA** and then click **OK**.
   7. **Commit** the changes.

STEP 5 | Create a certificate profile.

The GlobalProtect LSVPN portal and each gateway require a certificate profile that specifies which certificate to use to authenticate the satellites.

1. Select **Device** > **Certificate Management** > **Certificate Profile** and click **Add** and enter a profile **Name**.
2. Make sure **Username Field** is set to **None**.
3. In the **CA Certificates** field, click **Add**, select the Trusted Root CA certificate you imported in the previous step.
4. (Recommended) Enable use of CRL and/or OCSP to enable certificate status verification.
5. Click **OK** to save the profile.

STEP 6 | Commit your changes.

Click **Commit**.

Deploy Client Certificates to the GlobalProtect Satellites Using SCEP

As an alternative method for deploying client certificates to satellites, you can configure your GlobalProtect portal to act as a Simple Certificate Enrollment Protocol (SCEP) client to a SCEP server in your enterprise PKI. SCEP operation is dynamic in that the enterprise PKI generates a certificate when the portal requests it and sends the certificate to the portal.

When the satellite device requests a connection to the portal or gateway, it also includes its serial number with the connection request. The portal submits a CSR to the SCEP server using

https://t.me/learningnets

the settings in the SCEP profile and automatically includes the serial number of the device in the subject of the client certificate. After receiving the client certificate from the enterprise PKI, the portal transparently deploys the client certificate to the satellite device. The satellite device then presents the client certificate to the portal or gateway for authentication.

STEP 1 | Create a SCEP profile.

1. Select **Device** > **Certificate Management** > **SCEP** and then **Add** a new profile.
2. Enter a **Name** to identify the SCEP profile.
3. If this profile is for a firewall with multiple virtual systems capability, select a virtual system or **Shared** as the **Location** where the profile is available.

STEP 2 | (Optional) To make the SCEP-based certificate generation more secure, configure a SCEP challenge-response mechanism between the PKI and portal for each certificate request.

After you configure this mechanism, its operation is invisible, and no further input from you is necessary.

To comply with the U.S. Federal Information Processing Standard (FIPS), use a **Dynamic** SCEP challenge and specify a **Server URL** that uses HTTPS (see Step 7).

Select one of the following options:

- **None**—(Default) The SCEP server does not challenge the portal before it issues a certificate.
- **Fixed**—Obtain the enrollment challenge password from the SCEP server (for example, **`http://10.200.101.1/CertSrv/mscep_admin/`**) in the PKI infrastructure and then copy or enter the password into the Password field.
- **Dynamic**—Enter the SCEP **Server URL** where the portal-client submits these credentials (for example, **`http://10.200.101.1/CertSrv/mscep_admin/`**), and a username and OTP of your choice. The username and password can be the credentials of the PKI administrator.

STEP 3 | Specify the settings for the connection between the SCEP server and the portal to enable the portal to request and receive client certificates.

To identify the satellite, the portal automatically includes the device serial number in the CSR request to the SCEP server. Because the SCEP profile requires a value in the **Subject** field, you can leave the default **`$USERNAME`** token even though the value is not used in client certificates for LSVPN.

1. Configure the **Server URL** that the portal uses to reach the SCEP server in the PKI (for example, **`http://10.200.101.1/certsrv/mscep/`**).
2. Enter a string (up to 255 characters in length) in the **CA-IDENT Name** field to identify the SCEP server.
3. Select the **Subject Alternative Name Type**:
   - **RFC 822 Name**—Enter the email name in a certificate's subject or Subject Alternative Name extension.
   - **DNS Name**—Enter the DNS name used to evaluate certificates.
   - **Uniform Resource Identifier**—Enter the name of the resource from which the client will obtain the certificate.
   - **None**—Do not specify attributes for the certificate.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 4 | (Optional) Configure cryptographic settings for the certificate.

- Select the key length (**Number of Bits**) for the certificate. If the firewall is in FIPS-CC mode and the key generation algorithm is RSA. The RSA keys must be 2048 bits or larger.
- Select the **Digest for CSR** which indicates the digest algorithm for the certificate signing request (CSR): SHA1, SHA256, SHA384, or SHA512.

STEP 5 | (Optional) Configure the permitted uses of the certificate, either for signing or encryption.

- To use this certificate for signing, select the **Use as digital signature** check box. This enables the endpoint use the private key in the certificate to validate a digital signature.
- To use this certificate for encryption, select the **Use for key encipherment** check box. This enables the client use the private key in the certificate to encrypt data exchanged over the HTTPS connection established with the certificates issued by the SCEP server.

STEP 6 | (Optional) To ensure that the portal is connecting to the correct SCEP server, enter the **CA Certificate Fingerprint**. Obtain this fingerprint from the SCEP server interface in the Thumbprint field.

1. Enter the URL for the SCEP server's administrative UI (for example, **`http://<hostname or IP>/CertSrv/mscep_admin/`**).
2. Copy the thumbprint and enter it in the **CA Certificate Fingerprint** field.

STEP 7 | Enable mutual SSL authentication between the SCEP server and the GlobalProtect portal. This is required to comply with the U.S. Federal Information Processing Standard (FIPS).



FIPS-CC operation is indicated on the firewall login page and in its status bar.

Select the SCEP server's root **CA Certificate**. Optionally, you can enable mutual SSL authentication between the SCEP server and the GlobalProtect portal by selecting a **Client Certificate**.

STEP 8 | Save and commit the configuration.

1. Click **OK** to save the settings and close the SCEP configuration.
2. **Commit** the configuration.

The portal attempts to request a CA certificate using the settings in the SCEP profile and saves it to the firewall hosting the portal. If successful, the CA certificate is shown in **Device** > **Certificate Management** > **Certificates**.

STEP 9 | (Optional) If after saving the SCEP profile, the portal fails to obtain the certificate, you can manually generate a certificate signing request (CSR) from the portal.

1. Select **Device** > **Certificate Management** > **Certificates** > **Device Certificates** and then click **Generate**.
2. Enter a **Certificate Name**. This name cannot contain spaces.
3. Select the **SCEP Profile** to use to submit a CSR to your enterprise PKI.
4. Click **OK** to submit the request and generate the certificate.

 ©2022 Palo Alto Networks, Inc.

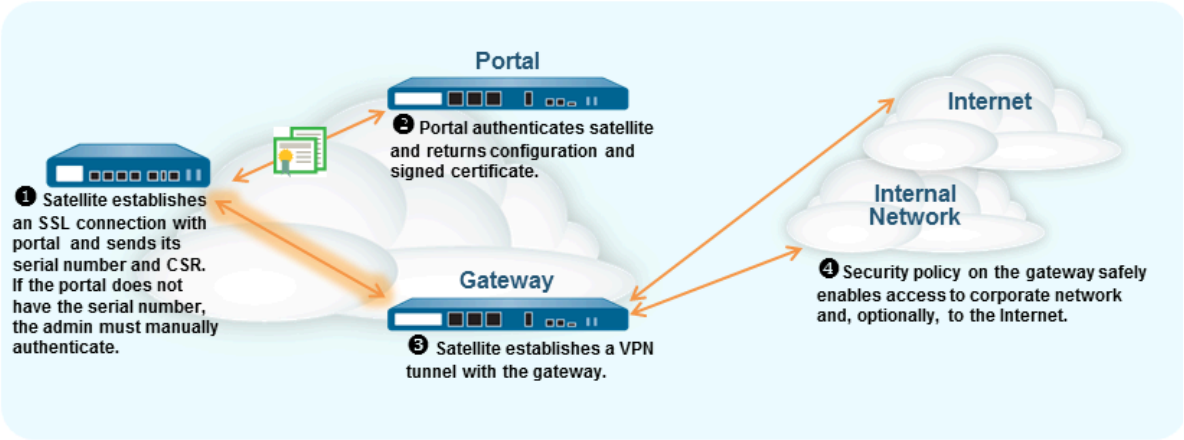
https://t.me/learningnets

Configure the Portal to Authenticate Satellites

In order to register with the LSVPN, each satellite must establish an SSL/TLS connection with the portal. After establishing the connection, the portal authenticates the satellite to ensure that is authorized to join the LSVPN. After successfully authenticating the satellite, the portal will issue a server certificate for the satellite and push the LSVPN configuration specifying the gateways to which the satellite can connect and the root CA certificate required to establish an SSL connection with the gateways.

There are two ways that the satellite can authenticate to the portal during its initial connection:

- **Serial number**—You can configure the portal with the serial number of the satellite firewalls that are authorized to join the LSVPN. During the initial satellite connection to the portal, the satellite presents its serial number to the portal and if the portal has the serial number in its configuration, the satellite will be successfully authenticated. You add the serial numbers of authorized satellites when you configure the portal. See Configure the Portal.
- **Username and password**—If you would rather provision your satellites without manually entering the serial numbers of the satellites into the portal configuration, you can instead require the satellite administrator to authenticate when establishing the initial connection to the portal. Although the portal will always look for the serial number in the initial request from the satellite, if it cannot identify the serial number, the satellite administrator must provide a username and password to authenticate to the portal. Because the portal will always fall back to this form of authentication, you must create an authentication profile in order to commit the portal configuration. This requires that you set up an authentication profile for the portal LSVPN configuration even if you plan to authenticate satellites using the serial number.



The following workflow describes how to set up the portal to authenticate satellites against an existing authentication service. GlobalProtect LSVPN supports external authentication using a local database, LDAP (including Active Directory), Kerberos, TACACS+, or RADIUS.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 1 | (External authentication only) Create a server profile on the portal.

The server profile defines how the firewall connects to an external authentication service to validate the authentication credentials that the satellite administrator enters.

If you use local authentication, skip this step and instead add a local user for the satellite administrator: see Add the user account to the local database.

Configure a server profile for the authentication service type:

- Add a RADIUS server profile.

You can use RADIUS to integrate with a Multi-Factor Authentication *service.*

- Add a TACACS+ server profile.
- Add a SAML IdP server profile.
- Add a Kerberos server profile.
- Add an LDAP server profile. If you use LDAP to connect to Active Directory (AD), create a separate LDAP server profile for every AD domain.

STEP 2 | Configure an authentication profile.

The authentication profile defines which server profile to use to authenticate satellites.

1. Select **Device** > **Authentication Profile** and click **Add**.
2. Enter a **Name** for the profile and then select the authentication **Type**. If the **Type** is an external service, select the **Server Profile** you created in the previous step. If you added a local user instead, set the **Type** to **Local Database**.
3. Click **OK** and **Commit**.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Configure GlobalProtect Gateways for LSVPN

Because the GlobalProtect configuration that the portal delivers to the satellites includes the list of gateways the satellite can connect to, it is a good idea to configure the gateways before configuring the portal.

Before you can configure the GlobalProtect gateway, you must complete the following tasks:

- Create Interfaces and Zones for the LSVPN on the interface where you will configure each gateway. You must configure both the physical interface and the virtual tunnel interface.
- Enable SSL Between GlobalProtect LSVPN Components by configuring the gateway server certificates, SSL/TLS service profiles, and certificate profile required to establish a mutual SSL/TLS connection from the GlobalProtect satellites to the gateway.

Configure each GlobalProtect gateway to participate in the LSVPN as follows:

STEP 1 | Add a gateway.

1. Select **Network** > **GlobalProtect** > **Gateways** and click **Add**.
2. In the **General** screen, enter a **Name** for the gateway. The gateway name should have no spaces and, as a best practice, should include the location or other descriptive information to help users and administrators identify the gateway.
3. (Optional) Select the virtual system to which this gateway belongs from the **Location** field.

STEP 2 | Specify the network information that enables satellite devices to connect to the gateway.

If you haven't created the network interface for the gateway, see Create Interfaces and Zones for the LSVPN for instructions.

1. Select the **Interface** that satellites will use for ingress access to the gateway.
2. Specify the **IP Address Type** and **IP address** for gateway access:
   - The IP address type can be **IPv4** (only), **IPv6** (only), or **IPv4 and IPv6**. Use **IPv4 and IPv6** if your network supports dual stack configurations, where IPv4 and IPv6 run at the same time.
   - The IP address must be compatible with the IP address type. For example, **`172.16.1/0`** for IPv4 addresses or **`21DA:D3:0:2F3B`** for IPv6 addresses. For dual stack configurations, enter both an IPv4 and IPv6 address.
3. Click **OK** to save changes.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 3 | Specify how the gateway authenticates satellites attempting to establish tunnels. If you haven't yet created an SSL/TLS Service profile for the gateway, see Deploy Server Certificates to the GlobalProtect LSVPN Components.

If you haven't set up the authentication profiles or certificate profiles, see Configure the Portal to Authenticate Satellites for instructions.

If you have not yet set up the certificate profile, see Enable SSL Between GlobalProtect LSVPN Components for instructions.

On the GlobalProtect Gateway Configuration dialog, select Authentication and then configure any of the following:

- To secure communication between the gateway and the satellites, select the **SSL/TLS Service Profile** for the gateway.
- To specify the authentication profile to use to authenticate satellites, **Add** a Client Authentication. Then, enter a **Name** to identify the configuration, select **OS**: **Satellite** to apply the configuration to all satellites, and specify the **Authentication Profile** to use to authenticate the satellite. You can also select a **Certificate Profile** for the gateway to use to authenticate satellite devices attempting to establish tunnels.

STEP 4 | Configure the tunnel parameters and enable tunneling.

1. On the GlobalProtect Gateway Configuration dialog, select **Satellite** > **Tunnel Settings**.
2. Select the **Tunnel Configuration** check box to enable tunneling.
3. Select the **Tunnel Interface** you defined to terminate VPN tunnels established by the GlobalProtect satellites when you performed the task to Create Interfaces and Zones for the LSVPN.
4. (Optional) If you want to preserve the Type of Service (ToS) information in the encapsulated packets, select **Copy TOS**.

   If there are multiple sessions inside the tunnel (each with a different TOS value), copying the TOS header can cause the IPSec packets to arrive out of order.

STEP 5 | (Optional) Enable tunnel monitoring.

Tunnel monitoring enables satellites to monitor its gateway tunnel connection, allowing it to failover to a backup gateway if the connection fails. Failover to another gateway is the only type of tunnel monitoring profile supported with LSVPN.

1. Select the **Tunnel Monitoring** check box.
2. Specify the **Destination IP Address** the satellites should use to determine if the gateway is active. You can specify an **IPv4** address, and **IPv6** address, or both. Alternatively, if you configured an IP address for the tunnel interface, you can leave this field blank and the tunnel monitor will instead use the tunnel interface to determine if the connection is active.
3. Select **Failover** from the **Tunnel Monitor Profile** drop-down (this is the only supported tunnel monitor profile for LSVPN).

STEP 6 | Select the IPSec Crypto profile to use when establishing tunnel connections.

The profile specifies the type of IPSec encryption and the authentication method for securing the data that will traverse the tunnel. Because both tunnel endpoints in an LSVPN are trusted

https://t.me/learningnets

firewalls within your organization, you can typically use the default (predefined) profile, which uses ESP as the IPSec protocol, group2 for the DH group, AES-128-CBC for encryption, and SHA-1 for authentication.

In the **IPSec Crypto Profile** drop-down, select **default** to use the predefined profile or select **New IPSec Crypto Profile** to define a new profile. For details on the authentication and encryption options, see Define IPSec Crypto Profiles.

STEP 7 | Configure the network settings to assign the satellites during establishment of the IPSec tunnel.



You can also configure the satellite to push the DNS settings to its local clients by configuring a DHCP server on the firewall hosting the satellite. In this configuration, the satellite will push DNS settings it learns from the gateway to the DHCP clients.

1. On the GlobalProtect Gateway Configuration dialog, select **Satellite** > **Network Settings**.
2. (Optional) If clients local to the satellite need to resolve FQDNs on the corporate network, configure the gateway to push DNS settings to the satellites in one of the following ways:
   - If the gateway has an interface that is configured as a DHCP client, you can set the **Inheritance Source** to that interface and assign the same settings received by the DHCP client to GlobalProtect satellites. You can also inherit the DNS suffix from the same source.
   - Manually define the **Primary DNS**, **Secondary DNS**, and **DNS Suffix** settings to push to the satellites.
3. To specify the **IP Pool** of addresses to assign the tunnel interface on the satellites when the VPN is established, click **Add** and then specify the IP address range(s) to use.
4. To define what destination subnets to route through the tunnel click **Add** in the **Access Route** area and then enter the routes as follows:
   - If you want to route all traffic from the satellites through the tunnel, leave this field blank.

     In this case, all traffic except traffic destined for the local subnet will be tunneled to the gateway.

   - To route only some traffic through the gateway (called *split tunneling*), specify the destination subnets that must be tunneled. In this case, the satellite will route traffic that is not destined for a specified access route using its own routing table. For example, you may choose to only tunnel traffic destined for your corporate network, and use the local satellite to safely enable Internet access.
   - If you want to enable routing between satellites, enter the summary route for the network protected by each satellite.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 8 | (Optional) Define what routes, if any, the gateway will accept from satellites.

By default, the gateway will not add any routes satellites advertise to its routing table. If you do not want the gateway to accept routes from satellites, you do not need to complete this step.

1. To enable the gateway to accept routes advertised by satellites, select **Satellite** > **Route Filter**.
2. Select the **Accept published routes** check box.
3. To filter which of the routes advertised by the satellites to add to the gateway routing table, click **Add** and then define the subnets to include. For example, if all the satellites are configured with subnet 192.168.x.0/24 on the LAN side, configuring a permitted route of 192.168.0.0/16 to enable the gateway to only accept routes from the satellite if it is in the 192.168.0.0/16 subnet.

STEP 9 | Save the gateway configuration.

1. Click **OK** to save the settings and close the GlobalProtect Gateway Configuration dialog.
2. **Commit** the configuration.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Configure the GlobalProtect Portal for LSVPN

The GlobalProtect portal provides the management functions for your GlobalProtect LSVPN. Every satellite system that participates in the LSVPN receives configuration information from the portal, including information about available gateways as well as the certificate it needs in order to connect to the gateways.

The following sections provide procedures for setting up the portal:

- GlobalProtect Portal for LSVPN Prerequisite Tasks
- Configure the Portal
- Define the Satellite Configurations

GlobalProtect Portal for LSVPN Prerequisite Tasks

Before configuring the GlobalProtect portal, you must complete the following tasks:

- ☐ Create Interfaces and Zones for the LSVPN on the interface where you will configure the portal.
- ☐ Enable SSL Between GlobalProtect LSVPN Components by creating an SSL/TLS service profile for the portal server certificate, issuing gateway server certificates, and configuring the portal to issue server certificates for the GlobalProtect satellites.
- ☐ Configure the Portal to Authenticate Satellites by defining the authentication profile that the portal will use to authenticate satellites if the serial number is not available.
- ☐ Configure GlobalProtect Gateways for LSVPN.

Configure the Portal

After you have completed the GlobalProtect Portal for LSVPN Prerequisite Tasks, configure the GlobalProtect portal as follows:

STEP 1 | Add the portal.

1. Select **Network** > **GlobalProtect** > **Portals** and click **Add**.
2. On the **General** tab, enter a **Name** for the portal. The portal name should not contain any spaces.
3. (Optional) Select the virtual system to which this portal belongs from the **Location** field.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 2 | Specify the network information to enable satellites to connect to the portal.

If you haven't yet created the network interface for the portal, see Create Interfaces and Zones for the LSVPN for instructions.

1. Select the **Interface** that satellites will use for ingress access to the portal.
2. Specify the **IP Address Type** and **IP address** for satellite access to the portal:
   - The IP address type can be **IPv4** (for IPv4 traffic only), **IPv6** (for IPv6 traffic only, or **IPv4 and IPv6**. Use **IPv4 and IPv6** if your network supports dual stack configurations, where IPv4 and IPv6 run at the same time.
   - The IP address must be compatible with the IP address type. For example, **`172.16.1/0`** for IPv4 addresses or **`21DA:D3:0:2F3B`** for IPv6 addresses. For dual stack configurations, enter both an IPv4 and IPv6 address.
3. Click **OK** to save changes.

STEP 3 | Specify an SSL/TLS Service profile to use to enable the satellite to establish an SSL/TLS connection to the portal.

If you haven't yet created an SSL/TLS service profile for the portal and issued gateway certificates, see Deploy Server Certificates to the GlobalProtect LSVPN Components.

1. On the GlobalProtect Portal Configuration dialog, select **Authentication**.
2. Select the **SSL/TLS Service Profile**.

STEP 4 | Specify an authentication profile and optional certificate profile for authenticating satellites.

*If the portal can't validate the serial numbers of connecting satellites, it will fall back to the authentication profile. Therefore, before you can save the portal configuration (by clicking **OK**), you must* Configure an authentication profile.

Add a Client Authentication, and then enter a **Name** to identify the configuration, select **OS**: **Satellite** to apply the configuration to all satellites, and specify the **Authentication Profile** to use to authenticate satellite devices. You can also specify a **Certificate Profile** for the portal to use to authenticate satellite devices.

STEP 5 | Continue with defining the configurations to push to the satellites or, if you have already created the satellite configurations, save the portal configuration.

Click **OK** to save the portal configuration or continue to Define the Satellite Configurations.

Define the Satellite Configurations

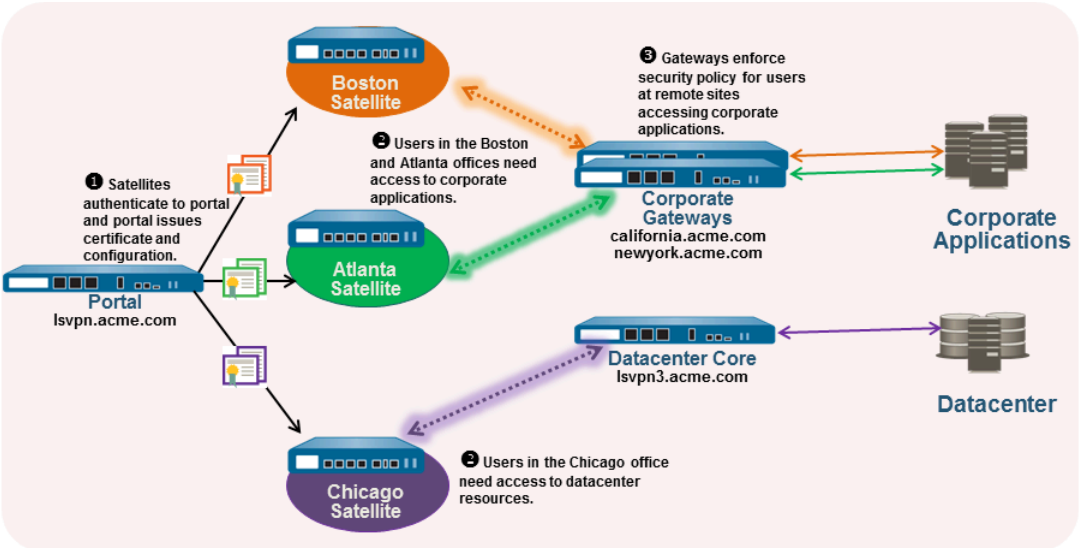
When a GlobalProtect satellite connects and successfully authenticates to the GlobalProtect portal, the portal delivers a satellite configuration, which specifies what gateways the satellite can connect to. If all your satellites will use the same gateway and certificate configurations, you can create a single satellite configuration to deliver to all satellites upon successful authentication. However, if you require different satellite configurations—for example if you want one group of satellites to connect to one gateway and another group of satellites to connect to a different gateway—you can create a separate satellite configuration for each. The portal will then use the enrollment username/group name or the serial number of the satellite to determine which satellite configuration to deploy. As with security rule evaluation, the portal looks for a match starting

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

from the top of the list. When it finds a match, it delivers the corresponding configuration to the satellite.

For example, the following figure shows a network in which some branch offices require VPN access to the corporate applications protected by your perimeter firewalls and another site needs VPN access to the data center.



Use the following procedure to create one or more satellite configurations.

STEP 1 | Add a satellite configuration.

The satellite configuration specifies the GlobalProtect LSVPN configuration settings to deploy to the connecting satellites. You must define at least one satellite configuration.

1. Select **Network** > **GlobalProtect** > **Portals** and select the portal configuration for which you want to add a satellite configuration and then select the **Satellite** tab.
2. In the Satellite section, click **Add**.
3. Enter a **Name** for the configuration.

   If you plan to create multiple configurations, make sure the name you define for each is descriptive enough to allow you to distinguish them.
4. To change how often a satellite should check the portal for configuration updates specify a value in the **Configuration Refresh Interval (hours)** field (range is 1-48; default is 24).

STEP 2 | Specify the satellites to which to deploy this configuration.

The portal uses the **Enrollment User/User Group** settings and/or **Devices** serial numbers to match a satellite to a configuration. Therefore, if you have multiple configurations, be sure to order them properly. As soon as the portal finds a match, it will deliver the configuration. Therefore, more specific configurations must precede more general ones. See Step 5 for instructions on ordering the list of satellite configurations.

Specify the match criteria for the satellite configuration as follows:

- To restrict this configuration to satellites with specific serial numbers, select the **Devices** tab, click **Add**, and enter serial number (you do not need to enter the satellite hostname; it

https://t.me/learningnets

will be automatically added when the satellite connects). Repeat this step for each satellite you want to receive this configuration.

- Select the **Enrollment User/User Group** tab, click **Add**, and then select the user or group you want to receive this configuration. Satellites that do not match on serial number will be required to authenticate as a user specified here (either an individual user or group member).

Before you can restrict the configuration to specific groups, you must Map Users to Groups.

STEP 3 | Specify the gateways that satellites with this configuration can establish VPN tunnels with.

Routes published by the gateway are installed on the satellite as static routes. The metric for the static route is 10x the routing priority. If you have more than one gateway, make sure to also set the routing priority to ensure that routes advertised by backup gateways have higher metrics compared to the same routes advertised by primary gateways. For example, if you set the routing priority for the primary gateway and backup gateway to 1 and 10 respectively, the satellite will use 10 as the metric for the primary gateway and 100 as the metric for the backup gateway.

1. On the **Gateways** tab, click **Add**.
2. Enter a descriptive **Name** for the gateway. The name you enter here should match the name you defined when you configured the gateway and should be descriptive enough identify the location of the gateway.
3. Enter the FQDN or IP address of the interface where the gateway is configured in the **Gateways** field. The address you specify must exactly match the Common Name (CN) in the gateway server certificate.
4. (Optional) If you are adding two or more gateways to the configuration, the **Routing Priority** helps the satellite pick the preferred gateway. Enter a value in the range of 1-25, with lower numbers having the higher priority (that is, the gateway the satellite will connect to if all gateways are available). The satellite will multiply the routing priority by 10 to determine the routing metric.

STEP 4 | Save the satellite configuration.

1. Click **OK** to save the satellite configuration.
2. If you want to add another satellite configuration, repeat the previous steps.

STEP 5 | Arrange the satellite configurations so that the proper configuration is deployed to each satellite.

- To move a satellite configuration up on the list of configurations, select the configuration and click **Move Up**.
- To move a satellite configuration down on the list of configurations, select the configuration and click **Move Down**.

STEP 6 | Specify the certificates required to enable satellites to participate in the LSVPN.

1. In the **Trusted Root CA** field, click **Add** and then select the CA certificate used to issue the gateway server certificates. The portal will deploy the root CA certificate you add

https://t.me/learningnets

here to all satellites as part of the configuration to enable the satellite to establish an SSL connection with the gateways. As a best practice, all of your gateways should use the same issuer.

2. Select the method of **Client Certificate** distribution:

   - **To store the client certificates on the portal**—select **Local** and select the Root CA certificate that the portal will use to issue client certificates to satellites upon successfully authenticating them from the **Issuing Certificate** drop-down.

     *If the root CA certificate used to issue your gateway server certificates is not on the portal, you can **Import** it now. See Enable SSL Between GlobalProtect LSVPN Components for details on how to import a root CA certificate.*

   - **To enable the portal to act as a SCEP client to dynamically request and issue client certificates**—select **SCEP** and then select the **SCEP** profile used to generate CSRs to your SCEP server.

     *If the you have not yet set up the portal to act as a SCEP client, you can add a **New** SCEP profile now. See Deploy Client Certificates to the GlobalProtect Satellites Using SCEP for details.*

STEP 7 | Save the portal configuration.

1. Click **OK** to save the settings and close the GlobalProtect Portal Configuration dialog.
2. **Commit** your changes.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Prepare the Satellite to Join the LSVPN

To participate in the LSVPN, the satellites require a minimal amount of configuration. Because the required configuration is minimal, you can pre-configure the satellites before shipping them to your branch offices for installation.

STEP 1 | Configure a Layer 3 Interface.

This is the physical interface the satellite will use to connect to the portal and the gateway. This interface must be in a zone that allows access outside of the local trust network. As a best practice, create a dedicated zone for VPN connections for visibility and control over traffic destined for the corporate gateways.

STEP 2 | Configure the logical tunnel interface for the tunnel to use to establish VPN tunnels with the GlobalProtect gateways.



IP addresses are not required on the tunnel interface unless you plan to use dynamic routing. However, assigning an IP address to the tunnel interface can be useful for troubleshooting connectivity issues.

1. Select **Network** > **Interfaces** > **Tunnel** and click **Add**.
2. In the **Interface Name** field, specify a numeric suffix, such as **.2**.
3. On the **Config** tab, expand the **Security Zone** drop-down and select an existing zone or create a separate zone for VPN tunnel traffic by clicking **New Zone** and defining a **Name** for new zone (for example *lsvpnsat*).
4. In the **Virtual Router** drop-down, select **default**.
5. (Optional) To assign an IP address to the tunnel interface:
   - For an IPv4 address, select **IPv4** and **Add** the IP address and network mask to assign to the interface, for example 203.0.11.100/24.
   - For an IPv6 address, select **IPv6**, **Enable IPv6 on the interface**, and **Add** the IP address and network mask to assign to the interface, for example 2001:1890:12f2:11::10.1.8.160/80.
6. To save the interface configuration, click **OK**.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 3 | If you generated the portal server certificate using a Root CA that is not trusted by the satellites (for example, if you used self-signed certificates), import the root CA certificate used to issue the portal server certificate.

The root CA certificate is required to enable the satellite to establish the initial connection with the portal to obtain the LSVPN configuration.

1. Download the CA certificate that was used to generate the portal server certificates. If you are using self-signed certificates, export the root CA certificate from the portal as follows:
    1. Select **Device** > **Certificate Management** > **Certificates** > **Device Certificates**.
    2. Select the CA certificate, and click **Export**.
    3. Select **Base64 Encoded Certificate (PEM)** from the **File Format** drop-down and click **OK** to download the certificate. (You do not need to export the private key.)
2. Import the root CA certificate you just exported onto each satellite as follows.
    1. Select **Device** > **Certificate Management** > **Certificates** > **Device Certificates** and click **Import**.
    2. Enter a **Certificate Name** that identifies the certificate as your client CA certificate.
    3. **Browse** to the **Certificate File** you downloaded from the CA.
    4. Select **Base64 Encoded Certificate (PEM)** as the **File Format** and then click **OK**.
    5. Select the certificate you just imported on the **Device Certificates** tab to open it.
    6. Select **Trusted Root CA** and then click **OK**.

STEP 4 | Configure the IPSec tunnel configuration.

1. Select **Network** > **IPSec Tunnels** and click **Add**.
2. On the **General** tab, enter a descriptive **Name** for the IPSec configuration.
3. Select the **Tunnel Interface** you created for the satellite.
4. Select **GlobalProtect Satellite** as the **Type**.
5. Enter the IP address or FQDN of the portal as the **Portal Address**.
6. Select the Layer 3 **Interface** you configured for the satellite.
7. Select the **IP Address** to use on the selected interface. You can select an **IPv4** address, an **IPv6** address, or both. Specify if you want **IPv6 preferred for portal registration**.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 5 | (Optional) Configure the satellite to publish local routes to the gateway.

Pushing routes to the gateway enables traffic to the subnets local to the satellite via the gateway. However, you must also configure the gateway to accept the routes as detailed in Configure GlobalProtect Gateways for LSVPN.

1. To enable the satellite to push routes to the gateway, on the **Advanced** tab select **Publish all static and connected routes to Gateway**.

   If you select this check box, the firewall will forward all static and connected routes from the satellite to the gateway. However, to prevent the creation of routing loops, the firewall will apply some route filters, such as the following:

   - Default routes
   - Routes within a virtual router other than the virtual router associated with the tunnel interface
   - Routes using the tunnel interface
   - Routes using the physical interface associated with the tunnel interface
2. (Optional) If you only want to push routes for specific subnets rather than all routes, click **Add** in the Subnet section and specify which subnet routes to publish.

STEP 6 | Save the satellite configuration.

1. Click **OK** to save the IPSec tunnel settings.
2. Click **Commit**.

STEP 7 | If required, provide the credentials to allow the satellite to authenticate to the portal.

This step is only required if the portal was unable to find a serial number match in its configuration or if the serial number didn't work. In this case, the satellite will not be able to establish the tunnel with the gateway(s).

1. Select **Network** > **IPSec Tunnels** and click the **Gateway Info** link in the Status column of the tunnel configuration you created for the LSVPN.
2. Click the **enter credentials** link in the **Portal Status** field and username and password required to authenticate the satellite to the portal.

   After the portal successfully authenticates to the portal, it will receive its signed certificate and configuration, which it will use to connect to the gateway(s). You should see the tunnel establish and the **Status** change to **Active**.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Verify the LSVPN Configuration

After configuring the portal, gateways, and satellites, verify that the satellites are able to connect to the portal and gateway and establish VPN tunnels with the gateway(s).

STEP 1 | Verify satellite connectivity with portal.

From the firewall hosting the portal, verify that satellites are successfully connecting by selecting **Network** > **GlobalProtect** > **Portal** and clicking **Satellite Info** in the Info column of the portal configuration entry.

STEP 2 | Verify satellite connectivity with the gateway(s).

On each firewall hosting a gateway, verify that satellites are able to establish VPN tunnels by selecting **Network** > **GlobalProtect** > **Gateways** and click **Satellite Info** in the Info column of the gateway configuration entry. Satellites that have successfully established tunnels with the gateway will display on the **Active Satellites** tab.

STEP 3 | Verify LSVPN tunnel status on the satellite.

On each firewall hosting a satellite, verify the tunnel status by selecting **Network** > **IPSec Tunnels** and verify active Status as indicated by a green icon.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

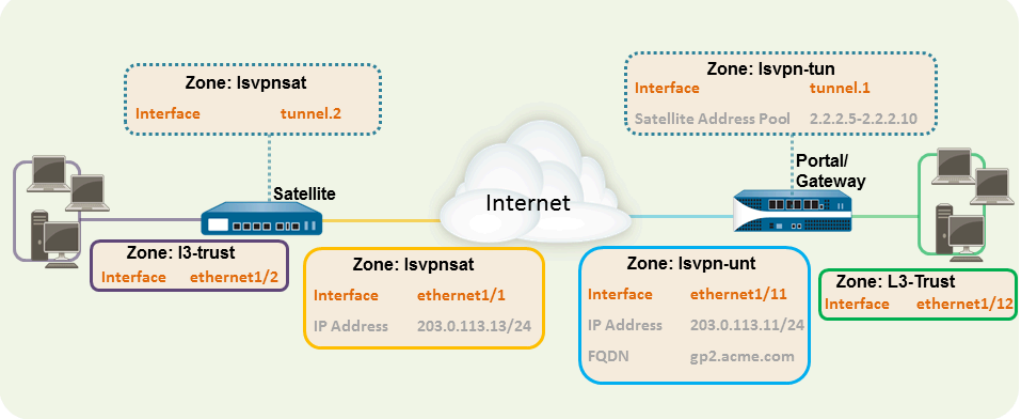
LSVPN Quick Configs

The following sections provide step-by-step instructions for configuring some common GlobalProtect LSVPN deployments:

- Basic LSVPN Configuration with Static Routing
- Advanced LSVPN Configuration with Dynamic Routing
- Advanced LSVPN Configuration with iBGP

Basic LSVPN Configuration with Static Routing

This quick config shows the fastest way to get up and running with LSVPN. In this example, a single firewall at the corporate headquarters site is configured as both a portal and a gateway. Satellites can be quickly and easily deployed with minimal configuration for optimized scalability.



The following workflow shows the steps for setting up this basic configuration:

STEP 1 | Configure a Layer 3 interface.

In this example, the Layer 3 interface on the portal/gateway requires the following configuration:

- **Interface**—ethernet1/11
- **Security Zone**—lsvpn-tun
- **IPv4**—203.0.113.11/24

STEP 2 | On the firewall(s) hosting GlobalProtect gateway(s), configure the logical tunnel interface that will terminate VPN tunnels established by the GlobalProtect satellites.

To enable visibility into users and groups connecting over the VPN, enable User-ID in the zone where the VPN tunnels terminate.

In this example, the Tunnel interface on the portal/gateway requires the following configuration:

- **Interface**—tunnel.1
- **Security Zone**—lsvpn-tun

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 3 | Create the Security policy rule to enable traffic flow between the VPN zone where the tunnel terminates (lsvpn-tun) and the trust zone where the corporate applications reside (L3-Trust).

See Create a Security Policy Rule.

STEP 4 | Assign an SSL/TLS Service profile to the portal/gateway. The profile must reference a self-signed server certificate.

The certificate subject name must match the FQDN or IP address of the Layer 3 interface you create for the portal/gateway.

1. On the firewall hosting the GlobalProtect portal, create the root CA certificate for signing the certificates of the GlobalProtect components. In this example, the root CA certificate, **lsvpn-CA**, will be used to issue the server certificate for the portal/gateway. In addition, the portal will use this root CA certificate to sign the CSRs from the satellites.
2. Create SSL/TLS service profiles for the GlobalProtect portal and gateways.

   Because the portal and gateway are on the same interface in this example, they can share an SSL/TLS Service profile that uses the same server certificate. In this example, the profile is named **lsvpnserver**.

STEP 5 | Create a certificate profile.

In this example, the certificate profile **lsvpn-profile** references the root CA certificate **lsvpn-CA**. The gateway will use this certificate profile to authenticate satellites attempting to establish VPN tunnels.

STEP 6 | Configure an authentication profile for the portal to use if the satellite serial number is not available.

1. Create one type of server profile on the portal:
   - Add a RADIUS server profile.

     You can use RADIUS to integrate with a Multi-Factor Authentication *service.*
   - Add a TACACS+ server profile.
   - Add a SAML IdP server profile.
   - Add a Kerberos server profile.
   - Add an LDAP server profile. If you use LDAP to connect to Active Directory (AD), create a separate LDAP server profile for every AD domain.
2. Configure an authentication profile. In this example, the profile **lsvpn-sat** is used to authenticate satellites.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 7 | Configure GlobalProtect Gateways for LSVPN.

Select **Network** > **GlobalProtect** > **Gateways** and **Add** a configuration. This example requires the following gateway configuration:

- **Interface**—ethernet1/11
- **IP Address**—203.0.113.11/24
- **SSL/TLS Server Profile**—lsvpnserver
- **Certificate Profile**—lsvpn-profile
- **Tunnel Interface**—tunnel.1
- **Primary DNS**/**Secondary DNS**—4.2.2.1/4.2.2.2
- **IP Pool**—2.2.2.111-2.2.2.120
- **Access Route**—10.2.10.0/24

STEP 8 | Configure the Portal.

Select **Network** > **GlobalProtect** > **Portal** and **Add** a configuration. This example requires the following portal configuration:

- **Interface**—ethernet1/11
- **IP Address**—203.0.113.11/24
- **SSL/TLS Server Profile**—lsvpnserver
- **Authentication Profile**—lsvpn-sat

STEP 9 | Define the Satellite Configurations.

On the **Satellite** tab in the portal configuration, **Add** a Satellite configuration and a Trusted Root CA and specify the CA the portal will use to issue certificates for the satellites. In this example the required settings are as following:

- **Gateway**—203.0.113.11
- **Issuing Certificate**—lsvpn-CA
- **Trusted Root CA**—lsvpn-CA

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 10 | Prepare the Satellite to Join the LSVPN.

The satellite configuration in this example requires the following settings:

Interface Configuration

- Layer 3 interface—ethernet1/1, 203.0.113.13/24
- Tunnel interface—tunnel.2
- Zone—lsvpnsat

Root CA Certificate from Portal

- lsvpn-CA

IPSec Tunnel Configuration

- **Tunnel Interface**—tunnel.2
- **Portal Address**—203.0.113.11
- **Interface**—ethernet1/1
- **Local IP Address**—203.0.113.13/24
- **Publish all static and connected routes to Gateway**—enabled

Advanced LSVPN Configuration with Dynamic Routing

In larger LSVPN deployments with multiple gateways and many satellites, investing a little more time in the initial configuration to set up dynamic routing will simplify the maintenance of gateway configurations because access routes will update dynamically. The following example configuration shows how to extend the basic LSVPN configuration to configure OSPF as the dynamic routing protocol.

Setting up an LSVPN to use OSPF for dynamic routing requires the following additional steps on the gateways and the satellites:

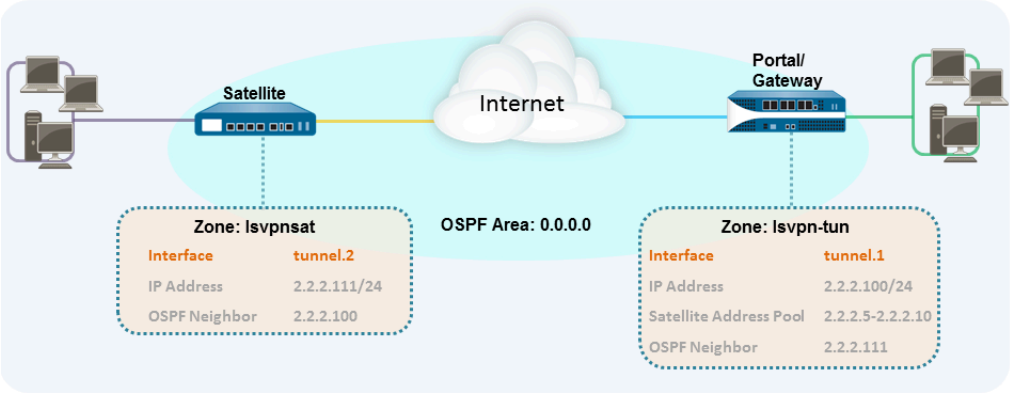
- Manual assignment of IP addresses to tunnel interfaces on all gateways and satellites.
- Configuration of OSPF point-to-multipoint (P2MP) on the virtual router on all gateways and satellites. In addition, as part of the OSPF configuration on each gateway, you must manually define the tunnel IP address of each satellite as an OSPF neighbor. Similarly, on each satellite, you must manually define the tunnel IP address of each gateway as an OSPF neighbor.

Although dynamic routing requires additional setup during the initial configuration of the LSVPN, it reduces the maintenance tasks associated with keeping routes up to date as topology changes occur on your network.

The following figure shows an LSVPN dynamic routing configuration. This example shows how to configure OSPF as the dynamic routing protocol for the VPN.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets



For a basic setup of a LSVPN, follow the steps in Basic LSVPN Configuration with Static Routing. You can then complete the steps in the following workflow to extend the configuration to use dynamic routing rather than static routing.

STEP 1 | Add an IP address to the tunnel interface configuration on each gateway and each satellite.

Complete the following steps on each gateway and each satellite:

1. Select **Network** > **Interfaces** > **Tunnel** and select the tunnel configuration you created for the LSVPN to open the Tunnel Interface dialog.

   If you have not yet created the tunnel interface, see Step 2 in Create Interfaces and Zones for the LSVPN.
2. On the **IPv4** tab, click **Add** and then enter an IP address and subnet mask. For example, to add an IP address for the gateway tunnel interface you would enter 2.2.2.100/24.
3. Click **OK** to save the configuration.

STEP 2 | Configure the dynamic routing protocol on the gateway.

To configure OSPF on the gateway:

1. Select **Network** > **Virtual Routers** and select the virtual router associated with your VPN interfaces.
2. On the **Areas** tab, click **Add** to create the backbone area, or, if it is already configured, click on the area ID to edit it.
3. If you are creating a new area, enter an **Area ID** on the **Type** tab.
4. On the **Interface** tab, click **Add** and select the tunnel **Interface** you created for the LSVPN.
5. Select **p2mp** as the **Link Type**.
6. Click **Add** in the Neighbors section and enter the IP address of the tunnel interface of each satellite, for example 2.2.2.111.
7. Click **OK** twice to save the virtual router configuration and then **Commit** the changes on the gateway.
8. Repeat this step each time you add a new satellite to the LSVPN.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 3 | Configure the dynamic routing protocol on the satellite.

To configure OSPF on the satellite:

1. Select **Network** > **Virtual Routers** and select the virtual router associated with your VPN interfaces.
2. On the **Areas** tab, click **Add** to create the backbone area, or, if it is already configured, click on the area ID to edit it.
3. If you are creating a new area, enter an **Area ID** on the **Type** tab.
4. On the **Interface** tab, click **Add** and select the tunnel **Interface** you created for the LSVPN.
5. Select **p2mp** as the **Link Type**.
6. Click **Add** in the Neighbors section and enter the IP address of the tunnel interface of each GlobalProtect gateway, for example 2.2.2.100.
7. Click **OK** twice to save the virtual router configuration and then **Commit** the changes on the gateway.
8. Repeat this step each time you add a new gateway.

STEP 4 | Verify that the gateways and satellites are able to form router adjacencies.

- On each satellite and each gateway, confirm that peer adjacencies have formed and that routing table entries have been created for the peers (that is, the satellites have routes to the gateways and the gateways have routes to the satellites). Select **Network** > **Virtual Router** and click the **More Runtime Stats** link for the virtual router you are using for the LSVPN. On the Routing tab, verify that the LSVPN peer has a route.
- On the **OSPF** > **Interface** tab, verify that the **Type** is **p2mp**.
- On the **OSPF** > **Neighbor** tab, verify that the firewalls hosting your gateways have established router adjacencies with the firewalls hosting your satellites and vice versa. Also verify that the **Status** is **Full**, indicating that full adjacencies have been established.

Advanced LSVPN Configuration with iBGP

This use case illustrates how GlobalProtect LSVPN securely connects distributed office locations with primary and disaster recovery data centers that house critical applications for users and how internal border gateway protocol (iBGP) eases deployment and upkeep. Using this method, you can extend up to 500 satellite offices connecting to a single gateway.

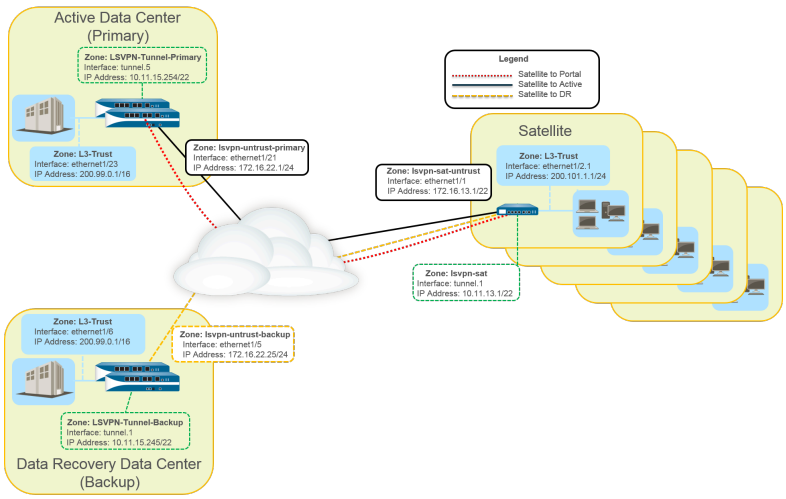
BGP is a highly scalable, dynamic routing protocol that is ideal for hub-and-spoke deployments such as LSVPN. As a dynamic routing protocol, it eliminates much of the overhead associated with access routes (static routes) by making it relatively easy to deploy additional satellite firewalls. Due to its route filtering capabilities and features such as multiple tunable timers, route dampening, and route refresh, BGP scales to a much higher number of routing prefixes with greater stability than other routing protocols like RIP and OSPF. In the case of iBGP, a peer group, which includes all the satellites and gateways in the LSVPN deployment, establishes adjacencies over the tunnel endpoints. The protocol then implicitly takes control of route advertisements, updates, and convergence.

In this example configuration, an active/passive HA pair of PA-5200 firewalls is deployed in the primary (active) data center and acts as the portal and primary gateway. The disaster recovery data

https://t.me/learningnets

center also has two PA-5200s in an active/passive HA pair acting as the backup LSVPN gateway. The portal and gateways serve 500 PA-220s deployed as LSVPN satellites in branch offices.

Both data center sites advertise routes but with different metrics. As a result, the satellites prefer and install the active data center's routes. However, the backup routes also exist in the local routing information base (RIB). If the active data center fails, the routes advertised by that data center are removed and replaced with routes from the disaster recovery data center's routes. The failover time depends on selection of iBGP times and routing convergence associated with iBGP.



The following workflow shows the steps for configuring this deployment:

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 1 | Create Interfaces and Zones for the LSVPN.

Portal and Primary gateway:

- **Zone**: LSVPN-Untrust-Primary
- Interface: ethernet1/21
- IPv4: 172.16.22.1/24
- **Zone**: L3-Trust
- **Interface**: ethernet1/23
- **IPv4**: 200.99.0.1/16

Backup gateway:

- **Zone**: LSVPN-Untrust-Primary
- **Interface**: ethernet1/5
- **IPv4**: 172.16.22.25/24
- **Zone**: L3-Trust
- **Interface**: ethernet1/6
- **IPv4**: 200.99.0.1/16

Satellite:

- **Zone**: LSVPN-Sat-Untrust
- **Interface**: ethernet1/1
- **IPv4**: 172.16.13.1/22
- **Zone**: L3-Trust
- **Interface**: ethernet1/2.1
- **IPv4**: 200.101.1.1/24



Configure the zones, interfaces, and IP addresses on each satellite. The interface and local IP address will be different for each satellite. This interface is used for the VPN connection to the portal and gateway.

STEP 2 | On the firewall(s) hosting GlobalProtect gateway(s), configure the logical tunnel interface that will terminate VPN tunnels established by the GlobalProtect satellites.

Primary gateway:

- **Interface**: tunnel.5
- **IPv4**: 10.11.15.254/22
- **Zone**: LSVPN-Tunnel-Primary

Backup gateway:

- **Interface**: tunnel.1
- **IPv4**: 10.11.15.245/22
- **Zone**: LSVPN-Tunnel-Backup

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 3 | Enable SSL Between GlobalProtect LSVPN Components.

The gateway uses the self-signed root certificate authority (CA) to issue certificates for the satellites in a GlobalProtect LSVPN. Because one firewall houses the portal and primary gateway, a single certificate is used for authenticating to the satellites. The same CA is used to generate a certificate for the backup gateway. The CA generates certificates that pushed to the satellites from the portal and then used by the satellites to authenticate to the gateways.

You must also generate a certificate from the same CA for the backup gateway, allowing it to authenticate with the satellites.

1. On the firewall hosting the GlobalProtect portal, create the root CA certificate for signing the certificates of the GlobalProtect components. In this example, the root CA certificate is called CA-cert.
2. Create SSL/TLS service profiles for the GlobalProtect portal and gateways. Because the GlobalProtect portal and primary gateway are the same firewall interface, you can use the same server certificate for both components.
   - **Root CA Certificate**: CA-Cert
   - **Certificate Name**: LSVPN-Scale
3. Deploy the self-signed server certificates to the gateways.
4. Import the root CA certificate used to issue server certificates for the LSVPN components.
5. Create a certificate profile.
6. Repeat steps 2 through 5 on the backup gateway with the following settings:
   - **Root CA Certificate**: CA-cert
   - **Certificate Name**: LSVPN-back-GW-cert

STEP 4 | Configure GlobalProtect Gateways for LSVPN.

1. Select **Network** > **GlobalProtect** > **Gateways** and click **Add**.
2. On the **General** tab, name the primary gateway **`LSVPN-Scale`**.
3. Under **Network Settings**, select **`ethernet1/21`** as the primary gateway interface and enter **`172.16.22.1/24`** as the IP address.
4. On the **Authentication** tab, select the LSVPN-Scale certificate created in 3.
5. Select **Satellite** > **Tunnel Settings** and select **Tunnel Configuration**. Set the **Tunnel Interface** to tunnel.5. All satellites in this use case connect to a single gateway, so a

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

single satellite configuration is needed. Satellites are matched based on their serial numbers, so no satellites will need to authenticate as a user.

6. On **Satellite** > **Network Settings**, define the pool of IP address to assign to the tunnel interface on the satellite once the VPN connection is established. Because this use case uses dynamic routing, the Access Routes setting remains blank.
7. Repeat steps 1 through 5 on the backup gateway with the following settings:
   - **Name**: LSVPN-backup
   - **Gateway interface**: ethernet1/5
   - **Gateway IP**: 172.16.22.25/24
   - **Server cert**: LSVPN-backup-GW-cert
   - **Tunnel interface**: tunnel.1

STEP 5 | Configure iBGP on the primary and backup gateways and add a redistribution profile to allow the satellites to inject local routes back to the gateways.

Each satellite office manages its own network and firewall, so the redistribution profile called ToAllSat is configured to redistribute local routes back to the GlobalProtect gateway.

1. Select **Network** > **Virtual Routers** and **Add** a virtual router.
2. On **Router Settings**, add the **Name** and **Interface** for the virtual router.
3. On **Redistribution Profile** and select **Add**.
   1. Name the redistribution profile **`ToAllSat`** and set the **Priority** to 1.
   2. Set Redistribute to **Redist**.
   3. **Add `ethernet1/23`** from the Interface drop-down.
   4. Click **OK**.
4. Select **BGP** on the Virtual Router to configure BGP.
   1. On **BGP** > **General**, select **Enable**.
   2. Enter the gateway IP address as the **Router ID** (**`172.16.22.1`**) and **`1000`** as the **AS Number**.
   3. In the Options section, select **Install Route**.
   4. On **BGP** > **Peer Group**, click **Add** a peer group with all the satellites that will connect to the gateway.
   5. On **BGP** > **Redist Rules**, **Add** the **`ToAllSat`** redistribution profile you created previously.
5. Click **OK**.
6. Repeat steps 1 through 5 on the backup gateway using **`ethernet1/6`** for the redistribution profile.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 6 | Prepare the Satellite to Join the LSVPN.

The configuration shown is a sample of a single satellite.

Repeat this configuration each time you add a new satellite to the LSVPN deployment.

1. Configure a tunnel interface as the tunnel endpoint for the VPN connection to the gateways.
2. Set the IPSec tunnel type to GlobalProtect Satellite and enter the IP address of the GlobalProtect Portal.
3. Select **Network** > **Virtual Routers** and **Add** a virtual router.
4. On **Router Settings**, add the **Name** and **Interface** for the virtual router.
5. Select **Virtual Router** > **Redistribution Profile** and **Add** a profile with the following settings.
   1. Name the redistribution profile **`ToLSVPNGW`** and set the **Priority** to 1.
   2. **Add** an **Interface** **`ethernet1/2.1`**.
   3. Click **OK**.
6. Select **BGP** > **General**, **Enable** BGP and configure the protocol as follows:
   1. Enter the gateway IP address as the **Router ID** (**`172.16.22.1`**) and **`1000`** as the **AS Number**.
   2. In the Options section, select **Install Route**.
   3. On **BGP** > **Peer Group**, **Add** a peer group containing all the satellites that will connect to the gateway.
   4. On **BGP** > **Redist Rules**, **Add** the **`ToLSVPNGW`** redistribution profile you created previously.
7. Click **OK**.

STEP 7 | Configure the GlobalProtect Portal for LSVPN.

Both data centers advertise their routes but with different routing priorities to ensure that the active data center is the preferred gateway.

1. Select **Network** > **GlobalProtect** > **Portals** and click **Add**.
2. On **General**, enter **`LSVPN-Portal`** as the portal name.
3. On **Network Settings**, select **`ethernet1/21`** as the **Interface** and select **`172.16.22.1/24`** as the **IP Address**.
4. On the **Authentication** tab, select the previously created primary gateway SSL/TLS Profile **`LSVPN-Scale`** from the **SSL/TLS Service Profile** drop-down menu.
5. On the **Satellite** tab, **Add** a satellite and **Name** it **`sat-config-1`**.
6. Set the **Configuration Refresh Interval** to **12**.
7. On **GlobalProtect Satellite** > **Devices**, add the serial number and hostname of each satellite device in the LSVPN.
8. On **GlobalProtect Satellite** > **Gateways**, add the name and IP address of each gateway. Set the routing priority of the primary gateway to 1 and the backup gateway to 10 to ensure that the active data center is the preferred gateway.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 8 | Verify the LSVPN Configuration.

STEP 9 | (Optional) Add a new site to the LSVPN deployment.

1. Select **Network** > **GlobalProtect** > **Portals** > **GlobalProtect Portal** > **Satellite Configuration** > **GlobalProtect Satellite** > **Devices** to add the serial number of the new satellite to the GlobalProtect portal.
2. Configure the IPSec tunnel on the satellite with the GlobalProtect Portal IP address.
3. Select **Network** > **Virtual Router** > **BGP** > **Peer Group** to add the satellite to the BGP Peer Group configuration on each gateway.
4. Select **Network** > **Virtual Router** > **BGP** > **Peer Group** to add the gateways to the BGP Peer Group configuration on the new satellite.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Policy

Policies allow you to enforce rules and take action. The different types of policy rules that you can create on the firewall are: Security, NAT, Quality of Service (QoS), Policy Based Forwarding (PBF), Decryption, Application Override, Authentication, Denial of Service (DoS), and Zone protection policies. All these different policies work together to allow, deny, prioritize, forward, encrypt, decrypt, make exceptions, authenticate access, and reset connections as needed to help secure your network. The following topics describe how to work with policy:

- Policy Types
- Security Policy
- Policy Objects
- Security Profiles
- Track Rules Within a Rulebase
- Enforce Policy Rule Description, Tag, and Audit Comment
- Move or Clone a Policy Rule or Object to a Different Virtual System
- Use an Address Object to Represent IP Addresses
- Use Tags to Group and Visually Distinguish Objects
- Use an External Dynamic List in Policy
- Register IP Addresses and Tags Dynamically
- Use Dynamic User Groups in Policy
- Use Auto-Tagging to Automate Security Actions
- Monitor Changes in the Virtual Environment
- CLI Commands for Dynamic IP Addresses and Tags
- Identify Users Connected through a Proxy Server
- Policy-Based Forwarding
- Test Policy Rules

https://t.me/learningnets

Policy Types

The Palo Alto Networks next-generation firewall supports a variety of policy types that work together to safely enable applications on your network.

For all policy types, when you Enforce Policy Rule Description, Tag, and Audit Comment, you can use the audit comment archive to view how a policy rule changed over time. The archive, which includes the audit comment history and the configuration logs, enables you to compare configuration versions and review who created or modified and why.

Policy Type	Description
Security	Determine whether to block or allow a session based on traffic attributes such as the source and destination security zone, the source and destination IP address, the application, user, and the service. For more details, see Security Policy.
NAT	Instruct the firewall which packets need translation and how to do the translation. The firewall supports both source address and/or port translation and destination address and/or port translation. For details, see NAT.
QoS	Identify traffic requiring QoS treatment (either preferential treatment or bandwidth-limiting) using a defined parameter or multiple parameters and assign it a class. For more details, see Quality of Service.
Policy Based Forwarding	Identify traffic that should use a different egress interface than the one that would normally be used based on the routing table. For more details, see Policy-Based Forwarding.
Decryption	Identify encrypted traffic that you want to inspect for visibility, control, and granular security. For more details, see Decryption.
Application Override	Identify sessions that you do not want processed by the App-ID engine, which is a Layer-7 inspection. Traffic matching an application override policy forces the firewall to handle the session as a regular stateful inspection firewall at Layer-4. For more details, see Manage Custom or Unknown Applications.
Authentication	Identify traffic that requires users to authenticate. For more details, see Authentication Policy.
DoS Protection	Identify potential denial-of-service (DoS) attacks and take protective action in response to rule matches. For more details, see DoS Protection Profiles.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Security Policy

Security policy protects network assets from threats and disruptions and helps to optimally allocate network resources for enhancing productivity and efficiency in business processes. On a Palo Alto Networks firewall, individual Security policy rules determine whether to block or allow a session based on traffic attributes, such as the source and destination security zone, the source and destination IP address, the application, the user, and the service.



To ensure that end users authenticate when they try to access your network resources, the firewall evaluates Authentication Policy *before Security policy.*

All traffic passing through the firewall is matched against a session and each session is matched against a Security policy rule. When a session match occurs, the firewall applies the matching Security policy rule to bidirectional traffic in that session (client to server and server to client). For traffic that doesn't match any defined rules, the default rules apply. The default rules—displayed at the bottom of the security rulebase—are predefined to allow all intrazone traffic (within a zone) and deny all interzone traffic (between zones). Although these rules are part of the predefined configuration and are read-only by default, you can override them and change a limited number of settings, including the tags, action (allow or block), log settings, and security profiles.

Security policy rules are evaluated left to right and from top to bottom. A packet is matched against the first rule that meets the defined criteria and, after a match is triggered, subsequent rules are not evaluated. Therefore, the more specific rules must precede more generic ones in order to enforce the best match criteria. Traffic that matches a rule generates a log entry at the end of the session in the traffic log if you enable logging for that rule. The logging options are configurable for each rule and can, for example, be configured to log at the start of a session instead of, or in addition to, logging at the end of a session.

After an administrator configures a rule, you can View Policy Rule Usage to determine when and how many times traffic matches the Security policy rule to determine its effectiveness. As your rulebase evolves, change and audit information get lost over time unless you archived this information at the time the rule is created or modified. You can Enforce Policy Rule Description, Tag, and Audit Comment to ensure that all administrators enter audit comments so that you can view the audit comment archive and review comments and configuration log history and can compare rule configuration versions for a selected rule. Together, you now have more visibility into and control over the rulebase.

- Components of a Security Policy Rule
- Security Policy Actions
- Create a Security Policy Rule

Components of a Security Policy Rule

The Security policy rule construct permits a combination of the required and optional fields as detailed in the following table:

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Required Optional	Field	Description
Required	**Name**	A label (up to 63 characters) that identifies the rule.
	UUID	The Universally Unique Identifier (UUID) is a distinct 32-character string that permanently identifies rules so that you can track a rule regardless of any changes to it, such as the name.
	Rule Type	Specifies whether the rule applies to traffic within a zone, between zones, or both: • **universal** (default)—Applies the rule to all matching interzone and intrazone traffic in the specified source and destination zones. For example, if you create a universal rule with source zones A and B and destination zones A and B, the rule would apply to all traffic within zone A, all traffic within zone B, and all traffic from zone A to zone B and all traffic from zone B to zone A. • **intrazone**—Applies the rule to all matching traffic within the specified source zones (you cannot specify a destination zone for intrazone rules). For example, if you set the source zone to A and B, the rule would apply to all traffic within zone A and all traffic within zone B, but not to traffic between zones A and B. • **interzone**—Applies the rule to all matching traffic between the specified source and destination zones. For example, if you set the source zone to A, B, and C and the destination zone to A and B, the rule would apply to traffic from zone A to zone B, from zone B to zone A, from zone C to zone A, and from zone C to zone B, but not traffic within zones A, B, or C.
	Source Zone	The zone from which the traffic originates.
	Destination Zone	The zone at which the traffic terminates. If you use NAT, make sure to always reference the post-NAT zone.
	Application	The application that you wish to control. The firewall uses App-ID, the traffic classification technology, to identify traffic on your network. App-ID provides application control and visibility in creating security policies that block unknown applications, while enabling, inspecting, and shaping those that are allowed.
	Action	Specifies an *Allow* or *Deny* action for the traffic based on the criteria you define in the rule. When you configure the firewall to deny traffic, it either resets the connection or silently drops packets. To provide a better user experience, you can configure granular options to deny traffic instead of silently dropping packets, which can cause some applications to break and appear

https://t.me/learningnets

Required Optional	Field	Description
		unresponsive to the user. For more details, see Security Policy Actions.
Optional	**Tag**	A keyword or phrase that allows you to filter security rules. This is handy when you have defined many rules and wish to then review those that are tagged with a keyword such as *IT-sanctioned applications* or *High-risk applications*.
	Description	A text field, up to 1024 characters, used to describe the rule.
	Source Address	Define host IP addresses, subnets, address objects (of type IP netmask, IP range, FQDN, or IP wildcard mask), address groups, or country-based enforcement. If you use NAT, make sure to always refer to the original IP addresses in the packet (i.e. the pre-NAT IP address).
	Destination Address	The location or destination for the packet. Define IP addresses, subnets, address objects (of type IP netmask, IP range, FQDN, or IP wildcard mask), address groups, or country-based enforcement. If you use NAT, make sure to always refer to the original IP addresses in the packet (i.e. the pre-NAT IP address).
	User	The user or group of users for whom the policy applies. You must have User-ID enabled on the zone. To enable User-ID, see User-ID Overview.
	URL Category	Using the URL Category as match criteria allows you to customize security profiles (Antivirus, Anti-Spyware, Vulnerability, File-Blocking, Data Filtering, and DoS) on a per-URL-category basis. For example, you can prevent.exe file download/upload for URL categories that represent higher risk while allowing them for other categories. This functionality also allows you to attach schedules to specific URL categories (allow social-media websites during lunch & after-hours), mark certain URL categories with QoS (financial, medical, and business), and select different log forwarding profiles on a per-URL-category-basis. Although you can manually configure URL categories on your firewall, to take advantage of the dynamic URL categorization updates available on Palo Alto Networks firewalls, you must purchase a URL filtering license.

©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Required Optional	Field	Description
		To block or allow traffic based on URL category, you must apply a URL Filtering profile to the security policy rules. Define the URL Category as Any and attach a URL Filtering profile to the security policy. See Set Up a Basic Security Policy for information on using the default profiles in your security policy.
	Service	Allows you to select a Layer 4 (TCP or UDP) port for the application. You can choose *any*, specify a port, or use *application-default* to permit use of the standards-based port for the application. For example, for applications with well-known port numbers such as DNS, the *application-default* option will match against DNS traffic only on TCP port 53. You can also add a custom application and define the ports that the application can use. *For inbound allow rules (for example, from untrust to trust), using application-default prevents applications from running on unusual ports and protocols. Application-default is the default option; while the firewall still checks for all applications on all ports, with this configuration, applications are only allowed on their standard ports/protocols.*
	Security Profiles	Provide additional protection from threats, vulnerabilities, and data leaks. Security profiles are evaluated only for rules that have an *allow* action.
	HIP Profile (for GlobalProtect)	Allows you to identify clients with Host Information Profile (HIP) and then enforce access privileges.
	Options	Allow you to define logging for the session, log forwarding settings, change Quality of Service (QoS) markings for packets that match the rule, and schedule when (day and time) the security rule should be in effect.

Security Policy Actions

For traffic that matches the attributes defined in a security policy, you can apply the following actions:

Action	Description
Allow (default)	Allows the traffic.

https://t.me/learningnets

Action	Description
Deny	Blocks traffic and enforces the default *Deny Action* defined for the application that is being denied. To view the deny action defined by default for an application, view the application details in **Objects** > **Applications** or check the application details in Applipedia.
Drop	Silently drops the traffic; for an application, it overrides the default deny action. A TCP reset is not sent to the host/application. For Layer 3 interfaces, to optionally send an ICMP unreachable response to the client, set Action: **Drop** and enable the **Send ICMP Unreachable** check box. When enabled, the firewall sends the ICMP code for *communication with the destination is administratively prohibited*—ICMPv4: Type 3, Code 13; ICMPv6: Type 1, Code 1.
Reset client	Sends a TCP reset to the client-side device.
Reset server	Sends a TCP reset to the server-side device.
Reset both	Sends a TCP reset to both the client-side and server-side devices.

A reset is sent only after a session is formed. If the session is blocked before a 3-way handshake is completed, the firewall will not send the reset.

For a TCP session with a reset action, the firewall does not send an ICMP Unreachable response.

For a UDP session with a drop or reset action, if the **ICMP Unreachable** check box is selected, the firewall sends an ICMP message to the client.

Create a Security Policy Rule

STEP 1 | (Optional) Delete the default Security policy rule.

By default, the firewall includes a security rule named *rule1* that allows all traffic from Trust zone to Untrust zone. You can either delete the rule or modify the rule to reflect your zone naming conventions.

STEP 2 | Add a rule.

1. Select **Policies** > **Security** and **Add** a new rule.
2. In the **General** tab, enter a descriptive **Name** for the rule.
3. Select a **Rule Type**.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 3 | Define the matching criteria for the source fields in the packet.

1. In the **Source** tab, select a **Source Zone**.
2. Specify a **Source IP Address** or leave the value set to **any**.

   *If you decide to **Negate** a region as a **Source Address**, ensure that all regions that contain private IP addresses are added to the **Source Address** to avoid connectivity loss between those private IP addresses.*
3. Specify a Source **User** or leave the value set to **any**.

STEP 4 | Define the matching criteria for the destination fields in the packet.

1. In the **Destination** tab, set the **Destination Zone**.
2. Specify a **Destination IP Address** or leave the value set to **any**.

   *If you decide to **Negate** a region as the **Destination Address**, ensure that all regions that contain private IP addresses are added to the **Destination Address** to avoid connectivity loss between those private IP addresses.*

   *As a best practice, use address objects as the **Destination Address** to enable access to only specific servers or specific groups of servers especially for commonly exploited services, such as DNS and SMTP. By restricting users to specific destination server addresses, you can prevent data exfiltration and command-and-control traffic from establishing communication through techniques such as DNS tunneling.*

STEP 5 | Specify the application that the rule will allow or block.

As a best practice, always use application-based security policy rules instead of port-based rules and always set the Service to application-default unless you are using a more restrictive list of ports than the standard ports for an application.

1. In the **Applications** tab, **Add** the **Application** you want to safely enable. You can select multiple applications or you can use application groups or application filters.
2. In the **Service/URL Category** tab, keep the Service set to **application-default** to ensure that any applications that the rule allows are allowed only on their standard ports.

STEP 6 | (Optional) Specify a URL category as match criteria for the rule.

In the **Service/URL Category** tab, select the **URL Category**.

If you select a URL category, only web traffic will match the rule and only if the traffic is destined for that specified category.

STEP 7 | Define what action you want the firewall to take for traffic that matches the rule.

In the **Actions** tab, select an **Action**. See Security Policy Actions for a description of each action.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 8 | Configure the log settings.

- By default, the rule is set to **Log at Session End**. You can disable this setting if you don't want any logs generated when traffic matches this rule or you can select **Log at Session Start** for more detailed logging.
- Select a **Log Forwarding** profile.



*As a best practice, do not select the check box to **Disable Server Response Inspection** (DSRI). Selecting this option prevents the firewall from inspecting packets from the server to the client. For the best security posture, the firewall must inspect both the client-to-server flows and the server-to-client flows to detect and prevent threats.*

STEP 9 | Attach security profiles to enable the firewall to scan all allowed traffic for threats.



Make sure you create best practice security profiles that help protect your network from both known and unknown threats.

In the **Actions** tab, select **Profiles** from the **Profile Type** drop-down and then select the individual security profiles to attach to the rule.

Alternatively, select **Group** from the **Profile Type** drop-down and select a security **Group Profile** to attach.

STEP 10 | Click **Commit** to save the policy rule to the running configuration on the firewall.

 ©2022 Palo Alto Networks, Inc.

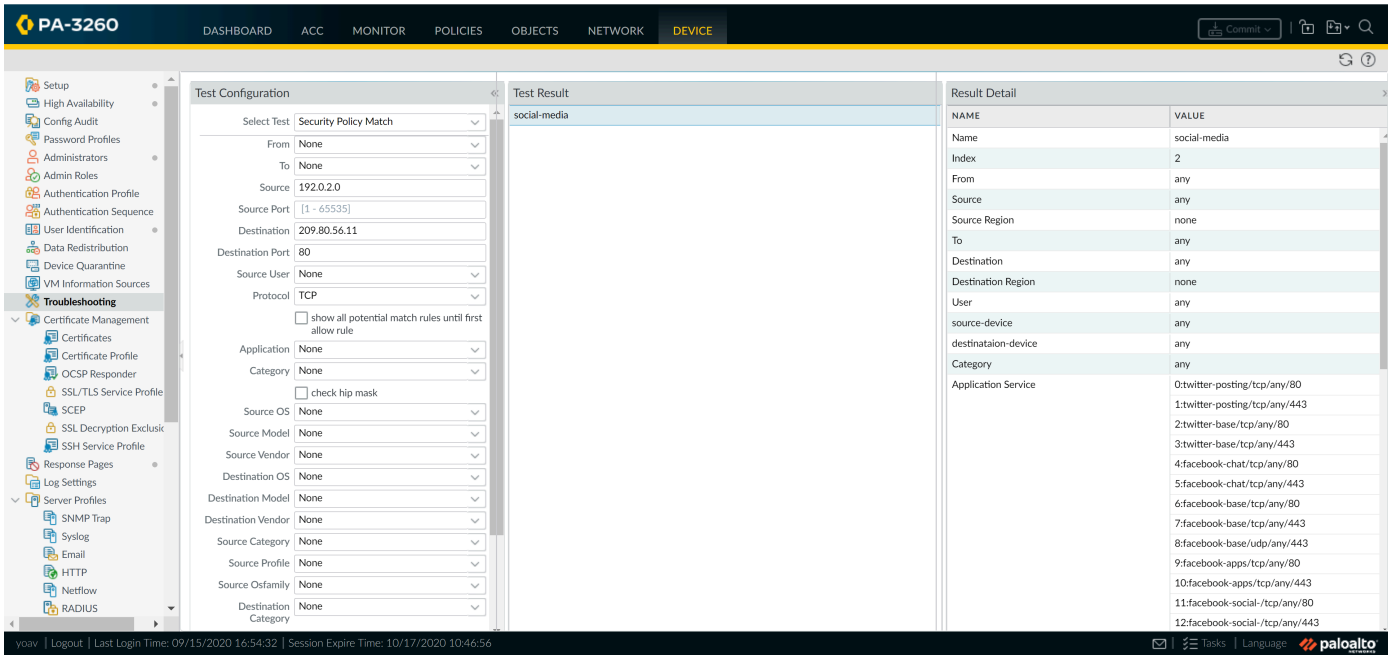
https://t.me/learningnets

STEP 11 | To verify that you have set up your basic security policies effectively, test whether your security policy rules are being evaluated and determine which security policy rule applies to a traffic flow.

The output displays the best rule that matches the source and destination IP address specified in the CLI command.

For example, to verify the policy rule that will be applied for a server in the data center with the IP address 208.90.56.11 when it accesses the Microsoft update server:

1. Select **Device** > **Troubleshooting**, and select **Security Policy Match** from the Select Test drop-down.
2. Enter the Source and Destination IP addresses.
3. Enter the Protocol.
4. **Execute** the security policy match test.



STEP 12 | After waiting long enough to allow traffic to pass through the firewall, View Policy Rule Usage to monitor the policy rule usage status and determine the effectiveness of the policy rule.

https://t.me/learningnets

Policy Objects

A *policy object* is a single object or a collective unit that groups discrete identities such as IP addresses, URLs, applications, or users. With policy objects that are a collective unit, you can reference the object in security policy instead of manually selecting multiple objects one at a time. Typically, when creating a policy object, you group objects that require similar permissions in policy. For example, if your organization uses a set of server IP addresses for authenticating users, you can group the set of server IP addresses as an *address group* policy object and reference the address group in the security policy. By grouping objects, you can significantly reduce the administrative overhead in creating policies.

If you need to export specific parts of the configuration for internal review or audit, you can Export Configuration Table Data as a PDF or CSV file.

You can create the following policy objects on the firewall:

Policy Object	Description
Address/Address Group, Region	Allow you to group specific source or destination addresses that require the same policy enforcement. The address object can include an IPv4 or IPv6 address (single IP, range, subnet), an IP wildcard address (IPv4 address/wildcard mask) or the FQDN. Alternatively, a region can be defined by the latitude and longitude coordinates or you can select a country and define an IP address or IP range. You can then group a collection of address objects to create an address group object. You can also use dynamic address groups to dynamically update IP addresses in environments where host IP addresses change frequently. *The predefined External Dynamic Lists (EDLs) on the firewall count toward the maximum number of address objects that a firewall model supports.*
User/User Group	Allow you to create a list of users from the local database, an external database, or match criteria and group them.
Application Group and Application Filter	An Application Filter allows you to filter applications dynamically. It allows you to filter, and save a group of applications using the attributes defined in the application database on the firewall. For example, you can Create an Application Filter by one or more attributes—category, sub-category, technology, risk, characteristics. With an application filter, when a content update occurs, any new applications that match your filter criteria are automatically added to your saved application filter. An Application Group allows you to create a static group of specific applications that you want to group together for a group of users or for

https://t.me/learningnets

Policy Object	Description
	a particular service, or to achieve a particular policy goal. See Create an Application Group.
Service/Service Groups	Allows you to specify the source and destination ports and protocol that a service can use. The firewall includes two pre-defined services —service-http and service-https— that use TCP ports 80 and 8080 for HTTP, and TCP port 443 for HTTPS. You can however, create any custom service on any TCP/UDP port of your choice to restrict application usage to specific ports on your network (in other words, you can define the default port for the application). *To view the standard ports used by an application, in **Objects** > **Applications** search for the application and click the link. A succinct description displays.*

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Security Profiles

While security policy rules enable you to allow or block traffic on your network, security profiles help you define an *allow but scan* rule, which scans allowed applications for threats, such as viruses, malware, spyware, and DDOS attacks. When traffic matches the allow rule defined in the security policy, the security profile(s) that are attached to the rule are applied for further content inspection rules such as antivirus checks and data filtering.



Security profiles are not used in the match criteria of a traffic flow. The security profile is applied to scan traffic after the application or category is allowed by the security policy.

The firewall provides default security profiles that you can use out of the box to begin protecting your network from threats. See Set Up a Basic Security Policy for information on using the default profiles in your security policy. As you get a better understanding about the security needs on your network, see Create Best Practice Security Profiles for the Internet Gateway to learn how you can create custom profiles.



For recommendations on the best-practice settings for security profiles, see Create Best Practice Security Profiles for the Internet Gateway.

You can add security profiles that are commonly applied together to Create a Security Profile Group; this set of profiles can be treated as a unit and added to security policies in one step (or included in security policies by default, if you choose to set up a default security profile group).

Profile Type	Description
Antivirus Profiles	Antivirus profiles protect against viruses, worms, and trojans as well as spyware downloads. Using a stream-based malware prevention engine, which inspects traffic the moment the first packet is received, the Palo Alto Networks antivirus solution can provide protection for clients without significantly impacting the performance of the firewall. This profile scans for a wide variety of malware in executables, PDF files, HTML and JavaScript viruses, including support for scanning inside compressed files and data encoding schemes. If you have enabled Decryption on the firewall, the profile also enables scanning of decrypted content. The default profile inspects all of the listed protocol decoders for viruses, and generates alerts for SMTP, IMAP, and POP3 protocols while blocking for FTP, HTTP, and SMB protocols. You can configure the action for a decoder or Antivirus signature and specify how the firewall responds to a threat event: • **Default**—For each threat signature and Antivirus signature that is defined by Palo Alto Networks, a default action is specified internally. Typically, the default action is an alert or a reset-both. The default action is displayed in parenthesis, for example default (alert) in the threat or Antivirus signature.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Profile Type	Description
	• **Allow**—Permits the application traffic.  *The **Allow** action does not generate logs related to the signatures or profiles.* • **Alert**—Generates an alert for each application traffic flow. The alert is saved in the threat log. • **Drop**—Drops the application traffic. • **Reset Client**—For TCP, resets the client-side connection. For UDP, drops the connection. • **Reset Server**—For TCP, resets the server-side connection. For UDP, drops the connection. • **Reset Both**—For TCP, resets the connection on both client and server ends. For UDP, drops the connection. Customized profiles can be used to minimize antivirus inspection for traffic between trusted security zones, and to maximize the inspection of traffic received from untrusted zones, such as the internet, as well as the traffic sent to highly sensitive destinations, such as server farms. The Palo Alto Networks WildFire system also provides signatures for persistent threats that are more evasive and have not yet been discovered by other antivirus solutions. As threats are discovered by WildFire, signatures are quickly created and then integrated into the standard Antivirus signatures that can be downloaded by Threat Prevention subscribers on a daily basis (sub-hourly for WildFire subscribers).
Anti-Spyware Profiles	Anti-Spyware profiles blocks spyware on compromised hosts from trying to phone-home or beacon out to external command-and-control (C2) servers, allowing you to detect malicious traffic leaving the network from infected clients. You can apply various levels of protection between zones. For example, you may want to have custom Anti-Spyware profiles that minimize inspection between trusted zones, while maximizing inspection on traffic received from an untrusted zone, such as internet-facing zones. When the firewall is managed by a Panorama management server, the ThreatID is mapped to the corresponding custom threat on the firewall to enable the firewall to generate a threat log populated with the configured custom ThreatID. You can define your own custom Anti-Spyware profiles, or choose one of the following predefined profiles when applying Anti-Spyware to a Security policy rule: • **Default**—Uses the default action for every signature, as specified by Palo Alto Networks when the signature is created. • **Strict**—Overrides the default action of critical, high, and medium severity threats to the block action, regardless of the action defined

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Profile Type	Description
	in the signature file. This profile still uses the default action for low and informational severity signatures. When the firewall detects a threat event, you can configure the following actions in an Anti-Spyware profile: • **Default**—For each threat signature and Anti-Spyware signature that is defined by Palo Alto Networks, a default action is specified internally. Typically the default action is an alert or a reset-both. The default action is displayed in parenthesis, for example default (alert) in the threat or Antivirus signature. • **Allow**—Permits the application traffic *The **Allow** action does not generate logs related to the signatures or profiles.* • **Alert**—Generates an alert for each application traffic flow. The alert is saved in the threat log. • **Drop**—Drops the application traffic. • **Reset Client**—For TCP, resets the client-side connection. For UDP, drops the connection. • **Reset Server**—For TCP, resets the server-side connection. For UDP, drops the connection. • **Reset Both**—For TCP, resets the connection on both client and server ends. For UDP, drops the connection. *In some cases, when the profile action is set to **reset-both**, the associated threat log might display the action as **reset-server**. This occurs when the firewall detects a threat at the beginning of a session and presents the client with a 503 block page. Because the block page disallows the connection, the client-side does not need to be reset and only the server-side connection is reset.* • **Block IP**— This action blocks traffic from either a source or a source-destination pair. It is configurable for a specified period of time. In addition, you can enable the DNS Sinkholing action in Anti-Spyware profiles to enable the firewall to forge a response to a DNS query for a known malicious domain, causing the malicious domain name to resolve to an IP address that you define. This feature helps to identify infected hosts on the protected network using DNS traffic. Infected hosts can then be easily identified in the traffic and threat logs because any host that attempts to connect to the sinkhole IP address are most likely infected with malware. Anti-Spyware and Vulnerability Protection profiles are configured similarly.

https://t.me/learningnets

Profile Type	Description
Vulnerability Protection Profiles	Vulnerability Protection profiles stop attempts to exploit system flaws or gain unauthorized access to systems. While Anti-Spyware profiles help identify infected hosts as traffic leaves the network, Vulnerability Protection profiles protect against threats entering the network. For example, Vulnerability Protection profiles help protect against buffer overflows, illegal code execution, and other attempts to exploit system vulnerabilities. The default Vulnerability Protection profile protects clients and servers from all known critical, high, and medium-severity threats. You can also create exceptions, which allow you to change the response to a specific signature. When the firewall is managed by a Panorama management server, the ThreatID is mapped to the corresponding custom threat on the firewall to enable the firewall to generate a threat log populated with the configured custom ThreatID. When the firewall detects a threat event, you can configure the following actions in an Anti-Spyware profile: • **Default**—For each threat signature and Anti-Spyware signature that is defined by Palo Alto Networks, a default action is specified internally. Typically the default action is an alert or a reset-both. The default action is displayed in parenthesis, for example default (alert) in the threat or Antivirus signature. • **Allow**—Permits the application traffic *The **Allow** action does not generate logs related to the signatures or profiles.* • **Alert**—Generates an alert for each application traffic flow. The alert is saved in the threat log. • **Drop**—Drops the application traffic. • **Reset Client**—For TCP, resets the client-side connection. For UDP, drops the connection. • **Reset Server**—For TCP, resets the server-side connection. For UDP, drops the connection. • **Reset Both**—For TCP, resets the connection on both client and server ends. For UDP, drops the connection. *In some cases, when the profile action is set to **reset-both**, the associated threat log might display the action as **reset-server**. This occurs when the firewall detects a threat at the beginning of a session and presents the client with a 503 block page. Because the block page disallows the connection, the client-side does not need to be reset and only the server-side connection is reset.*

https://t.me/learningnets

Profile Type	Description
	• **Block IP**— This action blocks traffic from either a source or a source-destination pair. It is configurable for a specified period of time.
URL Filtering Profiles	URL Filtering profiles enable you to monitor and control how users access the web over HTTP and HTTPS. The firewall comes with a default profile that is configured to block websites such as known malware sites, phishing sites, and adult content sites. You can use the default profile in a security policy, clone it to be used as a starting point for new URL filtering profiles, or add a new URL profile that will have all categories set to allow for visibility into the traffic on your network. You can then customize the newly added URL profiles and add lists of specific websites that should always be blocked or allowed, which provides more granular control over URL categories.
Data Filtering Profiles	Data filtering profiles prevent sensitive information such as credit card or social security numbers from leaving a protected network. The data filtering profile also allows you to filter on key words, such as a sensitive project name or the word confidential. It is important to focus your profile on the desired file types to reduce false positives. For example, you may only want to search Word documents or Excel spreadsheets. You may also only want to scan web-browsing traffic, or FTP. You can create custom data pattern objects and attach them to a Data Filtering profile to define the type of information on which you want to filter. Create data pattern objects based on: • **Predefined Patterns**—Filter for credit card and social security numbers (with or without dashes) using predefined patterns. • **Regular Expressions**—Filter for a string of characters. • **File Properties**—Filter for file properties and values based on file type. *If you're using a third-party, endpoint data loss prevention (DLP) solutions to populate file properties to indicate sensitive content, this option enables the firewall to enforce your DLP policy.* To get started, Set Up Data Filtering.
File Blocking Profiles	The firewall uses file blocking profiles to block specified file types over specified applications and in the specified session flow direction (inbound/outbound/both). You can set the profile to alert or block on upload and/or download and you can specify which applications will be subject to the file blocking profile. You can also configure custom block pages that will appear when a user attempts to download the

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Profile Type	Description
	specified file type. This allows the user to take a moment to consider whether or not they want to download a file. You can define your own custom File Blocking profiles, or choose one of the following predefined profiles when applying file blocking to a Security policy rule. The predefined profiles, which are available with content release version 653 and later, allow you to quickly enable best practice file blocking settings: • **basic file blocking**—Attach this profile to the Security policy rules that allow traffic to and from less sensitive applications to block files that are commonly included in malware attack campaigns or that have no real use case for upload/download. This profile blocks upload and download of PE files (.scr, .cpl, .dll, .ocx, .pif, .exe) , Java files (.class, .jar), Help files (.chm, .hlp) and other potentially malicious file types, including .vbe, .hta, .wsf, .torrent, .7z, .rar, .bat. Additionally, it prompts users to acknowledge when they attempt to download encrypted-rar or encrypted-zip files. This rule alerts on all other file types to give you complete visibility into all file types coming in and out of your network. • **strict file blocking**—Use this stricter profile on the Security policy rules that allow access to your most sensitive applications. This profile blocks the same file types as the other profile, and additionally blocks flash, .tar, multi-level encoding, .cab, .msi, encrypted-rar, and encrypted-zip files. Configure a file blocking profile with the following actions: • **Alert**—When the specified file type is detected, a log is generated in the data filtering log. • **Block**—When the specified file type is detected, the file is blocked and a customizable block page is presented to the user. A log is also generated in the data filtering log. • **Continue**—When the specified file type is detected, a customizable response page is presented to the user. The user can click through the page to download the file. A log is also generated in the data filtering log. Because this type of forwarding action requires user interaction, it is only applicable for web traffic. To get started, Set Up File Blocking.
WildFire Analysis Profiles	Use a WildFire analysis profile to enable the firewall to forward unknown files or email links for WildFire analysis. Specify files to be forwarded for analysis based on application, file type, and transmission direction (upload or download). Files or email links matched to the profile rule are forwarded either the WildFire public cloud or the WildFire private cloud (hosted with a WF-500 appliance), depending on the analysis location defined for the rule. If a profile rule is set to forward files to the WildFire public cloud, the firewall also

https://t.me/learningnets

Profile Type	Description
	forwards files that match existing antivirus signatures, in addition to unknown files. You can also use the WildFire analysis profiles to set up a WildFire hybrid cloud deployment. If you are using a WildFire appliance to analyze sensitive files locally (such as PDFs), you can specify for less sensitive files types (such as PE files) or file types that are not supported for WildFire appliance analysis (such as APKs) to be analyzed by the WildFire public cloud. Using both the WildFire appliance and the WildFire cloud for analysis allows you to benefit from a prompt verdict for files that have already been processed by the cloud, and for files that are not supported for appliance analysis, and frees up the appliance capacity to process sensitive content.
DoS Protection Profiles	DoS protection profiles provide detailed control for Denial of Service (DoS) protection policies. DoS policies allow you to control the number of sessions between interfaces, zones, addresses, and countries based on aggregate sessions or source and/or destination IP addresses. There are two DoS protection mechanisms that the Palo Alto Networks firewalls support. • **Flood Protection**—Detects and prevents attacks where the network is flooded with packets resulting in too many half-open sessions and/or services being unable to respond to each request. In this case the source address of the attack is usually spoofed. See DoS Protection Against Flooding of New Sessions. • **Resource Protection**— Detects and prevent session exhaustion attacks. In this type of attack, a large number of hosts (bots) are used to establish as many fully established sessions as possible to consume all of a system's resources. You can enable both types of protection mechanisms in a single DoS protection profile. The DoS profile is used to specify the type of action to take and details on matching criteria for the DoS policy. The DoS profile defines settings for SYN, UDP, and ICMP floods, can enable resource protect and defines the maximum number of concurrent connections. After you configure the DoS protection profile, you then attach it to a DoS policy. When configuring DoS protection, it is important to analyze your environment in order to set the correct thresholds and due to some of the complexities of defining DoS protection policies, this guide will not go into detailed examples.
Zone Protection Profiles	Zone Protection Profiles provide additional protection between specific network zones in order to protect the zones against attack. The profile must be applied to the entire zone, so it is important to carefully test the profiles in order to prevent issues that may arise

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Profile Type	Description
	with the normal traffic traversing the zones. When defining packets per second (pps) thresholds limits for zone protection profiles, the threshold is based on the packets per second that do not match a previously established session.
Security Profile Group	A security profile group is a set of security profiles that can be treated as a unit and then easily added to security policies. Profiles that are often assigned together can be added to profile groups to simplify the creation of security policies. You can also setup a default security profile group—new security policies will use the settings defined in the default profile group to check and control traffic that matches the security policy. Name a security profile group default to allow the profiles in that group to be added to new security policies by default. This allows you to consistently include your organization's preferred profile settings in new policies automatically, without having to manually add security profiles each time you create new rules. See Create a Security Profile Group and Set Up or Override a Default Security Profile Group.  *For recommendations on the best-practice settings for security profiles, see* Create Best Practice Security Profiles for the Internet Gateway.

Create a Security Profile Group

Use the following steps to create a security profile group and add it to a security policy.

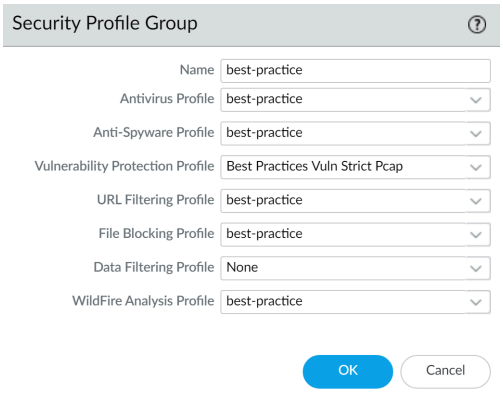
 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 1 | Create a security profile group.

If you name the group `default`, the firewall will automatically attach it to any new rules you create. This is a time saver if you have a preferred set of security profiles that you want to make sure get attached to every new rule.

1. Select **Objects** > **Security Profile Groups** and **Add** a new security profile group.
2. Give the profile group a descriptive **Name**, for example, Threats.
3. If the firewall is in Multiple Virtual System Mode, enable the profile to be **Shared** by all virtual systems.
4. Add existing profiles to the group.



5. Click **OK** to save the profile group.

STEP 2 | Add a security profile group to a security policy.

1. Select **Policies** > **Security** and **Add** or modify a security policy rule.
2. Select the **Actions** tab.
3. In the Profile Setting section, select **Group** for the **Profile Type**.
4. In the **Group Profile** drop-down, select the group you created (for example, select the best-practice group):



5. Click **OK** to save the policy and **Commit** your changes.

STEP 3 | Save your changes.

Click **Commit**.

Set Up or Override a Default Security Profile Group

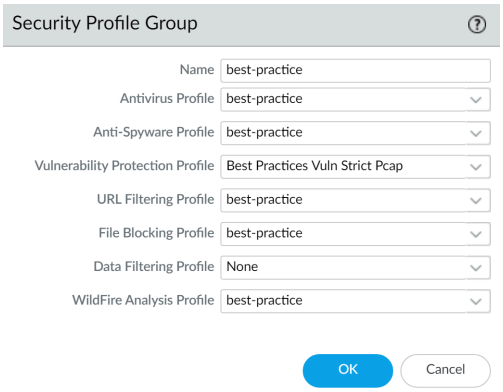
Use the following options to set up a default security profile group to be used in new security policies, or to override an existing default group. When an administrator creates a new security policy, the default profile group will be automatically selected as the policy's profile settings, and traffic matching the policy will be checked according to the settings defined in the profile group

https://t.me/learningnets

(the administrator can choose to manually select different profile settings if desired). Use the following options to set up a default security profile group or to override your default settings.

If no default security profile exists, the profile settings for a new security policy are set to ***None*** *by default.*

- Create a security profile group.
  1. Select **Objects** > **Security Profile Groups** and Add a new security profile group.
  2. Give the profile group a descriptive **Name**, for example, Threats.
  3. If the firewall is in Multiple Virtual System Mode, enable the profile to be **Shared** by all virtual systems.
  4. Add existing profiles to the group. For details on creating profiles, see Security Profiles.



  5. Click **OK** to save the profile group.
  6. Add the security profile group to a security policy.
  7. **Add** or modify a security policy rule and select the **Actions** tab.
  8. Select **Group** for the **Profile Type**.
  9. In the **Group Profile** drop-down, select the group you created (for example, select the Threats group):

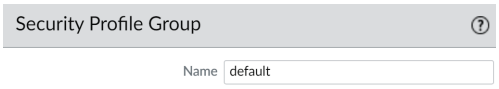


  10. Click **OK** to save the policy and **Commit** your changes.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

- Set up a default security profile group.
  1. Select **Objects** > **Security Profile Groups** and add a new security profile group or modify an existing security profile group.
  2. **Name** the security profile group **`default`**:

  3. Click **OK** and **Commit**.
  4. Confirm that the default security profile group is included in new security policies by default:
     1. Select **Policies** > **Security** and **Add** a new security policy.
     2. Select the **Actions** tab and view the **Profile Setting** fields:

        By default, the new security policy correctly shows the **Profile Type** set to Group and the default **Group Profile** is selected.

- Override a default security profile group.

  If you have an existing default security profile group, and you do not want that set of profiles to be attached to a new security policy, you can continue to modify the Profile Setting fields according to your preference. Begin by selecting a different Profile Type for your policy (**Policies** > **Security** > **Security Policy Rule** > **Actions**).

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Track Rules Within a Rulebase

To keep track of rules within a rulebase, you can refer to the *rule number*, which changes depending on the order of a rule in the rulebase. The rule number determines the order in which the firewall applies the rule.

The *universally unique identifier (UUID)* for a rule never changes even if you modify the rule, such as when you change the rule name. The UUID allows you to track the rule across rulebases even after you deleted the rule.

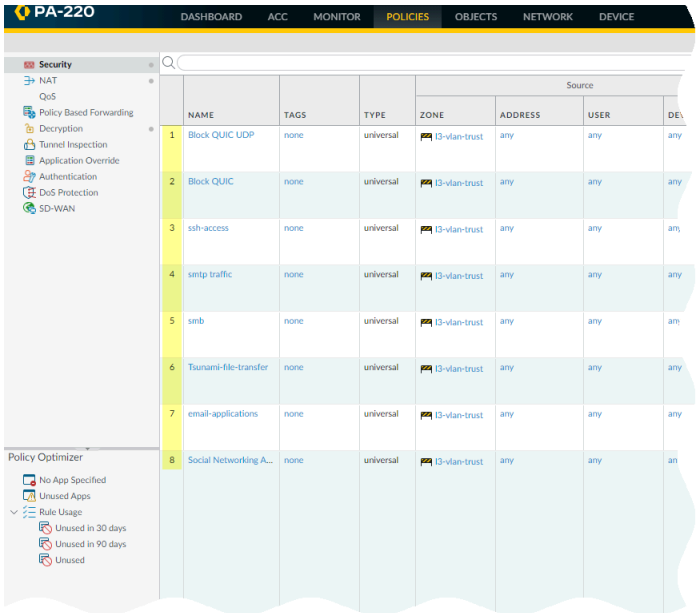
Rule Numbers

The firewall automatically numbers each rule within a rulebase; when you move or reorder rules, the numbers change based on the new order. When you filter the list of rules to find rules that match specific criteria, the firewall display each rule with its number in the context of the complete set of rules in the rulebase and its place in the evaluation order.

Panorama independently numbers pre-rules, post-rules, and default rules. When Panorama pushes rules to a firewall, the rule numbering reflects the hierarchy and evaluation order of shared rules, device group pre-rules, firewall rules, device group post-rules, and default rules. You can **Preview Rules** in Panorama to display an ordered list of the total number of rules on a firewall.

- View the numbered list of rules on the firewall.

  Select **Policies** and any rulebase under it. For example, **Policies** > **Security**. The left-most column in the table displays the rule number.

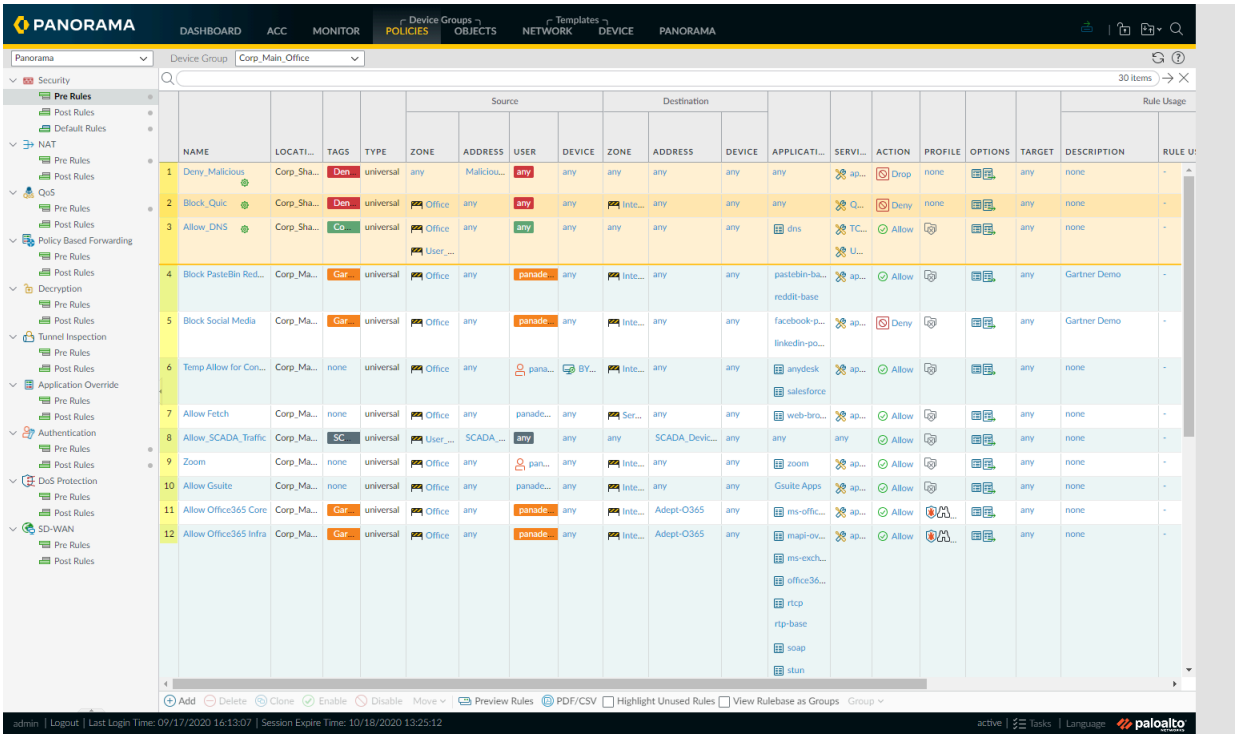


 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

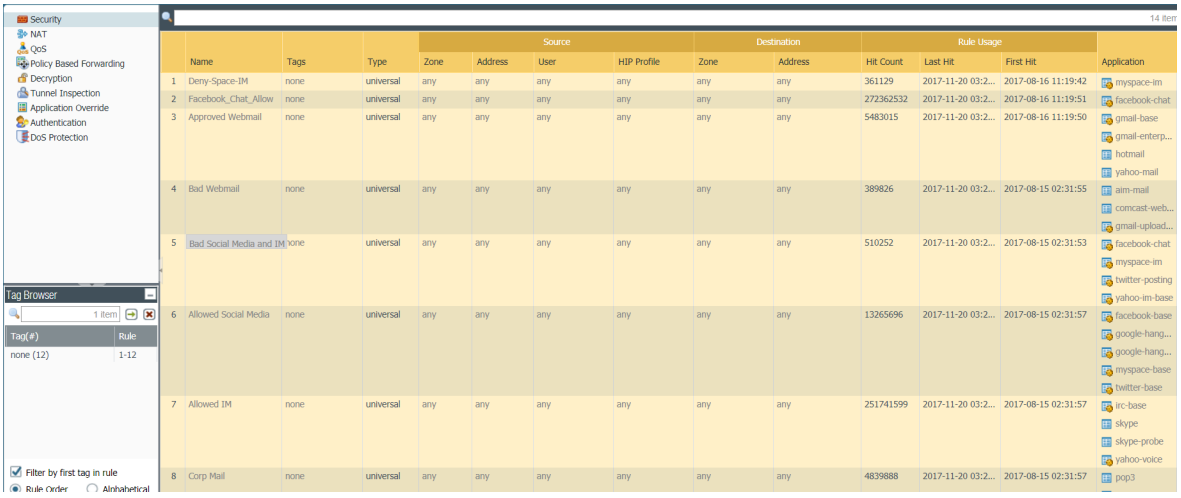
- View the numbered list of rules on Panorama.

  Select **Policies** and any rulebase under it. For example, **Policies** > **Security** > **Pre-rules**.



- After you push the rules from Panorama, view the complete list of rules with numbers on the firewall.

  From the web interface on the firewall, select **Policies** and pick any rulebase under it. For example, select **Policies** > **Security** and view the complete set of numbered rules that the firewall will evaluate.



https://t.me/learningnets

Rule UUIDs

The universally unique identifier (UUID) for a rule is a 32-character string (based on data such as the network address and the timestamp of creation) that the firewall or Panorama assigns to the rule. The UUID uses the format 8-4-4-4-12 (where 8, 4, and 12 represent the number of unique characters separated by hyphens). UUIDs identify rules for all policy rulebases. You can also use UUIDs to identify applicable rules in the following log types: Traffic, Threat, URL Filtering, WildFire Submission, Data Filtering, GTP, SCTP, Tunnel Inspection, Configuration, and Unified.

Using the UUID to search for a rule enables you to locate a specific rule you want to find among thousands of rules that may have similar or identical names. UUIDs also simplify automation and integration for rules in third-party systems (such as ticketing or orchestration) that do not support names.

In some cases, you may need to generate new UUIDs for existing rulebases. For example, if you want to export a configuration to another firewall, you need to *regenerate the UUIDs* for the rules as you import the configuration to ensure there are no duplicate UUIDs. If you regenerate UUIDs, you are no longer able to track those rules using their previous UUIDs and the hit data and app usage data for those rules are reset.

The firewall or Panorama assigns UUIDs when you:

- Create new rules
- Clone existing rules
- Override the default security rules
- Load a named configuration and regenerate UUIDs
- Load a named configuration containing new rules that are not in the running configuration
- Upgrade the firewall or Panorama to a PAN-OS 9.0 release

When you load a configuration that contains rules with UUIDs, the firewall considers rules to be the same if the rule name, rulebase, and virtual system all match. Panorama considers rules to be the same if the rule name, rulebase, and the device group all match.

Keep in mind the following important points for UUIDs:

- If you manage firewall policy from Panorama, UUIDs are generated on Panorama and therefore must be pushed from Panorama. If you do not push the configuration from Panorama prior to upgrading the firewalls to PAN-OS 9.0, the firewall upgrade will not succeed because it will not have the UUIDs.
- In addition, if you are upgrading an HA pair, upon upgrade to PAN-OS 9.0, each peer independently assigns UUIDs for each policy rule. Because of this, the peers will show as out of sync until you sync the configuration (**Dashboard** > **Widgets** > **System** > **High Availability** > **Sync to peer**).
- If you remove an existing high availability (HA) configuration after upgrading to PAN-OS 9.0, you must regenerate the UUIDs on one of the peers (**Device** > **Setup** > **Operations** > **Load named configuration snapshot** > **Regenerate UUIDs for the selected named configuration**) and commit the changes to prevent UUID duplication.
- All rules pushed from Panorama will share the same UUID; all rules local to a firewall will have different UUIDs. If you create a rule locally on the firewall after you push the rules from Panorama to the firewalls, the rule you created locally has its own UUID.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

- To replace an RMA Panorama, make sure you **Retain Rule UUIDs** when you load the named Panorama configuration snapshot. If you do not select this option, Panorama removes all previous rule UUIDs from the configuration snapshot and assigns new UUIDs to the rules on Panorama, which means it does not retain information associated with the previous UUIDs, such as the policy rule hit count.

 ©2022 Palo Alto Networks, Inc.

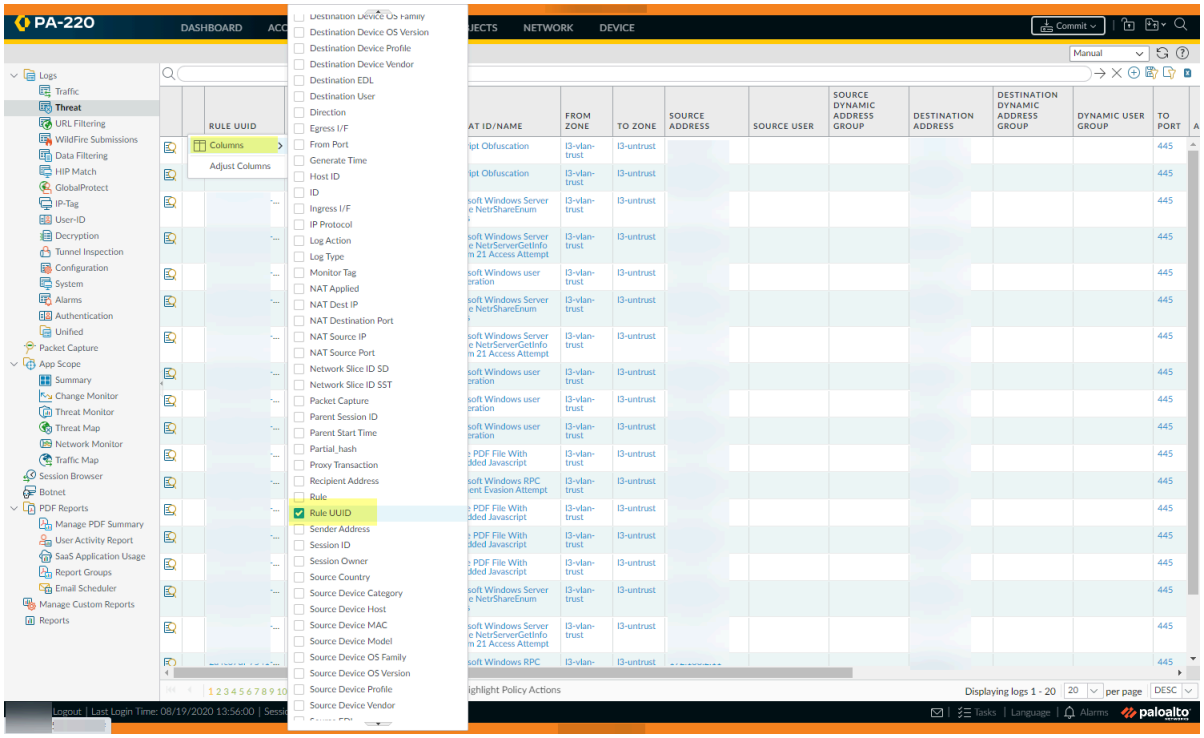
https://t.me/learningnets

- Display the Rule UUID column for logs and the UUID column for policy rules.

  To view the UUIDs, you must display the column, which does not display by default.

  - To display the UUID in logs:

    1. Select **Monitor** and then expand the column header ( ).
    2. Select **Columns**.
    3. Enable **Rule UUID**.

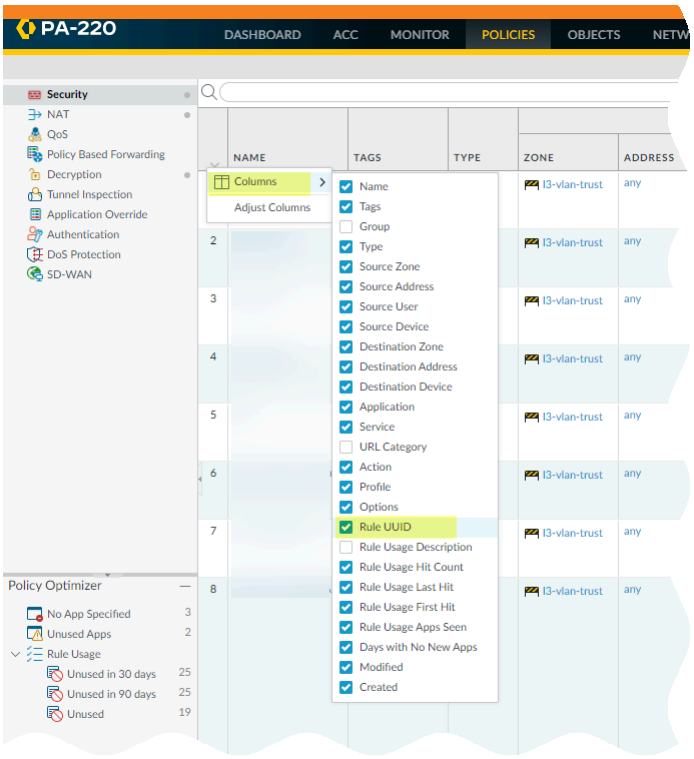


  - To display UUIDs on the policy rulebase:

    1. Select **Policies** and then expand the column header ( ).
    2. Select **Columns**.
    3. Enable **Rule UUID**.

       UUIDs are available for all policy rulebases.

https://t.me/learningnets

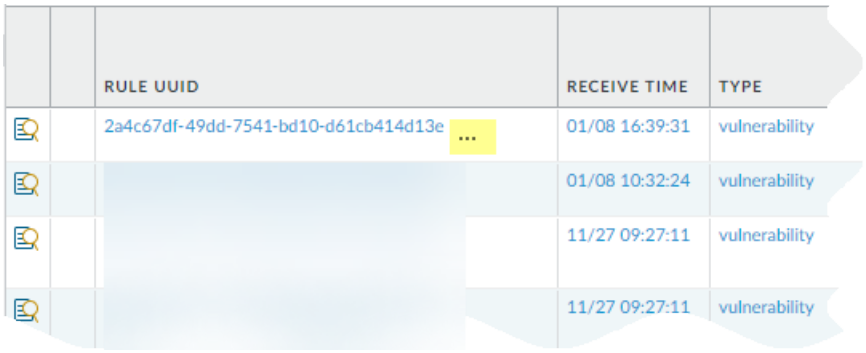


https://t.me/learningnets

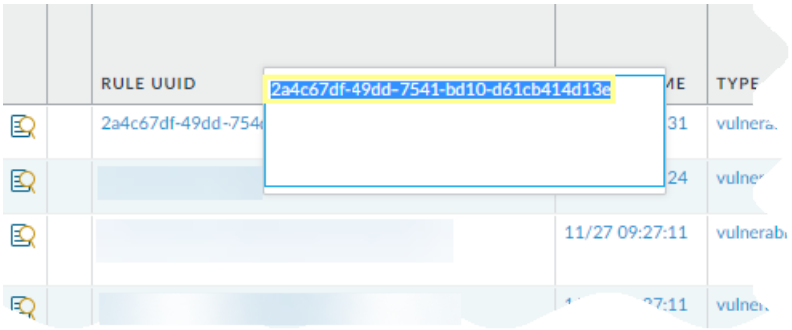
- Copy the UUID for a log or policy rule.

  Copying the UUID allows you to paste the UUID in to searches, the ACC, custom reports, filters, and anywhere else you want to locate a rule identified by that UUID.

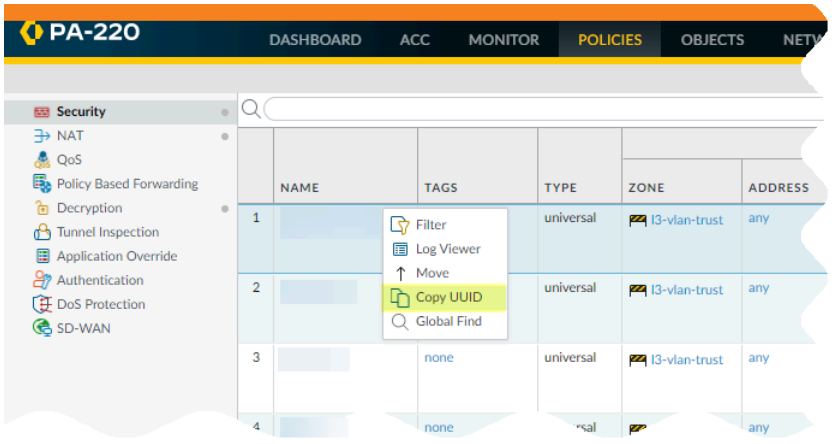
  1. Select the ellipses that display when you move your cursor over the entry in the Rule UUID column.



  2. Copy the UUID from the pop-up.



  You can also go to the **Policies** tab, click the arrow to the right of the rule name, and **Copy UUID**.



- Check the Configuration Logs to view UUIDs for deleted rules.

  To view the UUID for a deleted rule, select **Monitor** > **Logs** > **Configuration**.

©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Enforce Policy Rule Description, Tag, and Audit Comment

When creating or modifying rules, you can require a rule description, tag, and audit comment to ensure your policy rulebase is correctly organized and grouped, and to preserve important rule history for auditing purposes. By requiring a rule description, tag, and audit comment, you can simplify your policy rulebase review by ensuring that rules are appropriately grouped, and that the rule change history is tracked when creating or modifying a rule. For uniformity, you can set specific requirements for what the audit comment can include.

By default, enforcement of a description, tag, and audit comment is not enabled. You can specify whether a description, tag, audit comment, or any combination of these three is required to successfully add or modify a rule. The audit comment archive allows you to view the audit comments entered for a selected rule, review the configuration log history, and compare rule configuration versions.

STEP 1 | Launch the Web Interface.

STEP 2 | Select **Device** > **Setup** > **Management** and edit the Policy Rulebase Settings.

STEP 3 | Configure the settings you want to enforce. In this example, tags and audit comments are required for all policies.



Enforce audit comments for policy rules to capture the reason an administrator creates or modifies a rule. Requiring audit comments on policy rules helps maintain an accurate rule history for auditing purposes.

STEP 4 | Configure the Audit Comment Regular Expression to specify the audit comment format.

When administrators create or modify a rule, you can require they enter a comment those audit comments adhere to a specific format that fits your business and auditing needs by specifying letter and number expressions. For example, you can use this setting to specify regular expressions that match your ticketing number formats:

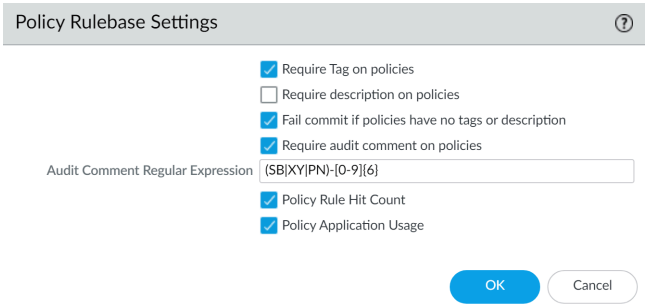
- **[0-9]{<Number of digits>}**—Requires the audit comment to contain a minimum number of digits that range from 0 to 9. For example, **[0-9]{6}** requires a minimum of six digit in a numerical expression with numbers 0 to 9.
- <**Letter Expression**>—Requires the audit comment to contain a letter expression. For example, **Reason for Change-** requires that the administrator begin the audit comment with this letter expression.
- <**Letter Expression**>-**[0-9]{<Number of digits>}**—Requires the audit comment to contain a predetermined character followed by a minimum number of digits that range from 0 to 9. For example, **SB-[0-9]{6}** requires the audit comment format to begin with **SB-**, followed by a minimum six digits in a numerical expression with values from 0 to 9. For example, **SB-012345**.
- **(**<**Letter Expression**>**)|(**<**Letter Expression**>**)|(**<**Letter Expression**>**)|-[0-9]{<Number of digits>}**—Requires the audit comment to contain a prefix using any one of the predetermined letter expressions with a minimum number of digits that range from 0 to 9. For example, **(SB|XY|PN)-[0-9]{6}** requires the

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

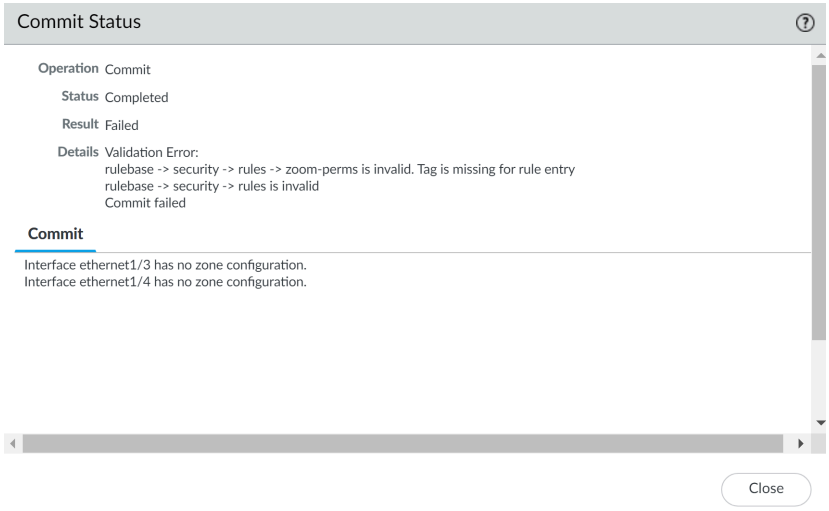
audit comment format to begin with **SB-**, **XY-**, or **PN-** followed by a minimum of six digits in a numerical expression with values from 0 to 9. For example, **SB-012345**, **XY-654321**, or **PN-012543**.

STEP 5 | Click **OK** to apply the new policy rulebase settings.



STEP 6 | **Commit** the changes.

After you commit the policy rulebase settings changes, modify the existing policy rule based on the rulebase settings you decided to enforce.

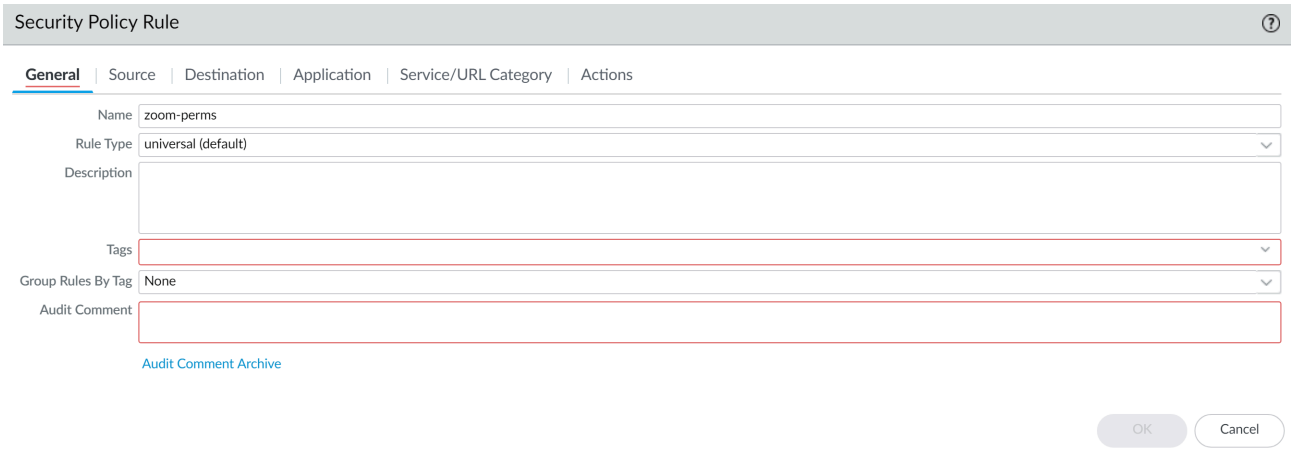


 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 7 | Verify that the firewall is enforcing the new policy rulebase settings.

1. Select **Policies** and **Add** a new rule.
2. Confirm that you must add a tag and enter an audit comment click **OK**.



 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Move or Clone a Policy Rule or Object to a Different Virtual System

On a firewall that has more than one virtual system (vsys), you can move or clone policy rules and objects to a different vsys or to the Shared location. Moving and cloning save you the effort of deleting, recreating, or renaming rules and objects. If the policy rule or object that you will move or clone from a vsys has references to objects in that vsys, move or clone the referenced objects also. If the references are to shared objects, you do not have to include those when moving or cloning. You can Use Global Find to Search the Firewall or Panorama Management Server for references.

When cloning multiple policy rules, the order by which you select the rules will determine the order they are copied to the device group. For example, if you have rules 1-4 and your selection order is 2-1-4-3, the device group where these rules will be cloned will display the rules in the same order you selected. However, you can reorganize the rules as you see fit once they have been successfully copied.

STEP 1 | Select the policy type (for example, **Policy** > **Security**) or object type (for example, **Objects** > **Addresses**).

STEP 2 | Select the **Virtual System** and select one or more policy rules or objects.

STEP 3 | Perform one of the following steps:

- Select **Move** > **Move to other vsys** (for policy rules).
- Click **Move** (for objects).
- Click **Clone** (for policy rules or objects).

STEP 4 | In the **Destination** drop-down, select the new virtual system or **Shared**.

STEP 5 | (Policy rules only) Select the **Rule order**:

- **Move top** (default)—The rule will come before all other rules.
- **Move bottom**—The rule will come after all other rules.
- **Before rule**—In the adjacent drop-down, select the rule that comes after the Selected Rules.
- **After rule**—In the adjacent drop-down, select the rule that comes before the Selected Rules.

STEP 6 | The **Error out on first detected error in validation** check box is selected by default. The firewall stops performing the checks for the move or clone action when it finds the first error, and displays just this error. For example, if an error occurs when the **Destination** vsys doesn't have an object that the policy rule you are moving references, the firewall will display the error and stop any further validation. When you move or clone multiple items at once, selecting this check box will allow you to find one error at a time and troubleshoot it. If you clear the check box, the firewall collects and displays a list of errors. If there are any errors in validation, the object is not moved or cloned until you fix all the errors.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 7 | Click **OK** to start the error validation. If the firewall displays errors, fix them and retry the move or clone operation. If the firewall doesn't find errors, the object is moved or cloned successfully. After the operation finishes, click **Commit**.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Use an Address Object to Represent IP Addresses

Create an address object on the firewall to group IP addresses or to specify an FQDN, and then reference the address object in a firewall policy rule, filter, or other function to avoid having to individually specify multiple IP addresses in the rule, filter, or other function.

Furthermore, you can reference the same address object in multiple policy rules, filters, or other functions without needing to specify the same individual addresses in each use. For example, you can create an address object that specifies an IPv4 address range and then reference the address object in a Security policy rule, a NAT policy rule, and a custom report log filter.

- Address Objects
- Create an Address Object

Address Objects

An address object is a set of IP addresses that you can manage in one place and then use in multiple firewall policy rules, filters, and other functions. There are four types of address objects: **IP Netmask**, **IP Range**, **IP Wildcard Mask**, and **FQDN**.

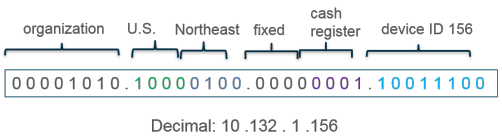
An address object of type **IP Netmask**, **IP Range**, or **FQDN** can specify IPv4 or IPv6 addresses. An address object of type **IP Wildcard Mask** can specify only IPv4 addresses.

An address object of type **IP Netmask** requires you to enter the IP address or network using slash notation to indicate the IPv4 network or the IPv6 prefix length. For example, 192.168.18.0/24 or 2001:db8:123:1::/64.

An address object of type **IP Range** requires you to enter the IPv4 or IPv6 range of addresses separated by a hyphen.

An address object of type **FQDN** (for example, paloaltonetworks.com) provides further ease of use because DNS provides the FQDN resolution to the IP addresses instead of you needing to know the IP addresses and manually updating them every time the FQDN resolves to new IP addresses.

An address object of type **IP Wildcard Mask** is useful if you define private IPv4 addresses to internal devices and your addressing structure assigns meaning to certain bits in the address. For example, the IP address of cash register 156 in the northeastern U.S. could be 10.132.1.156 based on these bit assignments:



An address object of type **IP Wildcard Mask** specifies which source or destination addresses are subject to a Security policy rule. For example, 10.132.1.1/0.0.2.255. A zero (0) bit in the mask indicates that the bit being compared must match the bit in the IP address that is covered by the zero. A one (1) bit in the mask (a wildcard bit) indicates that the bit being compared need not match the bit in the IP address. The following snippets of an IP address and wildcard mask illustrate how they yield four matches:

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets



After you Create an Address Object:

- You can reference an address object of type **IP Netmask**, **IP Range**, or **FQDN** in a policy rule for Security, Authentication, NAT, NAT64, Decryption, DoS Protection, Policy-Based Forwarding (PBF), QoS, Application Override, or Tunnel Inspection; or in a NAT address pool, VPN tunnel, path monitoring, External Dynamic List, Reconnaissance Protection, ACC global filter, log filter, or custom report log filter.
- You can reference an address object of type **IP Wildcard Mask** only in a Security policy rule.

Create an Address Object

Create Address Objects to represent one or more IP addresses and then reference the address objects in one or more policy rules, filters, or other firewall functions. If you want to change the set of addresses, you change an address object once rather than change multiple policy rules or filters, which reduces your operational overhead.

STEP 1 | Create an address object.

1. Select **Objects** > **Addresses** and **Add** an address object by **Name**. The name is case-sensitive, must be unique, and can be up to 63 characters (letters, numbers, spaces, hyphens, and underscores).
2. Select the **Type** of address object:
   - **IP Netmask**—Specify a single IPv4 or IPv6 address, an IPv4 network with slash notation, or an IPv6 address and prefix. For example, 192.168.80.0/24 or 2001:db8:123:1::/64. Optionally, click **Resolve** to see the associated FQDN (based on the DNS configuration of the firewall or Panorama). To change the address object type from **IP Netmask** to **FQDN**, select the FQDN and click **Use this FQDN**. The **Type** changes to **FQDN** and the FQDN you select appears in the text field.
   - **IP Range**—Specify a range of IPv4 addresses or IPv6 addresses separated by a hyphen. For example, 192.168.40.1-192.168.40.255 or 2001:db8:123:1::1-2001:db8:123:1::22.
   - **IP Wildcard Mask**—Specify an IP wildcard address (IPv4 address followed by a slash and a mask, which must begin with a 0). For example, 10.5.1.1/0.127.248.2. A zero (0) in the mask indicates the bit being compared must match the bit in the IP address that is covered by the zero. A one (1) in the mask (wildcard bit) indicates the bit being compared need not match the bit in the IP address covered by the one.
   - **FQDN**—Specify the domain name. The FQDN initially resolves at commit time. The firewall subsequently refreshes the FQDN based on the time-to-live (TTL) of the FQDN in DNS, as long as the TTL is greater than or equal to the **Minimum FQDN Refresh Time** you configure (or the default of 30 seconds). The FQDN is resolved by the system DNS server or a DNS proxy object, if a proxy is configured. Click **Resolve** to see the associated IP address (based on the DNS configuration of the

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

firewall or Panorama). To change the address object type from FQDN to IP Netmask, select an IP Netmask and click **Use this address**. The **Type** changes to **IP Netmask** and the IP address you select appears in the text field.

3. (Optional) Enter one or more Use Tags to Group and Visually Distinguish Objects to apply to the address object.
4. Click **OK**.

STEP 2 | **Commit** your changes.

STEP 3 | View logs filtered by address object, address group, or wildcard address.

1. For example, select **Monitor** > **Logs** > **Traffic** to view traffic logs.
2. Select ⊕ to add a log filter.
3. Select the **Address** attribute, the **in** Operator, and enter the name of the address object for which you want to view logs. Alternatively, enter an address group name or a wildcard address, such as 10.155.3.4/0.0.240.255.
4. Click **Apply**.

STEP 4 | View a custom report based on an address object.

1. Select **Monitor** > **Manage Custom Reports** and select a report that uses a Database such as Traffic Log.
2. Select **Filter Builder**.
3. Select an Attribute such as **Address**, **Destination Address** or **Source Address**, select an Operator, and enter the name of the address object for which you want to view the report.

STEP 5 | Use a filter in the ACC to view network activity based on a source IP address or destination IP address that uses an address object.

1. Select **ACC** > **Network Activity**.
2. View the **Source IP Activity—For Global Filters**, click ⊕ to add a filter and select one of the following: **Address** or **Source** > **Source Address** or **Destination** > **Destination Address** and select an address object.
3. View the **Destination IP Activity—For Global Filters**, click the ⊕ to add a filter and select one of the following: **Address** or **Source** > **Source Address** or **Destination** > **Destination Address** and select an address object.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Use Tags to Group and Visually Distinguish Objects

You can tag objects to group related items and add color to the tag in order to visually distinguish them for easy scanning. You can create tags for the following objects: address objects, address groups, user groups, zones, service groups, and policy rules.

The firewall and Panorama support both static tags and dynamic tags. Dynamic tags are registered from a variety of sources and are not displayed with the static tags because dynamic tags are not part of the configuration on the firewall or Panorama. See Register IP Addresses and Tags Dynamically for information on registering tags dynamically. The tags discussed in this section are statically added and are part of the configuration.

You can apply one or more tags to objects and to policy rules, up to a maximum of 64 tags per object. Panorama supports a maximum of 10,000 tags, which you can apportion across Panorama (shared and device groups) and the managed firewalls (including firewalls with multiple virtual systems).

- Create and Apply Tags
- Modify Tags
- View Rules by Tag Group

Create and Apply Tags

Use tags to identify the purpose of a rule or configuration object and to help you better organize your rulebase. To ensure that policy rules are properly tagged, see how to Enforce Policy Rule Description, Tag, and Audit Comment. Additionally, you can View Rules by Tag Group by first creating and then setting the tag as the Group tag.

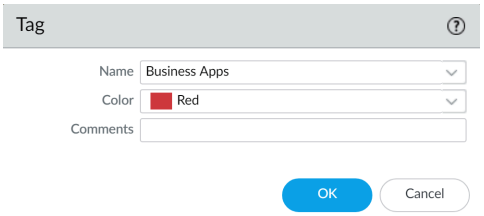
 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 1 | Create tags.

To tag a zone, you must create a tag with the same name as the zone. When the zone is attached in policy rules, the tag color automatically displays as the background color against the zone name.

1. Select **Objects** > **Tags**.
2. On Panorama or a multiple virtual system firewall, select the **Device Group** or the **Virtual System** to make the tag available.
3. **Add** a tag and enter a **Name** to identify the tag or select a zone **Name** to create a tag for a zone. The maximum length is 127 characters.
4. (Optional) Select **Shared** to create the object in a shared location for access as a shared object in Panorama or for use across all virtual systems in a multiple virtual system firewall.
5. (Optional) Assign a **Color** from the 17 predefined colors. By default, **Color** is **None**.



6. Click **OK** and **Commit** to save your changes.

STEP 2 | Apply tags to policy.

1. Select **Policies** and any rulebase under it.
2. **Add** a policy rule and use the tagged objects you created in Step 1.
3. Verify that the tags are in use.

				Source				Destination	
	NAME	TAGS	TYPE	ZONE	ADDRESS	USER	DEVICE	ZONE	ADDRESS
1	General Business Apps	Business Apps	universal	any	any	known-user	any	any	any

STEP 3 | Apply tags to an address object, address group, service, or service group.

1. Create the object.

   For example, to create a service group, select **Objects** > **Service Groups** > **Add**.
2. Select a tag (**Tags**) or enter a name in the field to create a new tag.

   To edit a tag or add color to the tag, see Modify Tags.

Modify Tags

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

- Select **Objects** > **Tags** to perform any of the following operations with tags:
  - Click the **Name** to edit the properties of a tag.
  - Select a tag in the table and **Delete** the tag from the firewall.
  - **Clone** a tag to duplicate it with the same properties. A numerical suffix is added to the tag name (for example, FTP-1).

  For details on creating tags, see Create and Apply Tags. For information on working with tags, see View Rules by Tag Group.

View Rules by Tag Group

View your policy rulebase as tag groups to visually group rules based on the tagging structure you created. In this view, you can perform operational procedures such as adding, deleting, and moving the rules in the selected tag group more easily. Viewing the rulebase as tag groups maintains the rule evaluation order and a single tag may appear multiple times throughout the rulebase to visually preserve the rule hierarchy.

You must create the tag before you can assign it as a group tag on a rule. Policy rules that are already tagged on upgrade to PAN-OS 9.0 have the first tag automatically assigned as the Group tag. Before you upgrade to PAN-OS 9.0, review the tagged rules in your rulebase to ensure rules are correctly grouped. You must manually edit each tag rule and configure the correct Group tag if your rules are grouped incorrectly after you upgrade to PAN-OS 9.0.

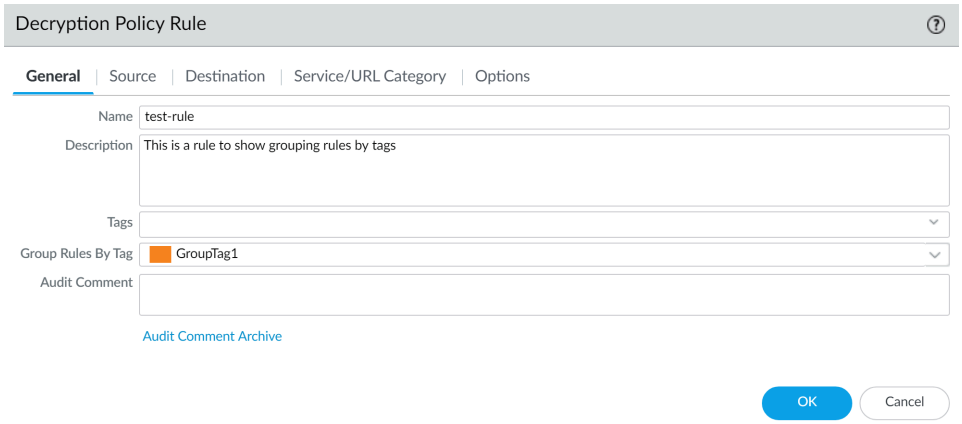


STEP 1 | Launch the Web Interface.

STEP 2 | Create and Apply Tags you want to use for grouping rules.

STEP 3 | Assign a policy rule to a tag group.

1. Create a policy rule. Refer to Policy for more information on creating policy rules.
2. In the **Group Rules by Tag** field, select the tag from the drop-down and click **OK**.



3. **Commit** your changes.

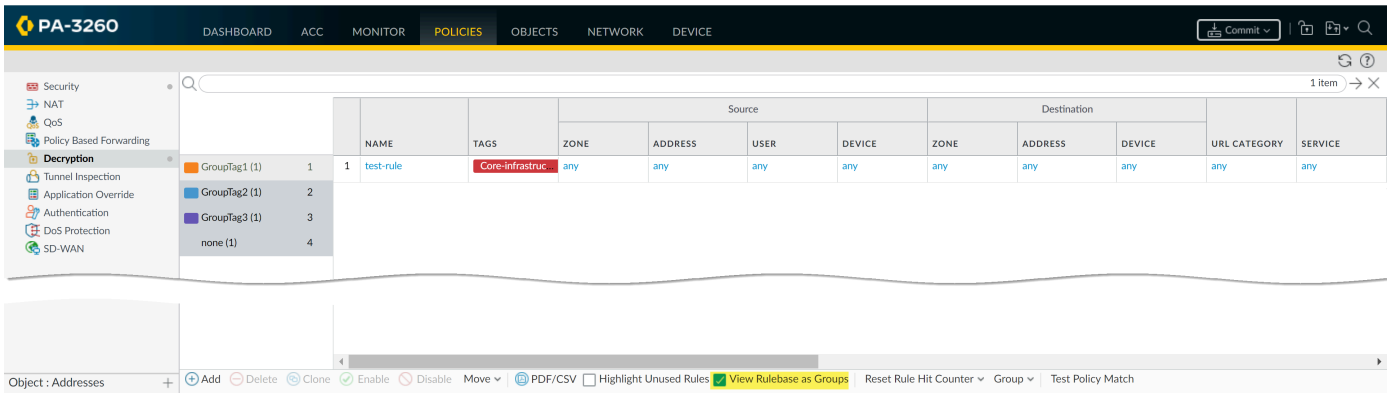
 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 4 | View your policy rulebase as groups.

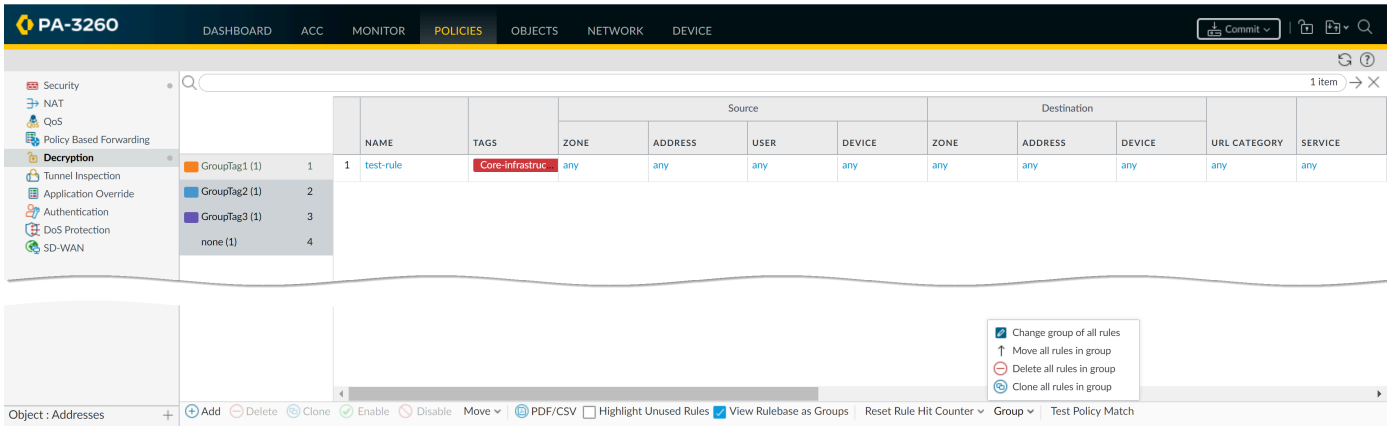
1. (Panorama only) From the **Device Group**, select the device group rulebase to view or view all Shared rules.
2. Click **Policies** and select the rulebase where you created the rules in Step 2.
3. Select the **View Rulebase as Groups** option (at the bottom).

*Rules not assigned a tag group display as **None**.*



STEP 5 | Perform Group operations as needed.

1. Click **Group** to perform group operations for rules in the selected tag group.
   - (Panorama only) **Move rules in group to a different rulebase or device group**—Move all policy rules in the selected tag group to the Pre-Rulebase or Post-Rulebase or move them to a different device group.
   - **Change group of all rules**—Move all rules in the selected tag group to a different tag group.
   - **Move all rules in group**—Move all rules in the selected tag group to change the rule priority order.
   - **Delete all rules in group**—Delete all rules in the selected tag group.
   - **Clone all rules in group**—Clone all rules in the selected tag group.



2. **Commit** your changes.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Use an External Dynamic List in Policy

An external dynamic list (formerly called dynamic block list) is a text file that you or another source hosts on an external web server so that the firewall can import objects—IP addresses, URLs, domains—to enforce policy on the entries in the list. As the list is updated, the firewall dynamically imports the list at the configured interval and enforces policy without the need to make a configuration change or a commit on the firewall.

- External Dynamic List
- Formatting Guidelines for an External Dynamic List
- Built-in External Dynamic Lists
- Configure the Firewall to Access an External Dynamic List
- Configure the Firewall to Access an External Dynamic List from the EDL Hosting Service
- Retrieve an External Dynamic List from the Web Server
- View External Dynamic List Entries
- Exclude Entries from an External Dynamic List
- Enforce Policy on an External Dynamic List
- Find External Dynamic Lists That Failed Authentication
- Disable Authentication for an External Dynamic List

External Dynamic List

An External Dynamic List is a text file that is hosted on an external web server so that the firewall can import objects—IP addresses, URLs, domains, International Mobile Equipment Identities (IMEIs), International Mobile Subscriber Identities (IMSIs)—included in the list and enforce policy. To enforce Security policy on the entries included in the external dynamic list, you must reference the list in a supported policy rule or profile. When multiple lists are referenced, you can prioritize the order of evaluation to make sure the most important EDLs are committed before capacity limits are reached. As you modify the list, the firewall dynamically imports the list at the configured interval and enforces policy without the need to make a configuration change or a commit on the firewall. If the web server is unreachable, the firewall uses the last successfully retrieved list for enforcing policy until the connection is restored with the web server. In cases where authentication to the EDL fails, the security policy stops enforcing the EDL. To retrieve the external dynamic list, the firewall uses the interface configured with the **Palo Alto Networks Services** service route.

The firewall supports these types of external dynamic lists:

- **Predefined IP Address**—A predefined IP address list is a type of IP address list that refers to the built-in, dynamic IP lists with fixed or “predefined” contents. These Built-In External Dynamic Lists—for bulletproof hosting providers, known malicious, and high-risk IP addresses—are automatically added to your firewall if you have an active Threat Prevention license. A predefined IP address list can also refer to an EDL that uses one of the built-in lists as a source. Because you can’t modify the contents of a predefined list, you can use a predefined list as a source for a different EDL if you want to add or exclude list entries.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

- **Predefined URL List**—This type of external dynamic list contains pre-populated URLs that applications use for background services, such as updates or Certificate Revocation List (CRL) checks, that the firewall can safely exclude from Authentication policy. Palo Alto Networks revises and maintains this type of external dynamic list, which is also known as an Authentication Portal Exclude List, through content updates.
- **IP Address**—The firewall typically enforces policy for a source or destination IP address that is defined as a static object on the firewall (see Enforce Policy on an External Dynamic List) If you need agility in enforcing policy for a list of source or destination IP addresses that emerge ad hoc, you can use an external dynamic list of type IP address as a source or destination address object in policy rules, and configure the firewall to deny or allow access to the IP addresses (IPv4 and IPv6 address, IP range and IP subnets) included in the list. You can also use an IP address EDL in the source or destination of an SD-WAN policy rule. The firewall treats an external dynamic list of type IP address as an address object; all the IP addresses included in a list are handled as one address object.
- **Domain**—This type of external dynamic list allows you to import custom domain names into the firewall to enforce policy using an Anti-Spyware profile or SD-WAN policy rule. An EDL in an Anti-Spyware profile is very useful if you subscribe to third-party threat intelligence feeds and want to protect your network from new sources of threat or malware as soon as you learn of a malicious domain. For each domain you include in the external dynamic list, the firewall creates a custom DNS-based spyware signature so that you can enable DNS sinkholing. The DNS-based spyware signature is of type spyware with medium severity and each signature is named **`Custom Malicious DNS Query <domain name>`**. You can also specify the firewall to include the subdomains of a specifed domain. For example, if your domain list includes paloaltonetworks.com, all lower level components of the domain name (e.g., *.paloaltonetworks.com) will also be included as part of the list. When this setting is enabled, each domain in a given list requires an additional entry, effectively doubling the number of entries used by the list. For details on configuring domain lists, see Configure DNS Sinkholing for a List of Custom Domains.
- **URL**—This type of external dynamic list gives you the agility to protect your network from new sources of threat or malware. The firewall handles an external dynamic list with URLs like a custom URL category and you can use this list in two ways:
  - As a match criterion in Security policy rules, Decryption policy rules, and QoS policy rules to allow, deny, decrypt, not decrypt, or allocate bandwidth for the URLs in the custom category.
  - In a URL Filtering profile where you can define more granular actions, such as continue, alert, or override, before you attach the profile to a Security policy rule (see Use an External Dynamic List in a URL Filtering Profile).
- **Equipment Identity**—You can reference an external dynamic list of IoT devices defined by International Mobile Equipment Identities (IMEIs) in a Security policy rule that controls traffic for equipment connected to a 5G or 4G network. Refer to the Mobile Network Infrastructure Getting Started for information about configuring Equipment ID security on supported firewall models.
- **Subscriber Identity**—You can reference an external dynamic list of International Mobile Subscriber Identities (IMSIs) in a Security policy rule that controls traffic for subscribers connected to a 5G or 4G network. Refer to the Mobile Network Infrastructure Getting Started for information about configuring Subscriber ID security on supported firewall models.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

On each firewall model, you can add a maximum of 30 custom EDLs with unique sources to a single policy rule to enforce policy. The external dynamic list limit is not applicable to Panorama. When using Panorama to manage a firewall that is enabled for multiple virtual systems, if you exceed the limit for the firewall, a commit error displays on Panorama. A source is a URL that includes the IP address or hostname, the path, and the filename for the external dynamic list. The firewall matches the URL (complete string) to determine whether a source is unique.

While the firewall does not impose a limit on the number of lists of a specific type, the following limits are enforced:

- IP address—The PA-5200 Series and the PA-7000 Series firewalls support a maximum of 150,000 total IP addresses; all other models support a maximum of 50,000 total IP addresses. No limits are enforced for the number of IP addresses per list. When the maximum supported IP address limit is reached on the firewall, the firewall generates a syslog message. The IP addresses in predefined IP address lists do not count toward the limit.
- URL and domain—The maximum number of URLs and domains supported varies by model. No limits are enforced for the number of URL or domain entries per list. Refer to the following table for specifics on your model:

Model	URL List Entry Limits	Domain List Entry Limits
PA-5200 Series, PA-7000 Series (upgraded with the PA-7000 20GXM NPC, PA-7000 20GQXM NPC, or the PA-7000 100G NPC). *PA-7000 appliances with mixed NPCs only support the standard capacities.*	250,000	4,000,000
VM-500, VM-700	100,000	2,000,000
PA-850, PA-820, PA-3200 Series	100,000	1,000,000
PA-7000 Series (and appliances upgraded with the PA-7000 20GQ NPC or the PA-7000 20G NPC), VM-300	100,000	500,000

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Model	URL List Entry Limits	Domain List Entry Limits
PA-220, VM-50, VM-50 (Lite), VM-100, VM-1000-HV	50,000	50,000

List entries only count toward the firewall limits if they belong to an external dynamic list that is referenced in policy.

- *When parsing the list, the firewall skips entries that do not match the list type, and ignores entries that exceed the maximum number supported for the model. To ensure that the entries do not exceed the limit, check the number of entries currently used in policy. Select **Objects** > **External Dynamic Lists** and click **List Capacities**.*
- *An external dynamic list must contain entries. If you want to stop using the list, remove the reference from the policy rule or profile instead leaving the list blank. If the list does not contain any entries, the firewall fails to refresh the list and continues to use the last information it retrieved.*
- *As a best practice, Palo Alto Networks recommends using shared EDLs when multiple virtual systems are used. Using individual EDLs with duplicate entries for each virtual system uses more memory, which might over-utilize firewall resources.*
- *EDL entry counts on firewalls operating multi-virtual systems take additional factors into account (such as DAGs, number of virtual systems, rules bases) to generate a more accurate capacity consumption listing. This might result in a discrepancy in capacity usage after upgrading from PAN-OS 8.x releases.*
- *Depending on the features enabled on the firewall, memory usage limits might be exceeded before EDL capacity limits are met due to memory allocation updates. As a best practice, Palo Alto Networks recommends reviewing EDL capacities and, when necessary, removing or consolidating EDLs into shared lists to minimize memory usage.*

Formatting Guidelines for an External Dynamic List

An external dynamic list of one type —IP address, URL or Domain—must include entries of that type only. The entries in a predefined IP address list comply with the formatting guidelines for IP address lists.

- IP Address List
- Domain List
- URL List

IP Address List

The external dynamic list can include individual IP addresses, subnet addresses (address/mask), or range of IP addresses. In addition, the block list can include comments and special characters such as ***** , **:** , **;** , **#**, or **/**. The syntax for each line in the list is **`[IP address, IP/Mask, or IP start range-IP end range] [space] [comment]`**.

Enter each IP address/range/subnet in a new line; URLs or domains are not supported in this list. A subnet or an IP address range, such as 92.168.20.0/24 or 192.168.20.40-192.168.20.50, count

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

as one IP address entry and not as multiple IP addresses. If you add comments, the comment must be on the same line as the IP address/range/subnet. The space at the end of the IP address is the delimiter that separates a comment from the IP address.

An example IP address list:

```
192.168.20.10/32
 2001:db8:123:1::1 #test IPv6 address
 192.168.20.0/24 ; test internal subnet
 2001:db8:123:1::/64 test internal IPv6 range
 192.168.20.40-192.168.20.50
```

For an IP address that is blocked, you can display a notification page only if the protocol is HTTP.

Domain List

You can use placeholder characters in domain lists to configure a single entry to match against multiple website subdomains, pages, including entire top-level domains, as well as matches to specific web pages.

Follow these guidelines when creating domain list entries:

- Enter each domain name in a new line; URLs or IP addresses are not supported in this list.
- Do not prefix the domain name with the protocol, http:// or https://.
- You can use an asterisk (*) to indicate a wildcard value.
- You can use a caret (^) to indicate an exact match value.
- The following characters are considered token separators: . / ? & = ; +

  Every string separated by one or two of these characters is a token. Use wildcard characters as token placeholders, indicating that a specific token can contain any value.
- Wildcard characters must be the only character within a token; however, an entry can contain multiple wildcards.
- Each domain entry can be up to 255 characters in length.

When to use the asterisk (*) wildcard:

Use an asterisk (*) wildcard to indicate one or multiple variable subdomains. For example, to specify enforcement for Palo Alto Network's website regardless of the domain extension used, which might be one or two subdomains depending on location, you would add the entry: **`*.paloaltonetworks.com`**. This entry would match to both docs.paloaltonetworks.com and support.paloaltonetworks.com.

You can also use this wildcard to indicate entire top-level domains. For example, to specify enforcement of a TLD named .work, you would add the entry **`*.work`**. This matches all websites ending with .work.

The (*) wildcard can only be prepended in domain entries.

Asterisk (*) examples

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

EDL Domain List Entry	Matching Sites
`*.company.com`	eng.tools.company.com support.tools.company.com tools.company.com docs.company.com
`*.click`	all websites ending with a top-level domain of .click.

When to use a caret (^) character:

Use carets (^) to indicate an exact match of a subdomain. For example, **`^paloaltonetworks.com`** matches only paloaltonetworks.com. This entry does not match to any other site.

Caret (^) examples

EDL Domain List Entry	Matching Site
`^company.com`	company.com
`^eng.company.com`	eng.company.com

URL List

See URL Category Exceptions.

Built-in External Dynamic Lists

With an active Threat Prevention license, Palo Alto Networks provides built-in IP address EDLs that you can use to protect against malicious hosts.

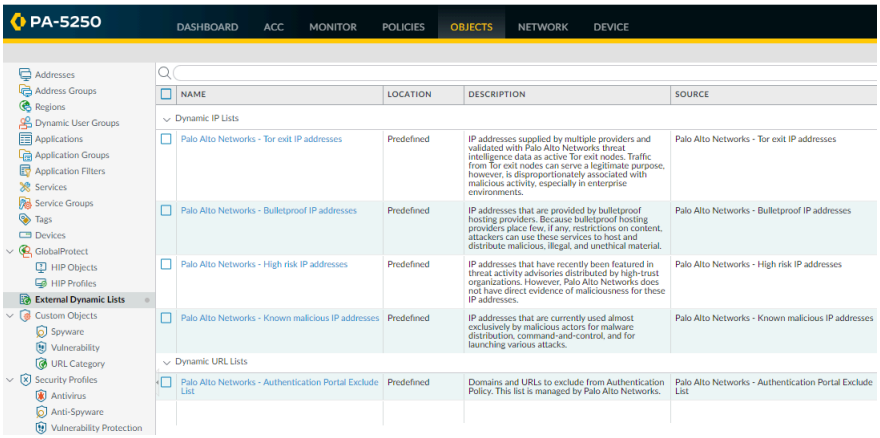
- **Palo Alto Networks Bulletproof IP Addresses**—Contains IP addresses provided by bulletproof hosting providers. Because bulletproof hosting providers place few, if any, restrictions on content, attackers frequently use these services to host and distribute malicious, illegal, and unethical material.
- **Palo Alto Networks High-Risk IP Addresses**—Contains malicious IP addresses from threat advisories issued by trusted third-party organizations. Palo Alto Networks compiles the list of threat advisories, but does not have direct evidence of the maliciousness of the IP addresses.
- **Palo Alto Networks Known Malicious IP Addresses**—Contains IP addresses that are verified malicious based on WildFire analysis, Unit 42 research, and data gathered from telemetry (Share Threat Intelligence with Palo Alto Networks). Attackers use these IP addresses almost exclusively to distribute malware, initiate command-and-control activity, and launch attacks.
- **Palo Alto Networks Tor Exit IP Addresses**—Contains IP addresses supplied by multiple providers and validated with Palo Alto Networks threat intelligence data as active

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Tor exit nodes. Traffic from Tor exit nodes can serve a legitimate purpose, however, is disproportionately associated with malicious activity, especially in enterprise environments.

The firewall receives updates for these feeds in content updates, allowing the firewall to automatically enforce policy based on the latest threat intelligence from Palo Alto Networks. You cannot modify the contents of the built-in lists. Use them as-is (see Enforce Policy on an External Dynamic List), or create a custom external dynamic list that uses one of the lists as a source (see Configure the Firewall to Access an External Dynamic List) and exclude entries from the list as needed.



Configure the Firewall to Access an External Dynamic List

You must establish the connection between the firewall and the source that hosts the external dynamic list before you can Enforce Policy on an External Dynamic List.

STEP 1 | (Optional) Customize the service route that the firewall uses to retrieve external dynamic lists.

Select **Device** > **Setup** > **Services** > **Service Route Configuration** > **Customize** and modify the **External Dynamic Lists** service route.

The firewall does not use the External Dynamic Lists service route to retrieve Built-in External Dynamic Lists; content updates modify or update the contents of those lists (active Threat Prevention license required).

https://t.me/learningnets

STEP 2 | Find an external dynamic list to use with the firewall.

- Create an external dynamic list and host it on a web server. Enter IP addresses, domains, or URLs in a blank text file. Each list entry must be on a separate line. For example:

  ```
  financialtimes.co.in

  www.wallaby.au/joey

  www.exyang.com/auto-tutorials/How-to-enter-Data-for-Success.aspx
  ```

  See the Formatting Guidelines for an External Dynamic List to ensure that the firewall does not skip list entries. To prevent commit errors and invalid entries, do not prefix http:// or https:// to any of the entries.
- Use an external dynamic list hosted by another source and verify that it follows the Formatting Guidelines for an External Dynamic List.

STEP 3 | Select **Objects** > **External Dynamic Lists**.

STEP 4 | Click **Add** and enter a descriptive **Name** for the list.

STEP 5 | (Optional) Select **Shared** to share the list with all virtual systems on a device that is enabled for multiple virtual systems. By default, the object is created on the virtual system that is currently selected in the **Virtual Systems** drop-down.

As a best practice, Palo Alto Networks recommends using shared EDLs when multiple virtual systems are used. Using individual EDLs with duplicate entries for each vsys uses more memory, which might over-utilize firewall resources.

STEP 6 | (Panorama only) Select **Disable override** to ensure that a firewall administrator cannot override settings locally on a firewall that inherits this configuration through a Device Group commit from Panorama.

STEP 7 | Select the list **Type** (for example, **URL List**).

Ensure that the list only includes entries for the list type. See Verify whether entries in the external dynamic list were ignored or skipped.

If you using a Domain List, you can optionally enable **Automatically expand to include subdomains** to also include the subdomains of a specified domain. For example, if your domain list includes paloaltonetworks.com, all lower level components of the domain name (e.g., *.paloaltonetworks.com) will also be included as part of the list. Keep in mind, when this setting is enabled, each domain in a given list requires an additional entry, effectively doubling the number of entries that are consumed.

STEP 8 | Enter the **Source** for the list you just created on the web server. The source must include the full path to access the list. For example, **`https://1.2.3.4/EDL_IP_2015`**.

- If you are creating a Predefined IP external dynamic list, select a Palo Alto Networks malicious IP address feed to use as a source.
- If you are creating a Predefined URL external dynamic list, select **panw-auth-portal-exclude-list** as a source.

 ©2022 Palo Alto Networks, Inc.

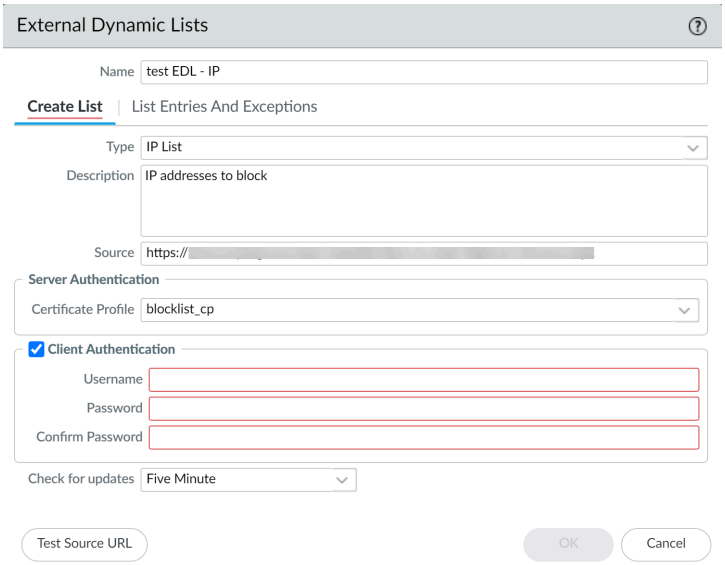
https://t.me/learningnets

STEP 9 | If the list source is secured with SSL (i.e. lists with an HTTPS URL), enable server authentication. Select a **Certificate Profile** or create a **New Certificate Profile** for authenticating the server that hosts the list. The certificate profile you select must have root certificate authority (CA) and intermediate CA certificates that match the certificates installed on the server you are authenticating.

Maximize the number of external dynamic lists that you can use to enforce policy. Use the same certificate profile to authenticate external dynamic lists from the same source URL. If you assign different certificate profiles to external dynamic lists from the same source URL, the firewall counts each list as a unique external dynamic list.

STEP 10 | Enable client authentication if the list source has an HTTPS URL and requires basic HTTP authentication for list access.

1. Select **Client Authentication**.
2. Enter a valid **Username** to access the list.
3. Enter the **Password** and **Confirm Password**.



STEP 11 | (Not available on Panorama or for Predefined URL EDLs) Click **Test Source URL** to verify that the firewall can connect to the web server.

*The **Test Source URL** function is not available when authentication is used for EDL access.*

STEP 12 | (Optional) Specify the frequency at which the firewall should **Check for updates** to the list. By default, the firewall retrieves the list once every hour and commits the changes.

The interval is relative to the last commit. So, for the five-minute interval, the commit occurs in 5 minutes if the last commit was an hour ago. To retrieve the list immediately, see Retrieve an External Dynamic List from the Web Server.

STEP 13 | Click **OK** and **Commit** your changes.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 14 | (Optional) EDLs are shown top to bottom, in order of evaluation. Use the directional controls at the bottom of the page to change the list order. This allows you to or order the lists to make sure the most important EDLs are committed before capacity limits are reached.

*You can only change the EDL order when **Group By Type** is deselected.*

STEP 15 | Enforce Policy on an External Dynamic List.

If the server or client authentication fails, the firewall ceases to enforce policy based on the last successfully retrieved external dynamic list. Find External Dynamic Lists That Failed Authentication *and view the reasons for authentication failure.*

Configure the Firewall to Access an External Dynamic List from the EDL Hosting Service

Configure the firewall to access an external dynamic list (EDL) from the EDL Hosting Service for Software-as-a-Service (SaaS) applications

- Create an External Dynamic List Using the EDL Hosting Service
- Convert the GlobalSign Root R1 Certificate to PEM Format

Create an External Dynamic List Using the EDL Hosting Service

Some Software-as-a-Service (SaaS) providers publish lists of IP addresses and URLs as destination endpoints for their SaaS applications. SaaS providers frequently update the SaaS applications destination endpoint lists as support grows and the service expands. This requires you to manually monitor the SaaS application endpoints for changes and manually update your policy configuration to ensure connectivity to these critical SaaS applications or set up an external tool to monitor and update your EDLs.

Configure an EDL using the EDL Hosting Service maintained by Palo Alto Networks to ease the operational burden of maintaining an EDL for a SaaS application. The EDL Hosting Service provides publicly available Feed URLs for SaaS application endpoints published by the SaaS application provider. Leveraging a Feed URL as the source in an EDL allows for dynamic enforcement of SaaS application traffic without the need for you to host and maintain your own EDL source.

Palo Alto Networks checks the application Feed URLs published by SaaS providers on a daily basis. For IP based feeds, Palo Alto Networks performs optimizations to combine entries from a continuous netmask and deduplication is performed if endpoints overlap across multiple areas. Additionally, the endpoints for the Microsoft 365 Common and Office Online SaaS application are always added to every Feed URL in the EDL Hosting Service.

Microsoft updates all Microsoft 365 Feed URLs at the end of each calendar month and provides a 30 day advanced notice prior to update. See the official Microsoft 365 Web Services page for more information. The EDL Hosting Service availability status and updates are posted to the Palo Alto Networks Cloud Services Status page.

STEP 1 | Visit the EDL Hosting Service and identify the Feed URL for your SaaS application.

Review the Microsoft 365 documentation for more information which Feed URL is best for your use case. Additionally, consider the SaaS application and location of users accessing

https://t.me/learningnets

the SaaS application when identifying a Feed URL to. For example, if you have a branch in Germany that only needs to access Exchange Online, select a Feed URL from the **Service Area: Exchange Online** for **Germany**.



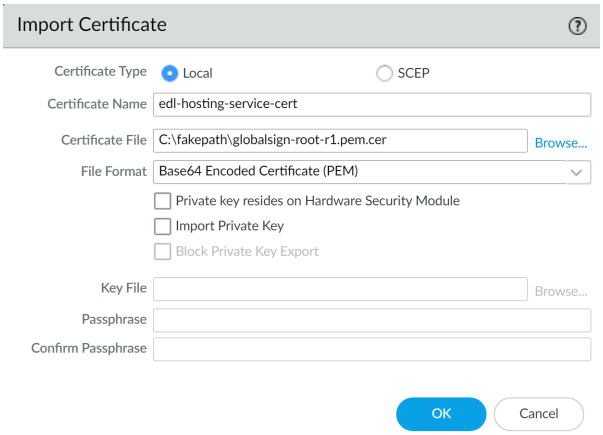
For a policy-based forwarding policy rule, use an IP-based Feed URL.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 2 | (Best Practices) Create a certificate profile to authenticate the EDL Hosting Service.

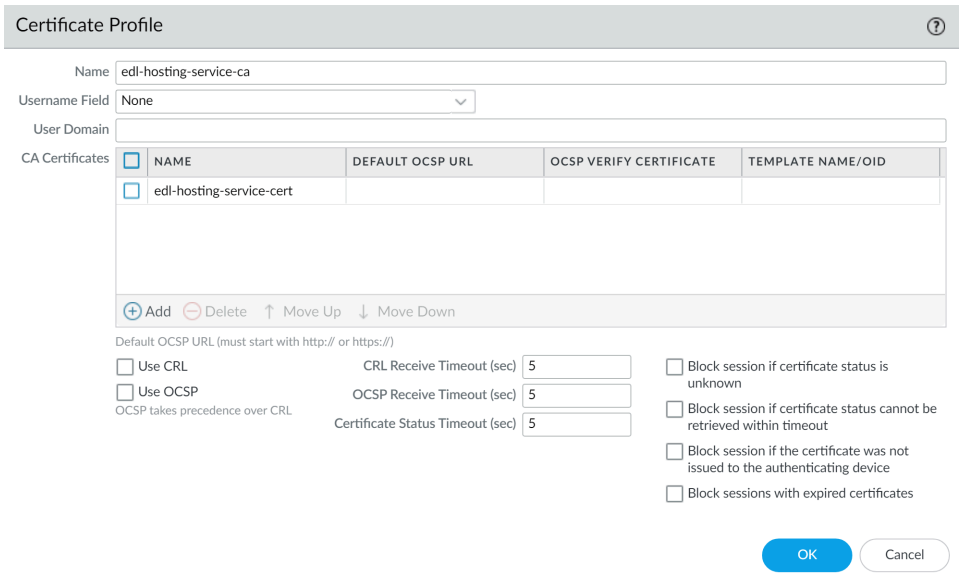
1. Download the GlobalSign Root R1 certificate.
2. Convert the GlobalSign Root R1 Certificate to PEM Format.
3. Launch the firewall web interface.
4. Import the GlobalSign Root R1 certificate.
   1. Select **Device** > **Certificate Management** > **Certificates** and **Import** a new certificate.
   2. For **Certificate Type**, select **Local**.
   3. Enter a descriptive **Certificate Name**.
   4. For the **Certificate File**, select **Browse** and select the certificate you converted in the previous step.
   5. For the **File Format**, select **Base64 Encoded Certificate (PEM)**.
   6. Click **OK**.



5. Create a certificate authority (CA) certificate profile.
   1. Select **Device** > **Certificate Management** > **Certificate Profile** and **Add** a new certificate profile.
   2. Enter a descriptive **Name**.
   3. For the **CA Certificates**, **Add** the certificate you imported in the previous step.
   4. Click **OK**.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets



6. **Commit**.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

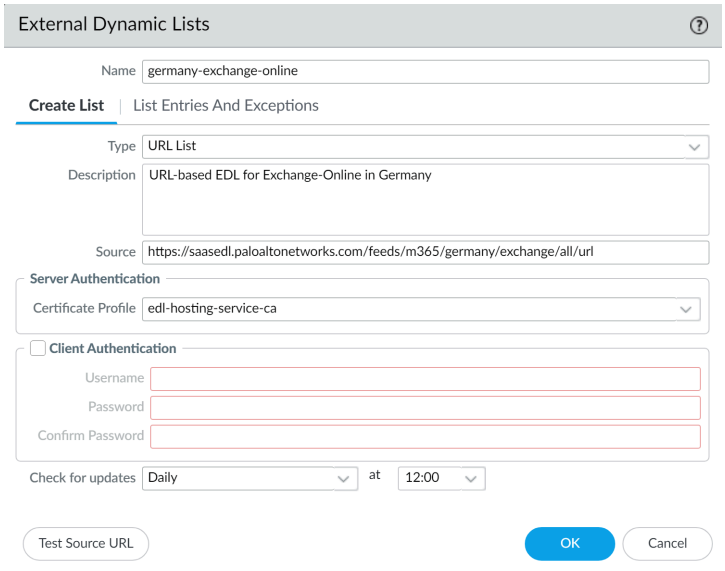
STEP 3 | Create an EDL using a Feed URL from the EDL Hosting Service.

1. Select **Objects** > **External Dynamic Lists** and **Add** a new EDL.
2. Enter a descriptive **Name** for the EDL.
3. Select the EDL **Type**.
   - For an IP-based EDL, select **IP List**.
   - For a URL-based EDL, select **URL List**.
4. (Optional) Enter a **Description for the EDL**
5. Enter the Feed URL as the EDL **Source**.

   Enforce all endpoints within a specific Feed URL. Adding an excluding a specific endpoint from a Feed URL can cause connectivity issues to the SaaS application.
6. (Best Practices) Select the **Certificate Profile** you created in the previous step.
7. Specify the frequency the firewall should **Check for updates** to match the update frequency of the Feed URL.

   For example, if the Feed URL is updated daily by Palo Alto Networks then configure the EDL to check for updates **Daily**.

   Palo Alto Networks displays the update frequency for each Feed URL in the EDL Hosting Service. Feed URLs are automatically updated with any new endpoints.
8. Click **Test Source URL** to verify that the firewall can access the Feed URL from the EDL Hosting Service.
9. Click **OK**.



https://t.me/learningnets

STEP 4 | Enforce Policy on an External Dynamic List.

When you enforce policy on an EDL from the EDL Hosting Service where the EDL is the source, be specific when configuring which users have access to the SaaS application to avoid over-provisioning access to the application.



Leverage App-ID *alongside EDLs in a policy rule for additional strict enforcement of SaaS application traffic.*

Convert the GlobalSign Root R1 Certificate to PEM Format

You must convert the GlobalSign Root R1 certificate to PEM format to create a certificate profile for authenticating the EDL Hosting Service. Creating the certificate profile to authenticate the EDL Hosting Service is a best practice when leveraging the EDL Hosting Service when you configure the firewall to access an external dynamic list from the EDL Hosting Service.

Refer to the appropriate procedure based on operating system of the device where you downloaded the GlobalSign Root R1 certificate.

STEP 1 | Download the GlobalSign Root R1 certificate if you have not already downloaded the certificate.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 2 | Convert the certificate.

- **Mac and Linux operating systems**

1. Open the terminal and convert the GlobalSign Root R1 certificate you downloaded.

```
admin:  openssl x509 -in <certificate-path>.crt -inform DER -out
 <target-export-path>.pem -outform PEM
```

```
admin-1@admin-1:~$ openssl x509 -in /home/admin-1/Downloads/Root-R1.crt -inform
DER -out /home/admin-1/Downloads/globalsign-root-r1.pem -outform PEM
```

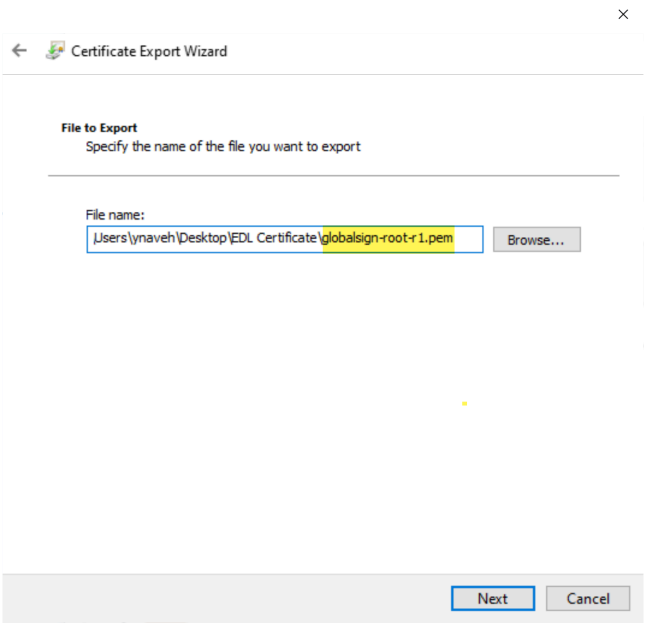
If no target export path is specified, the converted certificate is created on the device desktop.

- **Windows operating system**

1. Navigate to the location where you downloaded the GlobalSign Root1 certificate.
2. Double click and **Open** the certificate.
3. Click **Details** and **Copy to File**.

   Click **Next** when prompted to continue.
4. Select **Base-64 encoded x.509 (.CER)** and click **Next**
5. Click **Browse** to navigate to the location you want to copy the certificate and enter a name for the certificate that includes **`.pem`** appended to the end of file name. For example, **`globalsign-root-r1.pem`**

   Save the certificate. The **File Name** displayed shows the target export path and the certificate name you entered with `.cer` appended. Delete the appended`.cer`.



6. Click **Next** and **Finish** exporting the certificate.

https://t.me/learningnets

Retrieve an External Dynamic List from the Web Server

When you Configure the Firewall to Access an External Dynamic List, you can configure the firewall to retrieve the list from the web server on an hourly (default)five minute, daily, weekly, or monthly basis. If you have added or deleted IP addresses from the list and need to trigger an immediate refresh, use the following process to fetch the updated list.

STEP 1 | To retrieve the list on demand, select **Objects** > **External Dynamic Lists**.

STEP 2 | Select the list that you want to refresh, and click **Import Now**. The job to import the list is queued.

STEP 3 | To view the status of the job in the Task Manager, see Manage and Monitor Administrative Tasks.

STEP 4 | (Optional) After the firewall retrieves the list, View External Dynamic List Entries.

View External Dynamic List Entries

Before you Enforce Policy on an External Dynamic List, you can view the contents of an external dynamic list directly on the firewall to check if it contains certain IP addresses, domains, or URLs. The entries displayed are based on the version of the external dynamic list that the firewall most recently retrieved.

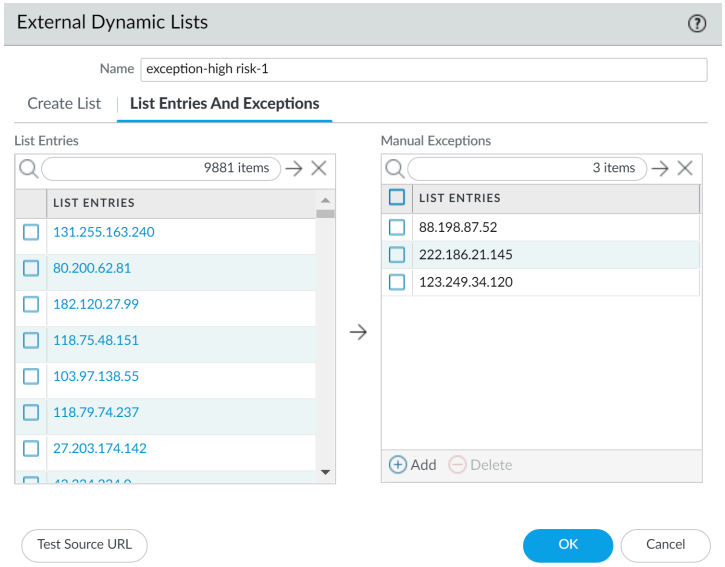
STEP 1 | Select **Objects** > **External Dynamic Lists**.

STEP 2 | Click the external dynamic list you want to view.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 3 | Click **List Entries and Exceptions** and view the objects that the firewall retrieved from the list.



The list might be empty if:

- The EDL has not yet been applied to a Security policy rule. To apply an EDL to a Security policy rule and populate the EDL, see Enforce Policy on an External Dynamic List.
- The firewall has not yet retrieved the external dynamic list. To force the firewall to retrieve an external dynamic list immediately, Retrieve an External Dynamic List from the Web Server.
- The firewall is unable to access the server that hosts the external dynamic list. Click **Test Source URL** to verify that the firewall can connect to the server.

STEP 4 | Enter an IP address, domain, or URL (depending on the type of list) in the filter field and Apply Filter (→) to check if it's in the list. Exclude Entries from an External Dynamic List based on which IP addresses, domains, and URLs you need to block or allow.

STEP 5 | (Optional) View the AutoFocus Intelligence Summary for a list entry. Hover over an entry to open the drop-down and then click **AutoFocus**.

Exclude Entries from an External Dynamic List

As you view the entries of an external dynamic list, you can exclude up to 100 entries from the list. The ability to exclude entries from an external dynamic list gives you the option to enforce policy on some (but not all) of the entries in a list. This is helpful if you cannot edit the contents of an external dynamic list (such as the Palo Alto Networks High-Risk IP Addresses feed) because it comes from a third-party source.

STEP 1 | View External Dynamic List Entries.

https://t.me/learningnets

STEP 2 | Select up to 100 entries to exclude from the list and click Submit (→) or manually **Add** a list exception.

- You cannot save your changes to the external dynamic list if you have duplicate entries in the Manual Exceptions list. To identify duplicate entries, look for entries with a red underline.
- A manual exception must match a list entry exactly. Additionally, you cannot exclude a specific IP address from within an IP address range. To exclude a specific IP address from an IP address range, you must add each IP address in the range as a list entry and then exclude the desired IP address.

  The firewall does not support excluding an individual IP address from an IP address range.

STEP 3 | Click **OK** and **Commit** to save your changes.

STEP 4 | (Optional) Enforce Policy on an External Dynamic List.

Enforce Policy on an External Dynamic List

Block or allow traffic based on IP addresses or URLs in an external dynamic list, or use a dynamic domain list with a DNS sinkhole to prevent access to malicious domains.

Tips for enforcing policy on the firewall with external dynamic lists:

- *When viewing external dynamic lists on the firewall (**Objects** > **External Dynamic Lists**), click **List Capacities** to compare how many IP addresses, domains, and URLs are currently used in policy with the total number of entries that the firewall supports for each list type.*
- Use Global Find to Search the Firewall or Panorama Management Server *for a domain, IP address, or URL that belongs to one or more external dynamic lists is used in policy. This is useful for determining which external dynamic list (referenced in a Security policy rule) is causing the firewall to block or allow a certain domain, IP address, or URL.*
- *Use the directional controls at the bottom of the page to change the evaluation order of EDLs. This allows you to or order the lists to make sure the most important entries in an EDL are committed before capacity limits are reached.*

  *You can only change the EDL order when **Group By Type** is deselected.*

- Configure DNS Sinkholing for a List of Custom Domains.
- Use an External Dynamic List in a URL Filtering Profile.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

- **Use an External Dynamic List of Type URL as Match Criteria in a Security Policy Rule.**
  1. Select **Policies** > **Security**.
  2. Click **Add** and enter a descriptive **Name** for the rule.
  3. In the **Source** tab, select the **Source Zone**.
  4. In the **Destination** tab, select the **Destination Zone**.
  5. In the **Service/URL Category** tab, click **Add** to select the appropriate external dynamic list from the URL Category list.
  6. In the **Actions** tab, set the **Action Setting** to **Allow** or **Deny**.
  7. Click **OK** and **Commit**.
  8. Verify whether entries in the external dynamic list were ignored or skipped.

     Use the following CLI command on a firewall to review the details for a list.

     ```
     request
     system external-list show type <domain | ip | url>
      name_of_list
     ```

     For example:

     ```
     request system
     external-list show type url EBL_ISAC_Alert_List
     ```

  9. Test that the policy action is enforced.
     1. View External Dynamic List Entries for the URL list, and attempt to access a URL from the list.
     2. Verify that the action you defined is enforced.
     3. To monitor the activity on the firewall:
        - Select **ACC** and add a URL Domain as a global filter to view the Network Activity and Blocked Activity for the URL you accessed.
        - Select **Monitor** > **Logs** > **URL Filtering** to access the detailed log view.

 ©2022 Palo Alto Networks, Inc.

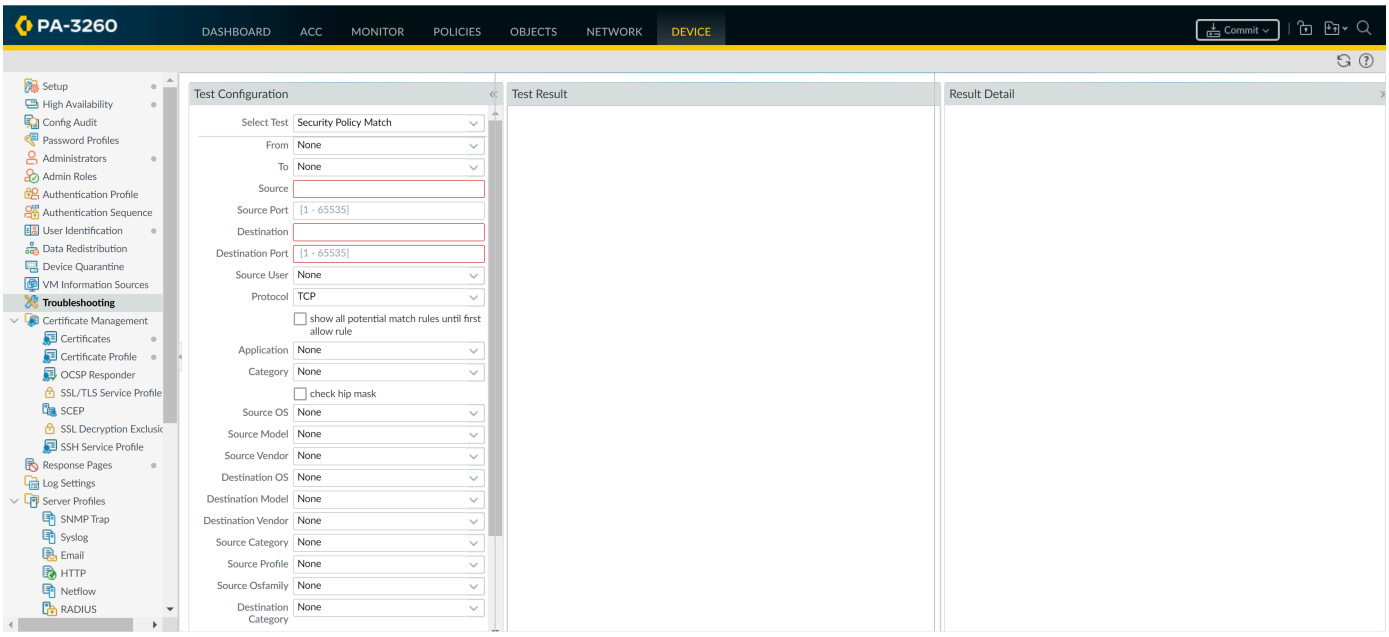
https://t.me/learningnets

- **Use an IP External Dynamic List or Predefined IP External Dynamic List as a Source or Destination Address Object in a Security Policy Rule.**

  This capability is useful if you deploy new servers and want to allow access to the newly deployed servers without requiring a firewall commit.

  1. Select **Policies** > **Security**.
  2. Click **Add** and give the rule a descriptive **Name**.
  3. In the **Source/Destination** tabs, set the external dynamic list to be used as the **Source/ Destination Address**(es).
  4. In the Service/ URL Category tab, make sure the **Service** is set to **application-default**.
  5. In the Actions tab, set the Action Setting to **Allow** or **Deny**.

     Create separate external dynamic lists if you want to specify allow and deny actions for specific IP addresses.
  6. Leave all the other options at the default values.
  7. Click **OK** to save the changes.
  8. **Commit** the changes.
  9. Test that the policy action is enforced.
     1. View External Dynamic List Entries for the external dynamic list, and attempt to access an IP address from the list.
     2. Verify that the action you defined is enforced.
     3. Select **Monitor** > **Logs** > **Traffic** and view the log entry for the session.
     4. To verify the policy rule that matches a flow, select **Device** > **Troubleshooting**, and execute a Security Policy Match test:



https://t.me/learningnets

- **Use a Predefined URL External Dynamic List to exclude benign domains that applications use for background traffic from Authentication policy.**

  When you select the **panw-auth-portal-exclude-list** EDL type, you can easily exclude from Authentication policy enforcement the domains that many applications use for background traffic, such as updates and other trusted services. This ensures that the firewall does not block the necessary traffic for these services and application maintenance is not interrupted.

  1. Select **Policies** > **Authentication**.
  2. On the **Service/URL Category** tab, select the Predefined URL EDL as the **URL Category**.
  3. On the **Actions** tab, select **`default-no-captive-portal`** as the **Authentication Enforcement**.
  4. Click **OK**.
  5. **Move** the rule to the top so that it is the first rule in the policy.
  6. **Commit** your changes.

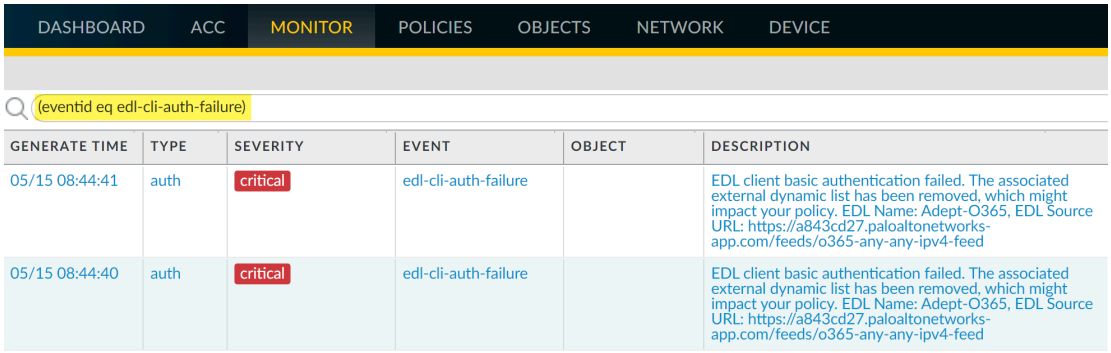
Find External Dynamic Lists That Failed Authentication

When an external dynamic list that requires SSL fails client or server authentication, the firewall generates a system log of critical severity. The log is critical because the firewall continues to enforce policy based on the last successful external dynamic list after it fails authentication, instead of using the latest version. Use the following process to view critical system log messages notifying you of authentication failure related to external dynamic lists.

STEP 1 | Select **Monitor** > **Logs** > **System**.

STEP 2 | Construct the following filters to view all messages related to authentication failure, and apply the filters. For more information, review the complete workflow to Filter Logs.

- Server authentication failure—**`(eventid eq tls-edl-auth-failure)`**
- Client authentication failure—**`(eventid eq edl-cli-auth-failure)`**



STEP 3 | Review the system log messages. The message description includes the name of the external dynamic list, the source URL for the list, and the reason for the authentication failure.

The server that hosts the external dynamic list fails authentication if the certificate is expired. If you have configured the certificate profile to check certificate revocation status via Certificate

https://t.me/learningnets

Revocation List (CRL) or Online Certificate Status Protocol (OCSP), the server may also fail authentication if:

- The certificate is revoked.
- The revocation status of the certificate is unknown.
- The connection times out as the firewall is attempting to connect to the CRL/OCSP service.

For more information on certificate profile settings, refer to the steps to Configure a Certificate Profile.

Verify that you added the root CA and intermediate CA of the server to the certificate profile configured with the external dynamic list. Otherwise, the firewall will not authenticate the list properly.

Client authentication fails if you have entered the incorrect username and password combination for the external dynamic list.

STEP 4 | (Optional) Disable Authentication for an External Dynamic List that failed authentication as a stop-gap measure until the list owner renews the certificate(s) of the server that hosts the list.

Disable Authentication for an External Dynamic List

Palo Alto Networks recommends that you enable authentication for the servers that host the external dynamic lists configured on your firewall. However, if you Find External Dynamic Lists That Failed Authentication and prefer to disable server authentication for those lists, you can do so through the CLI. The procedure below only applies to external dynamic lists secured with SSL (i.e., lists with an HTTPS URL); the firewall does not enforce server authentication on lists with an HTTP URL.

Disabling server authentication for an external dynamic list also disables client authentication. With client authentication disabled, the firewall will not be able to connect to an external dynamic list that requires a username and password for access.

STEP 1 | Launch the CLI and switch to configuration mode as follows:

```
username@hostname> configure
Entering configuration mode
[edit]
username@hostname#
```

The change from the > to the # symbol indicates that you are now in configuration mode.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 2 | Enter the appropriate CLI command for the list type:

- IP Address

```
set external-list <external dynamic list name> type ip
 certificate-profile None
```

- Domain

```
set external-list <external dynamic list name> type domain
 certificate-profile None
```

- URL

```
set external-list <external dynamic list name> type url
 certificate-profile None
```

STEP 3 | Verify that authentication is disabled for the external dynamic list.

Trigger a refresh for the list (see Retrieve an External Dynamic List from the Web Server). If the firewall retrieves the list successfully, server authentication is disabled.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Register IP Addresses and Tags Dynamically

To mitigate the challenges of scale, lack of flexibility, and performance, network architectures today allow for virtual machines (VMs) and applications to be provisioned, changed, and deleted on demand. This agility, though, poses a challenge for security administrators because they have limited visibility into the IP addresses of the dynamically provisioned VMs and the plethora of applications that can be enabled on these virtual resources.

Firewalls (hardware-based and VM-Series models) support the ability to register IP addresses, IP sets (IP ranges and subnets), and tags dynamically. The IP addresses and tags can be registered on the firewall directly or from Panorama. You can also automatically remove tags on the source and destination IP addresses included in a firewall log.

PAN-OS only supports IPv4 IP subnets and ranges in dynamic address groups.

You can enable the dynamic registration process using any of the following options:

- **User-ID agent for Windows**—In an environment where you've deployed the User-ID agent, you can enable the User-ID agent to monitor up to 100 VMware ESXi servers, vCenter Servers, or a combination of the two. As you provision or modify virtual machines on these VMware servers, the agent can retrieve the IP address changes and share them with the firewall.
- **VM Information Sources**—Enables you to monitor VMware ESXi, vCenter Server, AWS-VPCs, and Google Compute Engines natively on the firewall and to retrieve IP address changes when you provision or modify virtual machines on these sources. VM Information Sources option polls for a predefined set of attributes and does not require external scripts to register the IP addresses through the XML API. See Monitor Changes in the Virtual Environment.
- **Panorama Plugin**—You can enable a Panorama™ M-Series or virtual appliance to connect to your Azure or AWS public cloud environment and retrieve information on the virtual machines deployed within your subscription or VPC. Panorama then registers the VM information to the managed Palo Alto Networks firewalls that you configured for notification and then you can use these attributes to define dynamic address groups and attach them to Security policy rules to allow or deny traffic to and from these VMs.
- **VMware Service Manager** (Integrated NSX solutions only)—The integrated NSX solution is designed for automated provisioning and distribution of the Palo Alto Networks Next-Generation Security Operating Platform® and the delivery of dynamic context-based Security policies using Panorama. The NSX Manager updates Panorama with the latest information on the IP addresses, IP sets, and tags associated with the virtual machines deployed in this integrated solution. For information on this solution, see Set Up a VM-Series NSX Edition Firewall.
- **XML API**—The firewall and Panorama support an XML API that uses standard HTTP requests to send and receive data. You can use this API to register IP addresses and tags with the firewall or Panorama. You can make API calls directly from command-line utilities, such as cURL, or by using any scripting or application framework that supports REST-based services. Refer to the PAN-OS XML API Usage Guide for details.
- **Auto-Tag**—Tag the source or destination IP address automatically when a log is generated on the firewall and register the IP address and tag mapping to a User-ID agent on the firewall or on Panorama, or to a remote User-ID agent using an HTTP server profile. For example, whenever

https://t.me/learningnets

the firewall generates a threat log, you can configure the firewall to tag the source IP address in the threat log with a specific tag name. For more information, refer to Use Auto-Tagging to Automate Security Actions.

Additionally, you can configure the firewall to dynamically unregister a tag after a configured amount of time using a timeout. For example, you can configure the timeout to be the same duration as the DHCP lease timeout for the IP address. This allows the IP address-to-tag mapping to expire at the same time as the DHCP lease so that you don't unintentionally apply policy when the IP address is reassigned.

See Forward Logs to an HTTP(S) Destination.

For information on creating and using Dynamic Address Groups, see Use Dynamic Address Groups in Policy.

For the CLI commands for registering tags dynamically, see CLI Commands for Dynamic IP Addresses and Tags.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Use Dynamic User Groups in Policy

Dynamic user groups help you to create policy that provides auto-remediation for anomalous user behavior and malicious activity while maintaining user visibility. After you create the group and commit the changes, the firewall registers the users and associated tags then automatically updates the dynamic user group's membership. Because updates to dynamic user group membership are automatic, using dynamic user groups instead of static group objects allows you to respond to changes in user behavior or potential threats without manual policy changes.

To determine what users to include as members, a dynamic user group uses tags as filtering criteria. As soon as a user matches the filtering criteria, that user becomes a member of the dynamic user group. The tag-based filter uses logical *and* and *or* operators. Each tag is a metadata element or attribute-value pair that you register on the source statically or dynamically. Static tags are part of the firewall configuration, while dynamic tags are part of the runtime configuration. As a result, you don't need to commit updates to dynamic tags if they are already associated with a policy that you have committed on the firewall

To dynamically register tags, you can use:

- the XML API
- the User-ID agent
- Panorama
- the web interface on the firewall

The firewall redistributes the tags for the dynamic user group to the listening redistribution agents, which includes other firewalls, Panorama, or a Dedicated Log Collector, as well as Cortex applications.



To support redistribution for dynamic user group tags, all firewalls must use PAN-OS 9.1 to receive the tags from the registration sources.

The firewall redistributes the tags for the dynamic user group to the next hop and you can configure log forwarding to send the logs to a specific server. Log forwarding also allows you to use auto-tagging to automatically add or remove members of dynamic user groups based on events in the logs.

STEP 1 | Select **Objects** > **Dynamic User Groups** and **Add** a new dynamic user group.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 2 | Define the membership of the dynamic user group.

1. Enter a **Name** for the group.
2. (Optional) Enter a **Description** for the group.
3. **Add Match Criteria** using dynamic tags to define the members in the dynamic user group.
4. (Optional) Use the **And** or **Or** operators with the tag(s) that you want to use to filter for or match against.
5. Click **OK**.
6. (Optional) Select the **Tags** you want to assign to the group itself.

   *This tag displays in the **Tags** column in the **Dynamic User Group** list and defines the dynamic group object, not the members in the group.*
7. Click **OK** and **Commit** your changes.

   If you update the user group object filter, you must commit the changes to update the configuration.

STEP 3 | Depending on the log information that you want to use as match criteria, configure auto-tagging by creating a log forwarding profile or configuring the log settings.

- For Authentication, Data, Threat, Traffic, Tunnel Inspection, URL, and WildFire logs, create a log forwarding profile.
- For User-ID, GlobalProtect, and IP-Tag logs, configure the log settings.

STEP 4 | (Optional) To return dynamic user group members to their original groups after a specific duration of time, enter a **Timeout** value in minutes (default is 0, range is 0-43200).

STEP 5 | Use the dynamic user group in a policy to regulate traffic for the members of the group.

You will need to create at least two rules: one to allow initial traffic to populate the dynamic user group and one to deny traffic for the activity you want to prevent. To tag users, the rule to allow traffic must have a higher rule number in your rulebase than the rule that denies traffic.

1. Select the dynamic user group from Step 1 as the **Source User**.
2. Create the rule where the **Action** denies traffic to the dynamic user group members.
3. Create the rule that allows the traffic to populate the dynamic user group members.
4. If you configured a **Log Forwarding** profile in Step 3, select it to add it to the policy.
5. **Commit** your changes.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 6 | (Optional) Refine the group's membership and define the registration source for the user-to-tag mapping updates.

If the initial user-to-tag mapping retrieves users who should not be members or if it does not include users who should be, modify the members of the group to include the users for whom you want to enforce the policy and specify the source for the mappings.

1. In the **Users** column, select **more**.
2. **Register Users** to add them to the group and select the **Registration Source** for the tags and user-to-tag mappings.
   - **Local** (Default)—Register the tags and mappings for the dynamic user group members locally on the firewall.
   - **Panorama User-ID Agent**—Register the tags and mappings for the dynamic user group members on a User-ID agent connected to Panorama. If the dynamic user group originates from Panorama, the row displays in yellow and the group name, description, match criteria, and tags are read-only. However, you can still register or unregister users from the group.
   - **Remote device User-ID Agent**—Register the tags and mappings for the dynamic user group members on a remote User-ID agent. To select this option, you must first configure an HTTP server profile.
3. Select the **Tags** you want to register on the source using the tag(s) you used to configure the group.
4. (Optional) To return dynamic user group members to their original groups after a specific duration of time, enter a **Timeout** value in minutes (default is 0, range is 0-43200).
5. **Add** or **Delete** users as necessary.
6. (Optional) **Unregister Users** to remove their tags and user-to-tag mappings.

STEP 7 | Verify the firewall correctly populates the users in the dynamic user group.

1. Confirm the **Dynamic User Group** column in the Traffic, Threat, URL Filtering, WildFire Submissions, Data Filtering, and Tunnel Inspection logs displays the dynamic user groups correctly.
2. Use the **`show user group list dynamic`** command to display a list of all dynamic user groups as well as the total number of dynamic user groups.
3. Use the **`show object registered-user all`** command to display a list of users who are registered members of dynamic user groups.
4. Use the **`show user group name`** ***`group-name`*** command to display information about the dynamic user group, such as the source type.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Use Auto-Tagging to Automate Security Actions

Auto-tagging allows the firewall or Panorama to tag a policy object when it receives a log that matches specific criteria and establish IP address-to-tag or user-to-tag mapping. For example, when the firewall generates a threat log, you can configure the firewall to tag the source IP address or source user in the threat log with a specific tag name. You can then use these tags to automatically populate policy objects such as dynamic user groups or dynamic address groups, which can then be used to automate security actions in security, authentication, or decryption policies. For example, when you create a filter for the URL logs for `yes` in the **Credential Detected** column, you can apply a tag to the user that enforces an authentication policy that requires user to authenticate using multi-factor authentication (MFA).

Dynamic user groups do not support auto-tagging from HIP Match logs.

Redistribute the mappings across your network by registering the IP address-to-tag and user-to-tag mappings to a PAN-OS integrated User-ID agent on the firewall or Panorama or to a remote User-ID agent using an HTTP server profile. The firewall can automatically remove (unregister) a tag associated with an IP address or user when you configure a timeout as part of a built-in action for a log forwarding profile or as part of log forwarding settings. For example, if the firewall detects a user has potentially compromised credentials, you could configure the firewall to require MFA authentication for that user for a given period of time, then configure a timeout to remove the user from the MFA requirement group.

STEP 1 | Depending on the type of log you want to use for tagging, create a log forwarding profile or configure the log settings to define how you want the firewall or Panorama to handle logs.

- For Authentication, Data, Threat, Traffic, Tunnel Inspection, URL, and WildFire logs, create a log forwarding profile.
- For User-ID, GlobalProtect, and IP-Tag logs, configure the log settings.

STEP 2 | Define the match list criteria that determine when the firewall or Panorama adds the tag to the policy object.

For example, you can use a filter to configure a threshold or define a value (such as **`user eq "unknown"`** to identify users that the firewall has not yet mapped); when the firewall reaches that threshold or finds that value, the firewall adds the tag.

- To create a log forwarding profile, **Add** it and select the **Log Type** you want to monitor for match list criteria (**Objects** > **Log Forwarding**).
- To configure log settings, **Add** the log settings for the type of log you want to monitor for match list criteria (**Device** > **Log Settings**).

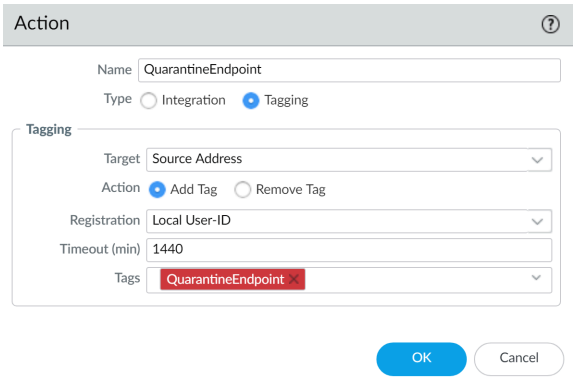
STEP 3 | Copy and paste a **Filter** value or use the **Filter Builder** to define the match criteria for the tag.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 4 | Add a built-in action to tag the policy object.

1. **Add** the **Built-in Actions** you want the firewall or Panorama to take when the logs contain an entry that meets the match list criteria.
2. **Name** the action.
3. Select the type of **Target** that you want to tag (**Destination Address**, **Source Address**, **User**, or **X-Forwarded-For Address**).
4. Confirm that **Add Tag** is the **Action**.
5. Select the **Registration** source for the tag to determine how the firewall or Panorama redistributes the IP address-to-tag mapping.
   - **Local User-ID**—Redistribute the IP address-to-tag mapping on the User-ID agent on the firewall or Panorama.
   - **Panorama User-ID**—Redistribute the IP address-to-tag mapping on Panorama.
   - **Remote User-ID**—Redistribute the IP address-to-tag mapping on another User-ID agent using an HTTP server profile. If you select this option, you must configure an HTTP server profile (see Step 5).
6. Enter or select the **Tags** you want to add to the policy object.
   You may need to click outside of the field or press Enter to enable the **OK** button.
7. Click **OK**.

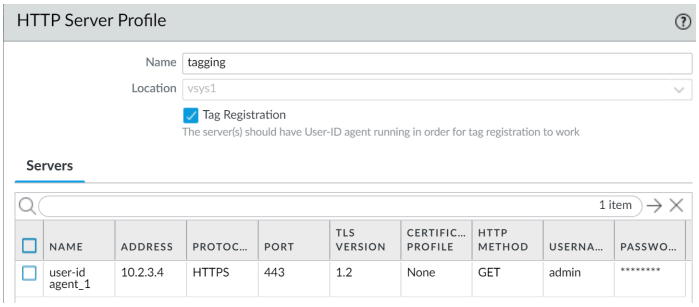


 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 5 | (Remote User-ID only) Configure an HTTP server profile to forward logs to a remote User-ID agent.

1. Select **Device** > **Server Profiles** > **HTTP**.
2. **Add** a profile and specify a **Name** for the server profile.
3. (Virtual systems only) Select the **Location**. The profile can be **Shared** across all virtual systems or can belong to a specific virtual system.
4. Select **Tag Registration** to enable the firewall to register the IP address and tag mapping with the User-ID agent on a remote firewall. With tag registration enabled, you cannot specify the payload format.
5. **Add** the server connection details to access the remote User-ID agent and click **OK**.



6. Select the log forwarding profile you created then select this server profile as the HTTP server profile for your **Remote User-ID** tag **Registration**.

STEP 6 | Define the policy objects to which you want to apply the tags.

1. Create or select one of the following policy objects: dynamic address groups, Use Dynamic User Groups in Policy, addresses, address groups, zones, policy rules, services, or service groups.
2. Enter the tags you want to apply to the object as the **Match** criteria.

   Confirm that the tag is identical to the tag in Step 4.

STEP 7 | Add the tagged policy objects to your policy.

This workflow uses a Security policy as an example, but you can also use tagged policy objects in Authentication policy.

1. Select **Policies** > **Security**.
2. Click **Add** and enter a **Name** and optionally a **Description** for the policy.
3. Add the **Source Zone** where the traffic originates.
4. Add the **Destination Zone** where the traffic terminates.
5. Select the **Source** object you created in Step 5.1.
6. Select whether the rule will **Allow** or **Deny** the traffic.

STEP 8 | If you configured a log forwarding profile, assign it to your Security policy.

You can assign one log forwarding profile for each policy but you can assign multiple methods and actions per profile. For an example, refer to Use Dynamic Address Groups in Policy.

STEP 9 | **Commit** your changes.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 10 | (Optional) Configure a timeout to remove the tag from the policy object after the specified time has elapsed.

Specify the amount of time (in minutes) that passes before the firewall removes the tag from the policy object. The range is from 0 to 43,200. If you set the timeout to zero, the IP address-to-tag mapping does not timeout and must be removed with an explicit action. If you set the timeout to the maximum of 43,200 minutes, the firewall removes the tag after 30 days.

*You cannot configure a Timeout with a **Remove Tag** action.*

1. Select the log forwarding profile.
2. **Add** or edit one of the **Built-in Actions**.
3. Specify the **Timeout** (in minutes). When the specified time has elapsed, the firewall or Panorama removes the tag.

   Set the IP-tag timeout to the same amount of time as the DHCP lease timeout for that IP address. This allows the IP address-to-tag mapping to expire at the same time as the DHCP lease so that you do not unintentionally apply policy when the IP address is reassigned.
4. Click **OK** and **Commit** your changes.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Monitor Changes in the Virtual Environment

To secure applications and prevent threats in an environment where new users and servers are constantly emerging, your security policy must be nimble. To be nimble, the firewall must be able to learn about new or modified IP addresses and consistently apply policy without requiring configuration changes on the firewall.

This capability is provided by the coordination between the **VM Information Sources** and **Dynamic Address Groups** features on the firewall. The firewall and Panorama provide an automated way to gather information on the virtual machine (or guest) inventory on each monitored source and create policy objects that stay in sync with the dynamic changes on the network.

- Enable VM Monitoring to Track Changes on the Virtual Network
- Attributes Monitored on Virtual Machines in Cloud Platforms
- Use Dynamic Address Groups in Policy

Enable VM Monitoring to Track Changes on the Virtual Network

VM information sources provides an automated way to gather information on the Virtual Machine (VM) inventory on each monitored source (host); the firewall can monitor the VMware ESXi, vCenter Server, AWS-VPC, Microsoft Azure VNet, and Google Cloud. As virtual machines (guests) are deployed or moved, the firewall collects a predefined set of attributes (or metadata elements) as tags; these tags can then be used to define Dynamic Address Groups (see Use Dynamic Address Groups in Policy) and matched against in policy.

You can directly configure the firewall or use Panorama templates to monitor up to 10 VM information sources. **VM Information Sources** offers easy configuration and enables you to monitor a predefined set of 16 metadata elements or attributes. See Attributes Monitored on Virtual Machines in Cloud Platforms for the list. By default, the traffic between the firewall and the monitored sources uses the management (MGT) port on the firewall.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

- *When monitoring ESXi hosts that are part of the VM-Series NSX edition solution, use Dynamic Address Groups instead of using VM Information Sources to learn about changes in the virtual environment. For the VM-Series NSX edition solution, the NSX Manager provides Panorama with information on the NSX security group to which an IP address belongs. The information from the NSX Manager provides the full context for defining the match criteria in a Dynamic Address Group because it uses the service profile ID as a distinguishing attribute and allows you to properly enforce policy when you have overlapping IP addresses across different NSX security groups. Up to a maximum of 32 tags (from vCenter server and NSX Manager) that can be registered to an IP address.*
- *For monitoring the virtual machines within your Azure deployment, instead of VM Monitoring Sources, you need to deploy the VM Monitoring script that runs on a virtual machine within the Azure public cloud. This script collects the IP address-to-tag mapping information for your Azure assets and publishes it to the firewalls and corresponding virtual systems you specify in the script.*
- *For Panorama version 8.1.3 and later, you can also use the Panorama plugin for AWS or Azure to retrieve VM Information and register it to the managed firewalls. See Attributes Monitored on Virtual Machines in Cloud Platformsfor details.*

STEP 1 | Enable VM Monitoring.

You can configure up to 10 VM information sources for each firewall, or for each virtual system on a multiple virtual systems capable firewall.

If your firewalls are configured in a high availability configuration:

- In an active/passive setup, only the active firewall monitors the VM sources.
- In an active/active setup, only the firewall with the priority value of primary monitors the VM sources.

1. Select **Device** > **VM Information Sources**. This example shows you how to add VMware ESX(i) or vCenter Server.
2. Click **Add** and enter the following information:
   - A **Name** to identify the source that you want to monitor.
   - Select the **Type** to indicate whether the source is an **AWS VPC**, a **Google Compute Engine** instance, a **VMware ESX(i)** server, or a **VMware vCenter** Server.

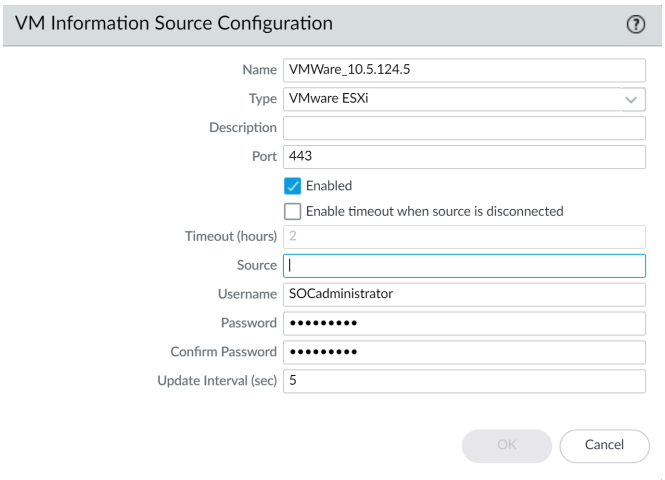
     The type chosen determines the fields displayed.
   - Enter the **Port** on which the source is listening.
   - To change the default value, select the check box to **Enable timeout when the source is disconnected** and specify the value. When the specified limit is reached or if the

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

host cannot be accessed or does not respond, the firewall will close the connection to the source.

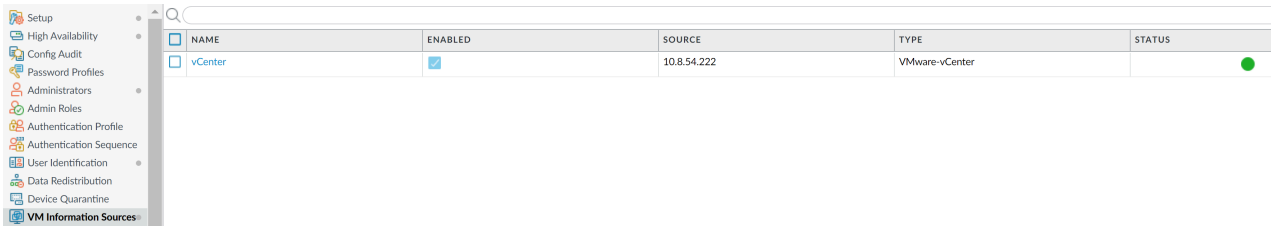
- Add the credentials (**Username** and **Password**) to authenticate to the server specified above.
- Define the **Source**—hostname or IP address.
- (Optional) Modify the **Update interval** to a value between 5-600 seconds. By default, the firewall polls every 5 seconds. The API calls are queued and retrieved within every 60 seconds, so updates may take up to 60 seconds plus the configured polling interval.



- Click **OK**, and **Commit** the changes.
- Verify that the connection **Status** displays as connected.

STEP 2 | Verify the connection status.

Verify that the connection **Status** displays as connected.



If the connection status is pending or disconnected, verify that the source is operational and that the firewall is able to access the source. If you use a port other than the MGT port for communicating with the monitored source, you must change the service route (**Device** > **Setup** > **Services**, click the **Service Route Configuration** link and modify the **Source Interface** for the **VM Monitor** service).

Attributes Monitored on Virtual Machines in Cloud Platforms

As you provision or remove virtual machines in the private or public cloud, you can use a Panorama plugin, a VM Monitoring script, or the VM Information Source on the next-gen firewall to monitor changes on virtual machines (VMs) deployed in the virtual environments.

https://t.me/learningnets

VM Information Sources—On a hardware or a VM-Series firewall you can monitor virtual machine instances and retrieves changes as you provision or modify the guests configured on the monitored sources—AWS, ESXi or vCenter Server, or AWS. For each firewall (and/or virtual system if your firewall has multiple virtual system capability), you can configure up to 10 sources.For information on how VM Information Sources and Dynamic Address Groups work synchronously and enable you to monitor changes in the virtual environment, refer to the VM-Series Deployment Guide .If your firewalls are configured in a high availability configuration:

- In an active/passive setup, only the active firewall monitors the VM information sources.
- In an active/active setup, only the primary firewall monitors the VM information sources.

Panorama Plugin—On a Panorama —hardware appliance or virtual appliance running version 8.1.3—you can install the plugin for Microsoft Azure and AWS. The plugin allows you to connect Panorama to your Azure public cloud subscriptions or AWS VPCs and retrieve the IP address-to-tag mapping for your virtual machines. Panorama then registers the VM information to the managed Palo Alto Networks® firewall(s) that you have configured for notification.

Use the following sections to review the options supported on each cloud vendor and the virtual machine attributes that you can monitor to create Dynamic Address Groups:

- VMware ESXi
- Amazon Web Services (AWS)
- Microsoft Azure
- Google

VMware ESXi

Each VM on a monitored ESXi or vCenter server must have VMware Tools installed and running. VMware Tools provide the capability to glean the IP address(es) and other values assigned to each VM.

When monitoring ESXi hosts that are part of the VM-Series NSX edition solution, use Dynamic Address Groups (instead of using VM Information Sources) to learn about changes in the virtual environment. For the VM-Series NSX edition solution, the NSX Manager provides Panorama with information on the NSX security group to which an IP address belongs. The information from the NSX Manager provides the full context for defining the match criteria in a Dynamic Address Group because it uses the service profile ID as a distinguishing attribute and allows you to properly enforce policy when you have overlapping IP addresses across different NSX security groups.

Up to 32 tags (from vCenter server and NSX Manager) can be registered to an IP address.

To collect the values assigned to the monitored VMs, use the VM Information Sources on the firewall to monitor the following predefined set of ESXi attributes:

Attributes Monitored on a VMware Source
UUID
Name

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Attributes Monitored on a VMware Source
Guest OS
VM State — the power state can be poweredOff, poweredOn, standBy, and unknown.
Annotation
Version
Network — Virtual Switch Name, Port Group Name, and VLAN ID
Container Name —vCenter Name, Data Center Object Name, Resource Pool Name, Cluster Name, Host, Host IP address.

Amazon Web Services (AWS)

As you provision or modify virtual machines in your AWS VPCs, you have two ways of monitoring these instances and retrieving the tags for use as match criteria in dynamic address groups.

- **VM Information Source**—On a next-gen firewall, you can monitor up to a total of 32 tags—14 pre-defined and 18 user-defined key-value pairs (tags). The following attributes (or tag names) are available as match criteria for dynamic address groups.
- **AWS Plugin on Panorama**—The Panorama plugin for AWS allows you to connect Panorama to your AWS VPCs and retrieve the IP address-to-tag mapping for your AWS virtual machines. Panorama then registers the VM information to the managed Palo Alto Networks® firewall(s) that you have configured for notification. With the plugin, Panorama can retrieve a total of 32 tags for each virtual machine, 11 predefined tags and up to 21 user-defined tags.

Attributes Monitored on the AWS-VPC	VM Information Source on the Firewall	AWS Plugin on Panorama
Architecture	Yes	No
Guest OS	Yes	No
AMI ID	Yes	Yes
IAM Instance Profile	No	Yes
Instance ID	Yes	No
Instance State	Yes	No
Instance Type	Yes	No
Key Name	Yes	Yes

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Attributes Monitored on the AWS-VPC	VM Information Source on the Firewall	AWS Plugin on Panorama
Owner ID	No	Yes
Placement—Tenancy	Yes	Yes
Placement—Group Name	Yes	Yes
Placement—Availability Zone	Yes	Yes
Private DNS Name	Yes	No
Public DNS Name	Yes	Yes
Subnet ID	Yes	Yes
Security Group ID	No	Yes
Security Group Name	No	Yes
VPC ID	Yes	Yes
Tag (key, value)	Yes; Up to a maximum of 18 user defined tags are supported. The user-defined tags are sorted alphabetically, and the first 18 tags are available for use on the firewalls.	Yes; Up to a maximum of 21 user defined tags are supported. The user-defined tags are sorted alphabetically, and the first 21 tags are available for use on Panorama and the firewalls.

Microsoft Azure

For VM Monitoring on Azure you need to retrieve the IP address-to-tag mapping for your Azure VMs and make it available as match criteria in dynamic address groups. The Panorama plugin for Microsoft Azure allows you to connect Panorama to your Azure public cloud subscriptions and retrieve the IP address-to-tag mapping for your Azure virtual machines. Panorama can retrieve a total of 26 tags for each virtual machine, 11 predefined tags and up to 15 user-defined tags and registers the VM information to the managed Palo Alto Networks® firewall(s) that you have configured for notification.

With the Panorama plugin for Azure, you can monitor the following set of virtual machine attributes within your Microsoft Azure deployment.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Attributes Monitored on Microsoft Azure	Azure Plugin on Panorama
VM Name	Yes
VM Size	No
Network Security Group Name	Yes
OS Type	Yes
OS Publisher	Yes
OS Offer	Yes
OS SKU	Yes
Subnet	Yes
VNet	Yes
Azure Region	Yes
Resource Group Name	Yes
Subscription ID	Yes
User Defined Tags	Yes Up to a maximum of 15 user defined tags are supported. The user-defined tags aresorted alphabetically, and the first 15 tags are available for use on Panorama and the firewalls.

Google

Using VM Information Sources on the next-gen firewall, you can monitor the following predefined set of Google Compute Engine (GCE) attributes.

High Availability is not supported on the firewalls.

Attributes Monitored on Google Compute Engine
Hostname of the VM
Machine type
Project ID

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Attributes Monitored on Google Compute Engine
Source (OS type)
Status
Subnetwork
VPC Network

Use Dynamic Address Groups in Policy

Dynamic address groups are used in policy. They allow you to create policy that automatically adapts to changes—adds, moves, or deletions of servers. It also enables the flexibility to apply different rules to the same server based on *tags* that define its role on the network, the operating system, or the different kinds of traffic it processes.

A dynamic address group uses tags as a filtering criteria to determine its members. The filter uses logical *and* and *or* operators. All IP addresses or address groups that match the filtering criteria become members of the dynamic address group. Tags can be defined statically on the firewall and/or registered (dynamically) to the firewall. The difference between static and dynamic tags is that static tags are part of the configuration on the firewall, and dynamic tags are part of the runtime configuration. This implies that a commit is not required to update dynamic tags; the tags must however be used by Dynamic Address Groups that are referenced in policy, and the policy must be committed on the firewall.

To dynamically register tags, you can use the XML API or the VM Monitoring agent on the firewall or on the User-ID agent. Each tag is a metadata element or attribute-value pair that is registered on the firewall or Panorama. For example, IP1 {tag1, tag2,.....tag32}, where the IP address and the associated tags are maintained as a list; each registered IP address can have up to 32 tags such as the operating system, the datacenter or the virtual switch to which it belongs. After receiving the API call, the firewall registers the IP address and associated tags, and automatically updates the membership information for the dynamic address group(s).

The maximum number of IP addresses that can be registered for each model is different. Use the following table for specifics on your model:

Model	Maximum number of dynamically registered IP addresses
M-Series and Panorama Virtual Appliances	500,000
PA-5200 Series, VM-7000 SMC-B Series	500,000
VM-500, VM-700	300,000
PA-3200 Series, VM-300	200,000

©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Model	Maximum number of dynamically registered IP addresses
PA-7000 Series	100,000
PA-850, VM-100	2,500
PA-820, PA-220, VM-50	1,000

An IP set, such as an IP range or subnet, is considered as a single registered IP address when counted towards the maximum number of registered IP addresses supported by each firewall model.

The following example shows how dynamic address groups can simplify network security enforcement. The example workflow shows how to:

- Enable the VM Monitoring agent on the firewall, to monitor the VMware ESX(i) host or vCenter Server and register VM IP addresses and the associated tags.
- Create dynamic address groups and define the tags to filter. In this example, two address groups are created. One that only filters for dynamic tags and another that filters for both static and dynamic tags to populate the members of the group.
- Validate that the members of the dynamic address group are populated on the firewall.
- Use dynamic address groups in policy. This example uses two different security policies:
  - A security policy for all Linux servers that are deployed as FTP servers; this rule matches on dynamically registered tags.
  - A security policy for all Linux servers that are deployed as web servers; this rule matches on a dynamic address group that uses static and dynamic tags.
- Validate that the members of the dynamic address groups are updated as new FTP or web servers are deployed. This ensure that the security rules are enforced on these new virtual machines too.

STEP 1 | Enable VM Source Monitoring.

See Enable VM Monitoring to Track Changes on the Virtual Network.

STEP 2 | Create dynamic address groups on the firewall.

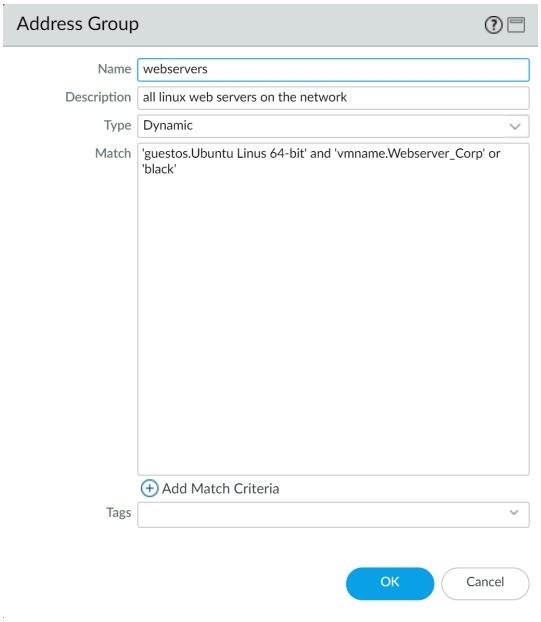
View the tutorial to see a big picture view of the feature.

1. Log in to the web interface of the firewall.
2. Select **Object** > **Address Groups**.
3. Click **Add** and enter a **Name** and a **Description** for the address group.
4. Select **Type** as **Dynamic**.
5. Define the match criteria. You can select dynamic and static tags as the match criteria to populate the members of the group. Click **Add Match Criteria**, and select the **And** or **Or**

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

operator and select the attributes that you would like to filter for or match against. and then click **OK**.

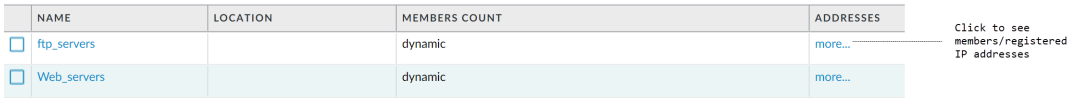


6. Click **Commit**.

STEP 3 | The match criteria for each dynamic address group in this example is as follows:

ftp_server: matches on the guest operating system “Linux 64-bit” and annotated as “ftp” ('guestos.Ubuntu Linux 64-bit' and 'annotation.ftp').

web-servers: matches on two criteria—the tag black or if the guest operating system is Linux 64-bit and the name of the server us Web_server_Corp. ('guestos.Ubuntu Linux 64-bit' and 'vmname.WebServer_Corp' or 'black')



STEP 4 | Use dynamic address groups in policy.

View the tutorial.

1. Select **Policies** > **Security**.
2. Click **Add** and enter a **Name** and a **Description** for the policy.
3. Add the **Source Zone** to specify the zone from which the traffic originates.
4. Add the **Destination Zone** at which the traffic is terminating.
5. For the **Destination Address**, select the Dynamic address group you just created.
6. Specify the action— **Allow** or **Deny**—for the traffic, and optionally attach the default security profiles to the rule.
7. Repeats Steps 1 through 6 to create another policy rule.
8. Click **Commit**.

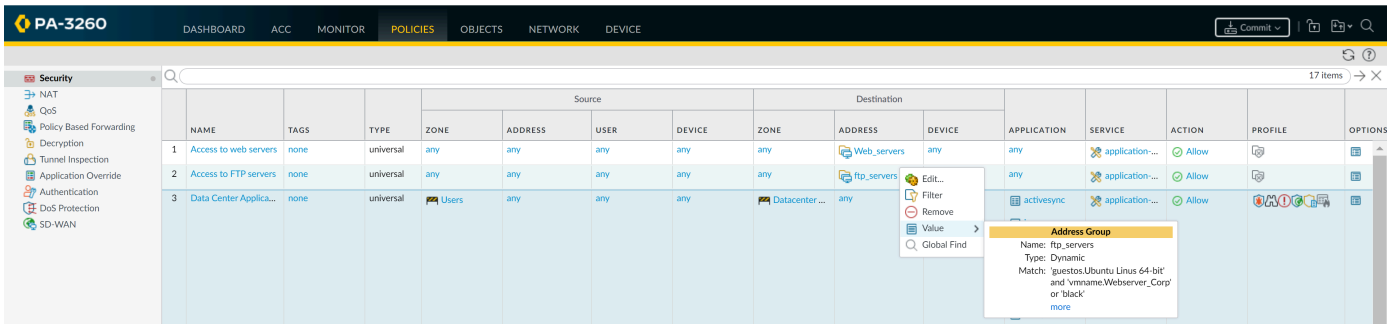
https://t.me/learningnets

STEP 5 | This example shows how to create two policies: one for all access to FTP servers and the other for access to web servers.

	NAME	TAGS	TYPE	Source ZONE	Source ADDRESS	Source USER	Source DEVICE	Destination ZONE	Destination ADDRESS	Destination DEVICE	APPLICATION	SERVICE	ACTION	PROFILE	OPTI
1	Access to web servers	none	universal	any	any	any	any	any	Web_servers	any	any	application-...	Allow		
2	Access to FTP servers	none	universal	any	any	any	any	any	ftp_servers	any	any	application-...	Allow		

STEP 6 | Validate that the members of the dynamic address group are populated on the firewall.

1. Select **Policies** > **Security**, and select the rule.
2. Select the drop-down arrow next to the address group link, and select **Value**. You can also verify that the match criteria is accurate.



3. Click the **more** link and verify that the list of registered IP addresses is displayed.

   Policy will be enforced for all IP addresses that belong to this address group, and are displayed here.

If you want to delete all registered IP addresses, use the CLI command ***debug object registered-ip clear all*** *and then reboot the firewall after clearing the tags.*

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

CLI Commands for Dynamic IP Addresses and Tags

The Command Line Interface on the firewall and Panorama give you a detailed view into the different sources from which tags and IP addresses are dynamically registered. It also allows you to audit registered and unregistered tags. The following examples illustrate the capabilities in the CLI.

Example	CLI Command
View all registered IP addresses that match the tag, `state.poweredOn` or that are not tagged as `vSwitch0`.	**show log iptag tag_name equal state.poweredOn** **show log iptag tag_name not-equal switch.vSwitch0**
View all dynamically registered IP addresses that were sourced by VM Information Source with name `vmware1` and tagged as `poweredOn`.	**show vm-monitor source source-name vmware1 tag state.poweredOn registered-ip all** registered IP Tags --- fe80::20c:29ff:fe69:2f76 "state.poweredOn" 10.1.22.100 "state.poweredOn" 2001:1890:12f2:11:20c:29ff:fe69:2f76" state.poweredOn" fe80::20c:29ff:fe69:2f80 "state.poweredOn" 192.168.1.102 "state.poweredOn" 10.1.22.105 "state.poweredOn" 2001:1890:12f2:11:2cf8:77a9:5435:c0d" state.poweredOn" fe80::2cf8:77a9:5435:c0d "state.poweredOn"
Clear all IP addresses and tags learned from a specific VM Monitoring source without disconnecting the source.	**debug vm-monitor clear source-name <*name*>**
Display IP addresses registered from all sources.	**show object registered-ip all**

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Example	CLI Command
Display the count for IP addresses registered from all sources.	**show object registered-ip all option count**
Clear IP addresses registered from all sources	**debug object registered-ip clear all**
Add or delete tags for a given IP address that was registered using the XML API.	**debug object registered-ip test [<register/unregister>] <ip/netmask><tag>**
View all tags registered from a specific information source.	**show vm-monitor source source-name vmware1 tag all** vlanId.4095 vswitch.vSwitch1 host-ip.10.1.5.22 portgroup.TOBEUSED hostname.panserver22 portgroup.VM Network 2 datacenter.ha-datacenter vlanId.0 state.poweredOn vswitch.vSwitch0 vmname.Ubuntu22-100 vmname.win2k8-22-105 resource-pool.Resources vswitch.vSwitch2 guestos.Ubuntu Linux 32-bit guestos.Microsoft Windows Server 2008 32-bit annotation. version.vmx-08 portgroup.VM Network vm-info-source.vmware1 uuid.564d362c-11cd-b27f-271f-c361604dfad7 uuid.564dd337-677a-eb8d-47db-293bd6692f76 Total: 22
View all tags registered from a specific data source, for example from the VM Monitoring Agent on the firewall, the XML API, Windows User-ID Agent or the CLI.	• To view tags registered from the CLI: **show log iptag datasource_type equal unknown**

https://t.me/learningnets

Example	CLI Command

	• To view tags registered from the XML API:
	`show log iptag datasource_type equal xml-api`
	• To view tags registered from VM Information sources:
	`show log iptag datasource_type equal vm-monitor`
	• To view tags registered from the Windows User-ID agent:
	`show log iptag datasource_type equal xml-api datasource_subtype equal user-id-agent`

View all tags that are registered for a specific IP address (across all sources).	`debug object registered-ip show tag-source ip ip_address tag all`

©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Enforce Policy on Endpoints and Users Behind an Upstream Device

If you have an upstream device, such as an explicit proxy server or load balance, deployed between the users on your network and the firewall, the firewall might see the upstream device IP address as the source IP address in HTTP/HTTPS traffic that the proxy forwards rather than the IP address of the client that requested the content. In many cases, the upstream device adds an X-Forwarded-For (XFF) header to HTTP requests that include the actual IPv4 or IPv6 address of the client that requested the content or from whom the request originated.

In such cases, you can configure the firewall to extract the IP address from the XFF field and map it to a user with User-ID or apply security policy based on the IP address.

- **Use X-Forwarded-For Header in User-ID**—This enables you enforce user-based policy to safely enable access to web-based applications for your users behind a proxy server. In addition, if User-ID is able to map the XFF IP address to a username, the firewall displays that username as the Source user in Traffic, Threat, WildFire Submissions, and URL Filtering logs for visibility into the web activity of users behind the proxy.
- **Use X-Forwarded-For Header in Security Policy**—This enables you to enforce security policy based on source IP address using the IP address in the XFF field of the HTTP header. Additionally, when policy is applied to traffic that includes an IP address in the XFF field, you can configure the Traffic, Threat, Data Filtering, and Wildfire Submission logs to assist in troubleshooting and remediation.

To ensure that attackers can't read and exploit the XFF values in web request packets that exit the firewall to retrieve content from an external server, you can also configure the firewall to strip the XFF values from outgoing packets. Using the XFF IP address for User-ID or in policy and stripping the XFF value are not mutually exclusive: if you configure both, the firewall zeroes out XFF values only after using them in policy enforcement and logging.

You cannot configure the firewall to use the IP address in the XFF field in User-ID and security policy at the same time.

- Use XFF Values for Policies and Logging Source Users
- Use XFF IP Address Values in Security Policy and Logging
- Use the IP Address in the XFF Header to Troubleshoot Events

Use XFF Values for Policy Based on Source Users

You can configure the firewall map the IP address in the XFF header to a username using User-ID so that you can have visibility into and user-based policy control over the web traffic of users behind a proxy server who cannot otherwise be identified. In order to map the IP addresses from the XFF headers to usernames, you must first Enable User-ID.

With this option enabled, the firewall uses the IP address in the XFF header for user mapping purposes only. The source IP address the firewall logs is still that of the proxy server, not that of the source user. When you see a log event attributed to a user that the firewall mapped using and IP address extracted from an XFF header, it can be difficult to track down the specific device associated with the event. To simplify debugging and troubleshooting of events attributed to users

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

behind the proxy server, you must also configure the firewall to populate the X-Forwarded-For column in the URL Filtering log with the IP address in the XFF header so that you can track down the specific user and device associated with an log event that is correlated with the URL Filtering log entry.

The XFF header your proxy server adds must contain the source IP address of the end user who originated the request. If the header contains multiple IP addresses, the firewall uses the first IP address only. If the header contains information other than an IP address, the firewall will not be able to perform user mapping.

Enabling the firewall to use the X-Forwarded-For headers to perform user mapping does not enable the firewall to use the client IP address in the XFF header as the source address in the logs; the logs still display the proxy server IP address as the source address. However, to simplify the debugging and troubleshooting process you can configure the firewall to Add XFF Values to URL Filtering Logs *to display the client IP address from the XFF header in the URL Filtering logs.*

STEP 1 | Enable the firewall to use XFF values in policies and in the source user fields of logs.

1. Select **Device** > **Setup** > **Content-ID** and edit the X-Forwarded-For Headers settings.
2. Select **Use X-Forwarded-For Header in User-ID**.

STEP 2 | Remove XFF values from outgoing web requests.

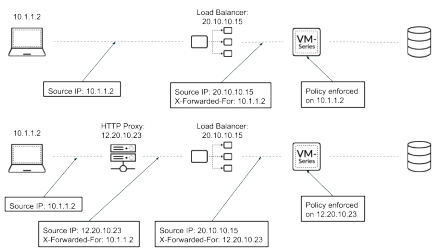
1. Select **Strip X-Forwarded-For Header**.
2. Click **OK** and **Commit**.

STEP 3 | Verify the firewall is populating the source user fields of logs.

1. Select a log type that has a source user field (for example, **Monitor** > **Logs** > **Traffic**).
2. Verify that the Source User column displays the usernames of users who access web applications.

Use XFF IP Address Values in Security Policy and Logging

You can configure the firewall to use the IP address in the X-Forwarded-For (XFF) field of the HTTP header to enforce security policy. If the packet passes through a single proxy server before reaching the firewall, the XFF field contains the IP address of the originating endpoint and the firewall can use that IP address to enforce security policy. However, if the packet passes through multiple upstream devices, the firewall uses the most-recently added IP address to enforce policy or use other features that rely on IP information.



- Use XFF Values in Policy
- Display XFF Values in Logs

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

- Display XFF Values in Reports

Use XFF Values in Policy

Complete the following procedure to use the client IP address in the XFF header when enforcing security policy.

In Microsoft Azure, by default, an application gateway inserts the original source IP address and port in the XFF header. To use XFF headers in policy on your firewall, you must configure the application gateway to omit the port from the XFF header. For more information, see Azure documentation.

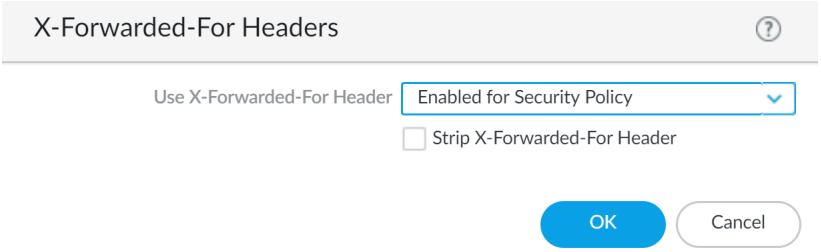
STEP 1 | Log in to your firewall.

STEP 2 | Select **Device** > **Setup** > **Content-ID** > **X-Forwarded-For Headers**.

STEP 3 | Click the edit icon.

STEP 4 | Select **Enabled for Security Policy** from the **Use X-Forwarded-For Header** drop-down.

You cannot enabled Use X-Forwarded-For Header for security policy and User-ID at the same time.



STEP 5 | (Optional) Select **Strip X-Forwarded-For Header**. Selecting this option removes the XFF header before the firewall forwards the request. This option does not disable the use of XFF headers; the firewall uses the XFF header for policy enforcement and logging.

STEP 6 | Click **OK**.

STEP 7 | **Commit** your changes.

Display XFF Values in Logs

In addition to XFF header usage in security policy, you can view the XFF IP address in various logs, reports, and the Application Command Center (ACC) to aide in monitoring and troubleshooting. You can add the X-Forwarded-For column in Traffic, Threat, Data Filtering, and Wildfire Submissions logs.

For non-URL Filtering logs, XFF IP logging is supported only when packet capture is no enabled.

To view the XFF IP address in your logs, complete the following steps.

https://t.me/learningnets

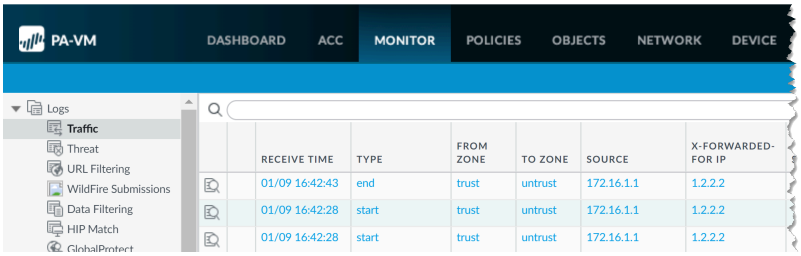
STEP 1 | Log in to your firewall.

STEP 2 | Select **Monitoring** > **Logs**.

STEP 3 | Select **Traffic**, **Threat**, **Data Filtering**, or **Wildfire Submissions**.

STEP 4 | Click the arrow to the right of any column header and select **Columns**.

STEP 5 | Select **X-Forwarded-For IP** to display the XFF IP in your log.



Display XFF Values in Reports

Predefined reports generate the firewall do not contain XFF values. To view XFF IP addresses in reports, the firewall includes built-in report templates that include XFF information.

STEP 1 | Log in to your firewall.

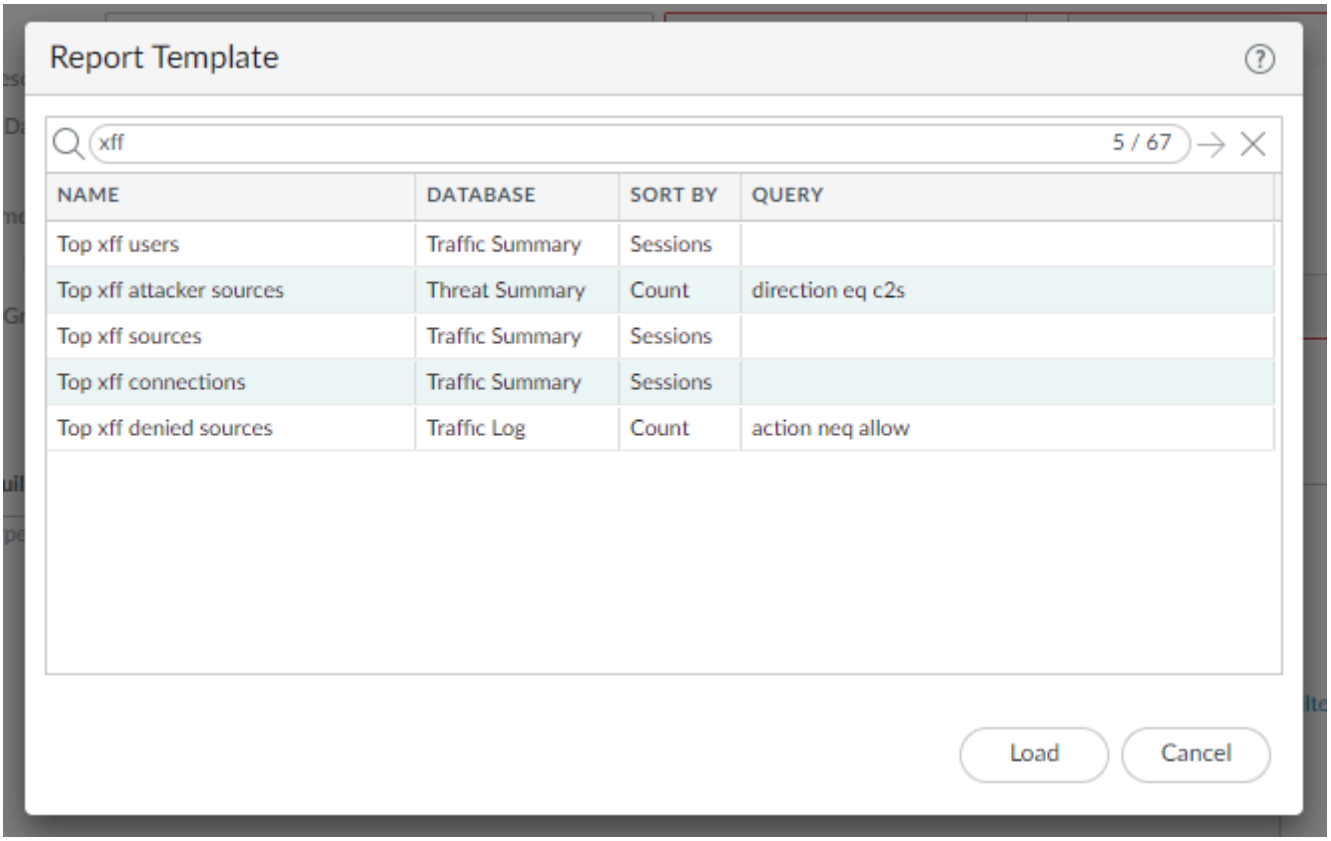
STEP 2 | Select **Monitor** > **Manage Custom Reports** > **Add**.

STEP 3 | Click **Load Template**.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 4 | Enter XFF into the search bar and click the search button to locate the built-in XFF report templates.



STEP 5 | Click **Load**.

STEP 6 | Configure your custom report **Time Frame**, **Sort By**, and **Group By** to display the XFF information in the manner best suited to your needs.

STEP 7 | (Optional) Click **Run Now** to generate your report on demand instead of, or in addition to, a **Scheduled Time**.

Use the IP Address in the XFF Header to Troubleshoot Events

By default, the firewall does not log the source address of a client behind a proxy server, even if you are using this address from the X-Forwarded-For (XFF) header for user mapping. Therefore, while you can identify the specific user associated with a log event, you will not be able to easily identify the source device that originated the log event. To simplify the debugging and troubleshooting of events for users behind a proxy server, you must enable the X-Forwarded-For option within HTTP Header Logging in the URL Filtering profile that you attach to security policy rules that allow access to web-based applications. With this option enabled, the firewall logs the IP address from the XFF header as the Source address for all traffic that matches the rule.

Enabling the firewall to use the XFF header as the Source address in URL Filtering logs does not enable user mapping of the source address. To populate the source user fields, see Use XFF Values for Policies and Logging Source Users.

https://t.me/learningnets

STEP 1 | Enable the X-Forwarded-For option within HTTP Header Logging in the URL Filtering profile.

1. Select **Objects** > **Security Profiles** > **URL Filtering** and select the URL Filtering profile you want to configure, or add a new one.

   You can't enable XFF logging in the default URL Filtering profile.

2. Select the **Settings** tab and select **X-Forwarded-For**.
3. Click **OK** to save the profile.

STEP 2 | Attach the URL Filtering profile to the security policy rule(s) that enable access to web applications.

1. Select **Policies** > **Security** and click the rule.
2. Select the **Actions** tab, set the **Profile Type** to **Profiles**, and select the **URL Filtering** profile you just configured for X-Forwarded-For HTTP Header Logging.
3. Click **OK** and **Commit**.

STEP 3 | Verify the firewall is logging XFF values.

1. Select **Monitor** > **Logs** > **URL Filtering**.
2. View the XFF values in one of the following ways:
   - To display the XFF value for a single URL Filtering log—Click the spyglass icon for the log to displays its details. The HTTP Headers section displays the X-Forwarded-For value.
   - To display the XFF values for all URL Filtering logs—Open the drop-down in any column header, select **Columns**, and select **X-Forwarded-For**. The page then displays an X-Forwarded-For column.

STEP 4 | Use the XFF field in the URL Filtering log to troubleshoot a log event in another log type.

Although only the URL Filtering logs display the IP address of the source user in the X-Forwarded-For column of the logs, if you notice an event associated with HTTP/HTTPS traffic but that you cannot identify the source IP address because it is that of the proxy server, you can use the X-Forwarded-For value in a correlated URL Filtering log to help you identify the source address associated with the log event. To do this:

1. Find an event you want investigate in a Traffic, Threat, or WildFire Submissions logs that is showing the IP address of the proxy server as the source address.
2. Click the spyglass icon for the log to display its details and look for an associated URL Filtering log at the bottom of the Detailed Log Viewer window.
3. Select the header row and then select **X-Forwarded-For** from the **Columns** drop-down to display this value. The IP address in this column of the X-Forwarded-For column represents the IP address of the source user behind the proxy server. Use this IP address to track down the device that triggered the event you are investigating.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Policy-Based Forwarding

Normally, the firewall uses the destination IP address in a packet to determine the outgoing interface. The firewall uses the routing table associated with the virtual router to which the interface is connected to perform the route lookup. Policy-Based Forwarding (PBF) allows you to override the routing table, and specify the outgoing or *egress* interface based on specific parameters such as source or destination IP address, or type of traffic.

- PBF
- Create a Policy-Based Forwarding Rule
- Use Case: PBF for Outbound Access with Dual ISPs

PBF

PBF rules allow traffic to take an alternative path from the next hop specified in the route table, and are typically used to specify an egress interface for security or performance reasons. Let's say your company has two links between the corporate office and the branch office: a cheaper internet link and a more expensive leased line. The leased line is a high-bandwidth, low-latency link. For enhanced security, you can use PBF to send applications that aren't encrypted traffic, such as FTP traffic, over the private leased line and all other traffic over the internet link. Or, for performance, you can choose to route business-critical applications over the leased line while sending all other traffic, such as web browsing, over the cheaper link.

- Egress Path and Symmetric Return
- Path Monitoring for PBF
- Service Versus Applications in PBF

Egress Path and Symmetric Return

Using PBF, you can direct traffic to a specific interface on the firewall, drop the traffic, or direct traffic to another virtual system (on systems enabled for multiple virtual systems).

In networks with asymmetric routes, such as in a dual ISP environment, connectivity issues occur when traffic arrives at one interface on the firewall and leaves from another interface. If the route is asymmetrical, where the forward (SYN packet) and return (SYN/ACK) paths are different, the firewall is unable to track the state of the entire session and this causes a connection failure. To ensure that the traffic uses a symmetrical path, which means that the traffic arrives at and leaves from the same interface on which the session was created, you can enable the *Symmetric Return* option.

With symmetric return, the virtual router overrides a routing lookup for return traffic and instead directs the flow back to the MAC address from which it received the SYN packet (or first packet). However, if the destination IP address is on the same subnet as the ingress/egress interface's IP address, a route lookup is performed and symmetric return is not enforced. This behavior prevents traffic from being silently discarded.



To determine the next hop for symmetric returns, the firewall uses an Address Resolution Protocol (ARP) table. The maximum number of entries that this ARP table supports is limited by the firewall model and the value is not user configurable. To determine the limit for your model, use the CLI command: ***show pbf return-mac all***.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Path Monitoring for PBF

Path monitoring allows you to verify connectivity to an IP address so that the firewall can direct traffic through an alternate route, when needed. The firewall uses ICMP pings as *heartbeats* to verify that the specified IP address is reachable.

A monitoring profile allows you to specify the threshold number of heartbeats to determine whether the IP address is reachable. When the monitored IP address is unreachable, you can either disable the PBF rule or specify a *fail-over* or *wait-recover* action. Disabling the PBF rule allows the virtual router to take over the routing decisions. When the fail-over or wait-recover action is taken, the monitoring profile continues to monitor whether the target IP address is reachable, and when it comes back up, the firewall reverts back to using the original route.

The following table lists the difference in behavior for a path monitoring failure on a new session versus an established session.

Behavior of a session on a monitoring failure	If the rule stays enabled when the monitored IP address is unreachable	If rule is disabled when the monitored IP address is unreachable
For an established session	**wait-recover**—Continue to use egress interface specified in the PBF rule	**wait-recover**—Continue to use egress interface specified in the PBF rule
	fail-over—Use path determined by routing table (no PBF)	**fail-over**—Use path determined by routing table (no PBF)
For a new session	**wait-recover**—Use path determined by routing table (no PBF)	**wait-recover**—Check the remaining PBF rules. If no match, use the routing table
	fail-over—Use path determined by routing table (no PBF)	**fail-over**—Check the remaining PBF rules. If no match, use the routing table

Service Versus Applications in PBF

PBF rules are applied either on the first packet (SYN) or the first response to the first packet (SYN/ACK). This means that a PBF rule may be applied before the firewall has enough information to determine the application. Therefore, application-specific rules are not recommended for use with PBF. Whenever possible, use a service object, which is the Layer 4 port (TCP or UDP) used by the protocol or application.

However, if you specify an application in a PBF rule, the firewall performs *App-ID caching*. When an application passes through the firewall for the first time, the firewall does not have enough information to identify the application and therefore cannot enforce the PBF rule. As more packets arrive, the firewall determines the application and creates an entry in the App-ID cache and retains this App-ID for the session.When a new session is created with the same destination IP address, destination port, and protocol ID, the firewall could identify the application as the

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

same from the initial session (based on the App-ID cache) and apply the PBF rule. Therefore, a session that is not an exact match and is not the same application, can be forwarded based on the PBF rule.

Further, applications have dependencies and the identity of the application can change as the firewall receives more packets. Because PBF makes a routing decision at the start of a session, the firewall cannot enforce a change in application identity. YouTube, for example, starts as web-browsing but changes to Flash, RTSP, or YouTube based on the different links and videos included on the page. However with PBF, because the firewall identifies the application as web-browsing at the start of the session, the change in application is not recognized thereafter.



You cannot use custom applications, application filters, or application groups in PBF rules.

Create a Policy-Based Forwarding Rule

Use a PBF rule to direct traffic to a specific egress interface on the firewall and override the default path for the traffic.

https://t.me/learningnets

STEP 1 | Create a Policy-Based Forwarding (PBF) rule.

When creating a PBF rule, you must specify a name for the rule, a source zone or interface, and an egress interface. All other components are either optional or have a default value.

You can specify the source and destination addresses using an IP address, an address object, or an FQDN.

1. Select **Policies** > **Policy Based Forwarding** and **Add** a PBF policy rule.
2. Give the rule a descriptive name (**General**).
3. Select **Source** and configure the following:
   1. Select the **Type** (**Zone** or **Interface**) to which you will apply the forwarding policy and specify the relevant zone or interface. If you want to enforce symmetric return, you must select a source interface.

      Only Layer 3 interfaces support PBF; loopback interfaces do not support PBF.

   2. (Optional) Specify the **Source Address** to which the PBF rule applies. For example, a specific IP address or subnet IP address from which you want to forward traffic to the interface or zone specified in this rule.

      *Click **Negate** to exclude one or more **Source Addresses** from the PBF rule. For example, if your PBF rule directs all traffic from the specified zone to the internet, **Negate** allows you to exclude internal IP addresses from the PBF rule.*

      The evaluation order is top down. A packet is matched against the first rule that meets the defined criteria; after a match is triggered, subsequent rules are not evaluated.
   3. (Optional) **Add** and select the **Source User** or groups of users to whom the policy applies.
4. Select **Destination/Application/Service** and configure the following:
   1. **Destination Address**—By default, the rule applies to **Any** IP address. Click **Negate** to exclude one or more destination IP addresses from the PBF rule.
   2. **Add** any **Application** and **Service** that you want to control using PBF.

      We do not recommend application-specific rules for use with PBF because PBF rules may be applied before the firewall has enough information to determine the application. Whenever possible, use a service object, which is the Layer 4 port (TCP or UDP) used by the protocol or application. For more details, see Service Versus Applications in PBF.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 2 | Specify how to forward packets that match the rule.

If you are configuring PBF in a multi-VSYS environment, you must create separate PBF rules for each virtual system (and create the appropriate Security policy rules to enable the traffic).

1. Select **Forwarding**.
2. Set the **Action** to take when matching a packet:
   - **Forward**—Directs the packet to the specified **Egress Interface**.
   - **Forward to VSYS** (On a firewall enabled for multiple virtual systems)—Select the virtual system to which to forward the packet.
   - **Discard**—Drops the packet.
   - **No PBF**—Excludes packets that match the criteria for source, destination, application, or service defined in the rule. Matching packets use the route table instead of PBF; the firewall uses the route table to exclude the matched traffic from the redirected port.
3. To trigger the specified **Action** at a daily, weekly, or non-recurring frequency, create and attach a **Schedule**.
4. For **Next Hop**, select one of the following:
   - **IP Address**—Enter an IP address or select an address object of type IP Netmask to which the firewall forwards matching packets. An IPv4 address object must have a /32 netmask and an IPv6 address object must have a /128 netmask.
   - **FQDN**—Enter an FQDN (or select or create an address object of type FQDN) to which the firewall forwards matching packets. The FQDN can resolve to an IPv4 address, an IPv6 address, or both. If the FQDN resolves to both IPv4 and IPv6 addresses, then the PBF rule has two next hops: one IPv4 address and one IPv6 address. You can use the same PBF rule for both IPv4 and IPv6 traffic. IPv4 traffic is forwarded to the IPv4 next hop; IPv6 traffic is forwarded to the IPv6 next hop.

     This FQDN must resolve to an IP address that belongs to the same subnet as the interface you configured for PBF; otherwise, the firewall rejects the resolution and the FQDN remains unresolved.

     The firewall uses only one IP address (from each IPv4 or IPv6 family type) from the DNS resolution of the FQDN. If the DNS resolution returns more than one address, the firewall uses the preferred IP address that matches the IP family type (IPv4 or IPv6) configured for the next hop. The preferred IP address is the first address the DNS server returns in its initial response. The firewall retains this address as preferred as long as the address appears in subsequent responses, regardless of order.
   - **None**—No next hop mean the destination IP address of the packet is used as the next hop. Forwarding fails if the destination IP address is not in the same subnet as the egress interface.
5. (Optional) Enable monitoring to verify connectivity to a target IP address or to the **Next Hop** IP address if no IP address is specified. Select **Monitor** and attach a monitoring

https://t.me/learningnets

Profile (default or custom) that specifies the action when the monitored address is unreachable.

- You can **Disable this rule if nexthop/monitor ip is unreachable**.
- Enter a target **IP Address** to monitor.

The **Egress Interface** can have both IPv4 and IPv6 addresses and the **Next Hop** FQDN can resolve to both IPv4 and IPv6 addresses. In this case:

1. If the egress interface has both IPv4 and IPv6 addresses and the next hop FQDN resolves to only one address family type, the firewall monitors the resolved IP address. If the FQDN resolves to both IPv4 and IPv6 addresses but the egress interface has only one address family type address, the firewall monitors the resolved next hop address that matches the address family of the egress interface.
2. If both the egress interface and next hop FQDN have both IPv4 and IPv6 addresses, the firewall monitors the IPv4 next hop address.
3. If the egress interface has one address family address and the next hop FQDN resolves to a different address family address, the firewall does not monitor anything.

6. (Required for asymmetric routing environments; otherwise, optional) **Enforce Symmetric Return** and **Add** one or more IP addresses in the **Next Hop Address List**. You can add up to 8 next-hop IP addresses; tunnel and PPoE interfaces are not available as a next-hop IP address.

   Enabling symmetric return ensures that return traffic (such asfrom the Trust zone on the LAN to the internet) is forwarded out through the same interface through which traffic ingresses from the internet.

STEP 3 | **Commit** your changes. The PBF rule is in effect.

	Source			Destination			Forwarding			Monitoring	
NAME	ZONE/INTERFACE	ADDRESS	USER	ADDRESS	SERVICE	ACTION	EGRESS I/F	NEXT HOP	ENFORCE SYMMETRIC RETURN	PROFILE	DISABLE IF UNREACHABL
pdf2	ethernet1/3	any	any	HQ-subnet	service-http	forward	ethernet1/1.100	192.168.100.2	false	none	false

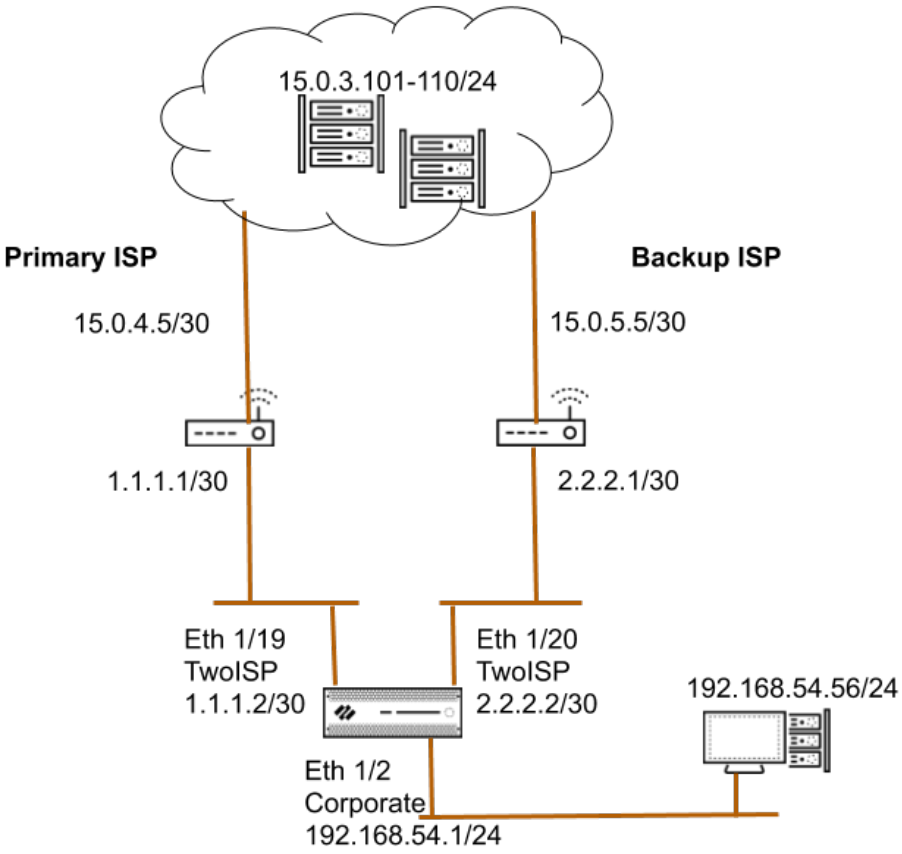
Use Case: PBF for Outbound Access with Dual ISPs

In this use case, the branch office has a dual ISP configuration and implements PBF for redundant internet access. The backup ISP is the default route for traffic from the client to the web servers. In order to enable redundant internet access without using an internetwork protocol such as BGP, we use PBF with destination interface-based source NAT and static routes, and configure the firewall as follows:

- Enable a PBF rule that routes traffic through the primary ISP, and attach a monitoring profile to the rule. The monitoring profile triggers the firewall to use the default route through the backup ISP when the primary ISP is unavailable.
- Define Source NAT rules for both the primary and backup ISP that instruct the firewall to use the source IP address associated with the egress interface for the corresponding ISP. This ensures that the outbound traffic has the correct source IP address.
- Add a static route to the backup ISP, so that when the primary ISP is unavailable, the default route comes into effect and the traffic is directed through the backup ISP.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets



 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 1 | Configure the ingress and the egress interfaces on the firewall.

Egress interfaces can be in the same zone.

1. Select **Network** > **Interfaces** and select the interface you want to configure.

   The interface configuration on the firewall used in this example is as follows:

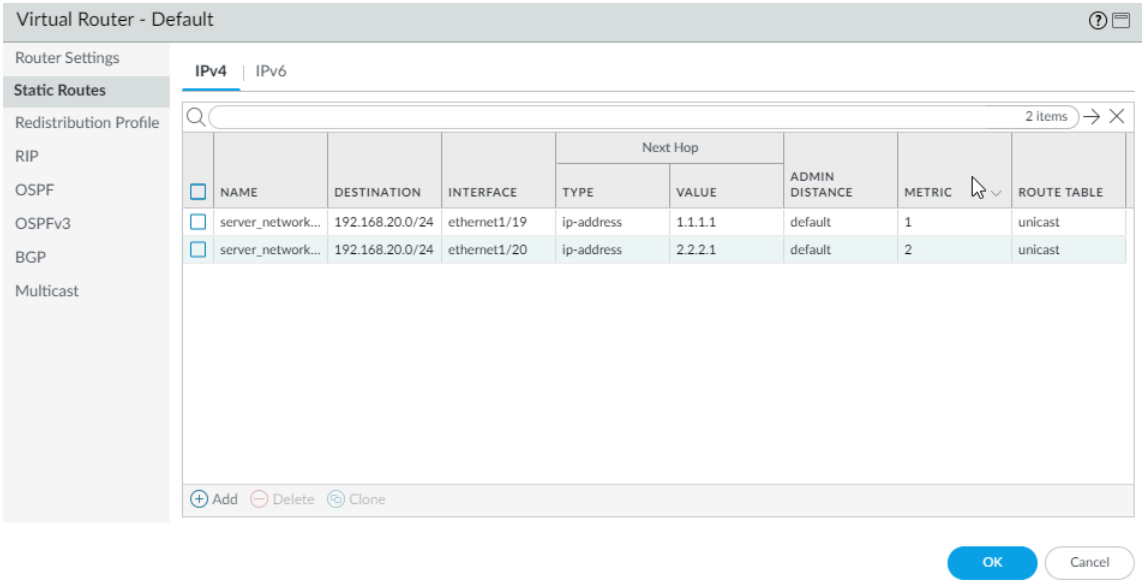
   - Ethernet 1/19 connected to the primary ISP:
     - Zone: TwoISP
     - IP Address: 1.1.1.2/30
     - Virtual Router: Default
   - Ethernet 1/20 connected to the backup ISP:
     - Zone: TwoISP
     - IP Address: 2.2.2.2/30
     - Virtual Router: Default
   - Ethernet 1/2 is the ingress interface, used by the network clients to connect to the internet:
     - Zone: Corporate
     - IP Address: 192.168. 54.1/24
     - Virtual Router: Default
2. To save the interface configuration, click **OK**.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 2 | On the virtual router, add a static route to the backup ISP.

1. Select **Network** > **Virtual Router** and select the **default** link to open the Virtual Router dialog.
2. Select **Static Routes** and click **Add**. Enter a **Name** for the route and specify the **Destination** IP address for which you are defining the static route. In this example, we use 0.0.0.0/0 for all traffic.
3. Select the **IP Address** radio button and set the **Next Hop** IP address for your router that connects to the backup internet gateway (you cannot use a domain name for the next hop). In this example, 2.2.2.1.
4. Specify a cost metric for the route.



5. Click **OK** twice to save the virtual router configuration.

 ©2022 Palo Alto Networks, Inc.

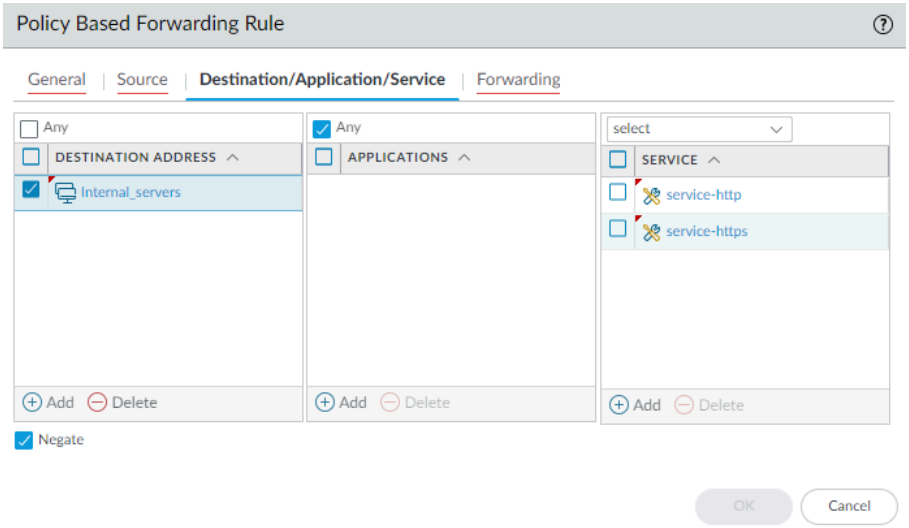
https://t.me/learningnets

STEP 3 | Create a PBF rule that directs traffic to the interface that is connected to the primary ISP.

Make sure to exclude traffic destined to internal servers/IP addresses from PBF. Define a negate rule so that traffic destined to internal IP addresses is not routed through the egress interface defined in the PBF rule.

1. Select **Policies** > **Policy Based Forwarding** and click **Add**.
2. Give the rule a descriptive **Name** in the **General** tab.
3. In the **Source** tab, set the **Source Zone**; in this example, the zone is Corporate.
4. In the **Destination/Application/Service** tab, set the following:
   1. In the Destination Address section, **Add** the IP addresses or address range for servers on the internal network or create an address object for your internal servers. Select **Negate** to exclude the IP addresses or address object listed above from using this rule.
   2. In the Service section, **Add** the **service-http** and **service-https** services to allow HTTP and HTTPS traffic to use the default ports. For all other traffic that is allowed by security policy, the default route will be used.

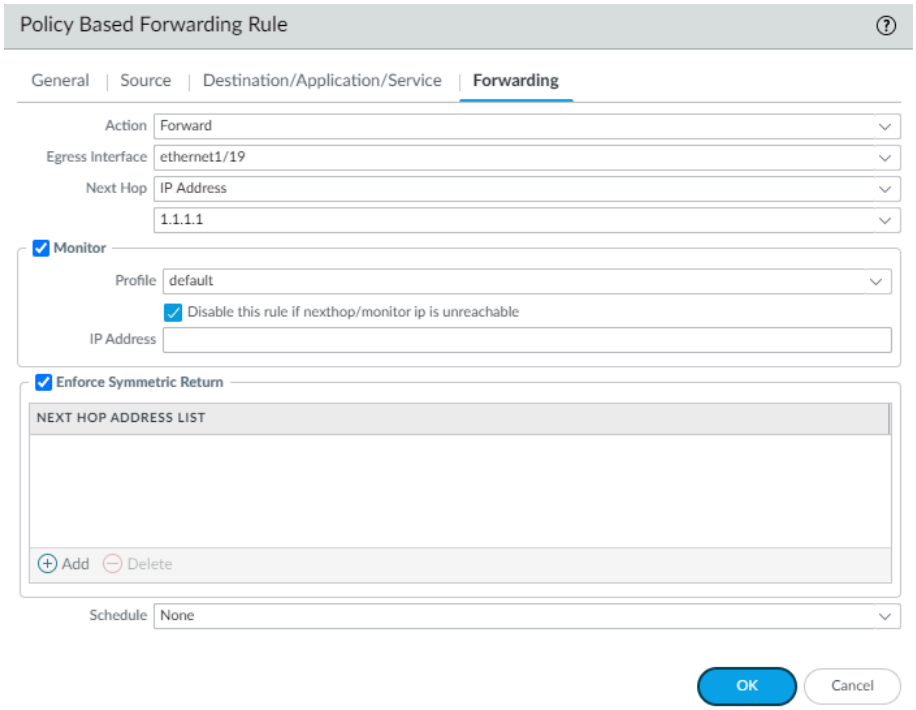
   *To forward all traffic using PBF, set the Service to **Any**.*



https://t.me/learningnets

STEP 4 | Specify where to forward traffic.

1. In the **Forwarding** tab, specify the interface to which you want to forward traffic and enable path monitoring.
2. To forward traffic, set the **Action** to **Forward**, and select the **Egress Interface** and specify the **Next Hop**. In this example, the egress interface is ethernet1/19, and the next hop IP address is 1.1.1.1 (you cannot use a FQDN for the next hop).



3. Enable **Monitor** and attach the default monitoring profile to trigger a failover to the backup ISP. In this example, we do not specify a target IP address to monitor. The firewall will monitor the next hop IP address; if this IP address is unreachable, the firewall will direct traffic to the default route specified on the virtual router.
4. (Required if you have asymmetric routes) Select **Enforce Symmetric Return** to ensure that return traffic from the Corporate zone to the internet is forwarded out on the same interface through which traffic ingressed from the internet.
5. NAT ensures that the traffic from the internet is returned to the correct interface/IP address on the firewall.
6. Click **OK** to save the changes.

		Source			Destination				Forwarding			Monitoring		
	NAME	ZONE/INTERFACE	ADDRESS	USER	ADDRESS	APPLICATION	SERVICE	ACTION	EGRESS I/F	NEXT HOP	ENFORCE SYMMETRIC RETURN	PROFILE	TARGET	DISABLE IF UNR
1	pbf_rule_source_zone	Corporate	192.168.10.2	any	any	any	service-http service-https	forward	ethernet1/19	1.1.1.1	true	default	none	true

https://t.me/learningnets

STEP 5 | Create NAT rules based on the egress interface and ISP. These rules ensure that the correct source IP address is used for outbound connections.

1. Select **Policies** > **NAT** and click **Add**.
2. In this example, the NAT rule we create for each ISP is as follows:

   NAT for Primary ISP

   In the **Original Packet** tab,

   Source Zone: Corporate

   Destination Zone: TwoISP

   In the **Translated Packet** tab, under Source Address Translation

   Translation Type: Dynamic IP and Port

   Address Type: Interface Address

   Interface: ethernet1/19

   IP Address: 1.1.1.2/30

   NAT for Backup ISP

   In the **Original Packet** tab,

   Source Zone: Corporate

   Destination Zone: TwoISP

   In the **Translated Packet** tab, under Source Address Translation

   Translation Type: Dynamic IP and Port

   Address Type: Interface Address

   Interface: ethernet1/20

   IP Address: 2.2.2.2/30

			Original Packet						Translated Packet	
	NAME	TAGS	SOURCE ZONE	DESTINATION ZONE	DESTINATION INTERFACE	SOURCE ADDRESS	DESTINATION ADDRESS	SERVICE	SOURCE TRANSLATION	DESTINATION TRANSL
1	NAT for Primary ISP	none	Corporate	TwoISP	any	any	any	any	dynamic-ip-and-port ethernet1/19 1.1.1.2/30	none
2	NAT for Backup ISP	none	Corporate	TwoISP	any	any	any	any	dynamic-ip-and-port ethernet1/20 2.2.2.2/30	none

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 6 | Create security policy to allow outbound access to the internet.

To safely enable applications, create a simple rule that allows access to the internet and attach the security profiles available on the firewall.

1. Select **Policies** > **Security** and click **Add**.
2. Give the rule a descriptive **Name** in the **General** tab.
3. In the **Source** tab, set the **Source Zone** to Corporate.
4. In the **Destination** tab, Set the **Destination Zone** to TwoISP.
5. In the **Service/ URL Category** tab, leave the default **application-default**.
6. In the **Actions** tab, complete these tasks:
   1. Set the **Action Setting** to **Allow**.
   2. Attach the default profiles for Antivirus, Anti-Spyware, Vulnerability Protection and URL Filtering, under **Profile Setting.**
7. Under **Options**, verify that logging is enabled at the end of a session. Only traffic that matches a security rule is logged.

	NAME	TAGS	TYPE	Source				Destination			APPLICATION	SERVICE	ACTION
				ZONE	ADDRESS	USER	DEVICE	ZONE	ADDRESS	DEVICE			
1	Copr2ISP	none	universal	Corporate	any	any	any	TwoISP	any	any	any	any	Allow

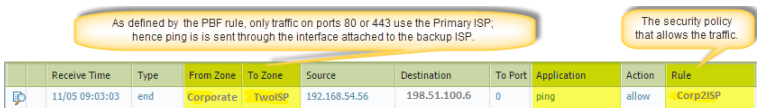
STEP 7 | Save the policies to the running configuration on the firewall.

Click **Commit**.

STEP 8 | Verify that the PBF rule is active and that the primary ISP is used for internet access.

1. Launch a web browser and access a web server. On the firewall, check the traffic log for web-browsing activity.
2. From a client on the network, use the ping utility to verify connectivity to a web server on the internet, and check the traffic log on the firewall.

```
C:\Users\pm-user1>ping 198.51.100.6
Pinging 198.51.100.6 with 32 bytes of data:
Reply from 198.51.100.6: bytes=32 time=34ms TTL=117
Reply from 198.51.100.6: bytes=32 time=13ms TTL=117
Reply from 198.51.100.6: bytes=32 time=25ms TTL=117
Reply from 198.51.100.6: bytes=32 time=3ms TTL=117
Ping statistics for 198.51.100.6:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milliseconds:
    Minimum = 3ms, Maximum = 34ms, Average = 18ms
```



3. To confirm that the PBF rule is active, use the following CLI command:

```
admin@PA-NGFW> show pbf rule all
Rule       ID    Rule State Action   Egress IF/VSYS  NextHop
```

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

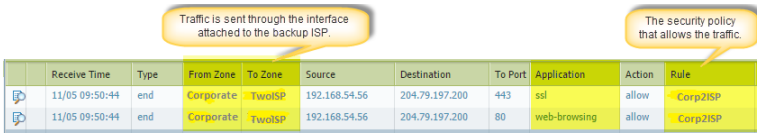
```
========== === ========== ====== ==============
Use ISP-Pr 1 Active     Forward ethernet1/1 1.1.1.1
```

STEP 9 | Verify that the failover to the backup ISP occurs and that the Source NAT is correctly applied.

1. Unplug the connection to the primary ISP.
2. Confirm that the PBF rule is inactive with the following CLI command:

```
admin@PA-NGFW> show pbf rule all
Rule       ID    Rule State Action   Egress IF/VSYS  NextHop
========== === ========== ====== ============== ===
Use ISP-Pr 1 Disabled Forward ethernet1/19     1.1.1.1
```

3. Access a web server, and check the traffic log to verify that traffic is being forwarded through the backup ISP.



4. View the session details to confirm that the NAT rule is working properly.

```
admin@PA-NGFW> show session all
--------------------------------------------------------------
ID Application    State   Type Flag Src[Sport]/Zone/Proto
 (translated IP[Port]) Vsys Dst[Dport]/Zone (translated
 IP[Port])
--------------------------------------------------------------
87212 ssl ACTIVE  FLOW  NS   192.168.54.56[53236]/Corporate/6
  (2.2.2.2[12896]) vsys1 204.79.197.200[443]/TwoISP
 (204.79.197.200[443])
```

5. Obtain the session identification number from the output and view the session details.

The PBF rule is not used and hence is not listed in the output.

```
admin@PA-NGFW> show session id 87212
Session           87212
c2s flow:
                source:      192.168.54.56 [Corporate]
                dst:         204.79.197.200
                proto:       6
                sport:       53236           dport:      443
                state:       ACTIVE          type:       FLOW
                src user:    unknown
                dst user:    unknown
s2c flow:
                source:      204.79.197.200 [TwoISP]
                dst:         2.2.2.2
                proto:       6
                sport:       443             dport:
 12896
```


https://t.me/learningnets

```
                state:       ACTIVE          type:       FLOW
                src user:    unknown
                dst user:    unknown
start time                    : Wed Nov5 11:16:10 2014
        timeout                          : 1800 sec
        time to live                     : 1757 sec
        total byte count(c2s)            : 1918
        total byte count(s2c)            : 4333
        layer7 packet count(c2s)         : 10
        layer7 packet count(s2c)         : 7
        vsys                             : vsys1
        application                      : ssl
        rule                             : Corp2ISP
        session to be logged at end      : True
        session in session ager          : True
        session synced from HA peer      : False
        address/port translation         : source
        nat-rule                         : NAT-Backup ISP(vsys1)
        layer7 processing                : enabled
        URL filtering enabled            : True
        URL category                     : search-engines
        session via syn-cookies          : False
        session terminated on host       : False
        session traverses tunnel         : False
        authentication portal session         : False
        ingress interface                : ethernet1/2
        egress interface                 : ethernet1/20
        session QoS rule                 : N/A (class 4)
```

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Test Policy Rules

Test the policy rules in your running configuration to ensure that your policies appropriately allow and deny traffic and access to applications and websites in compliance with your business needs and requirements. You can test and verify that your policy rules are allowing and denying the correct traffic by executing policy match tests for your firewalls directly from the web interface.

STEP 1 | Launch the Web Interface.

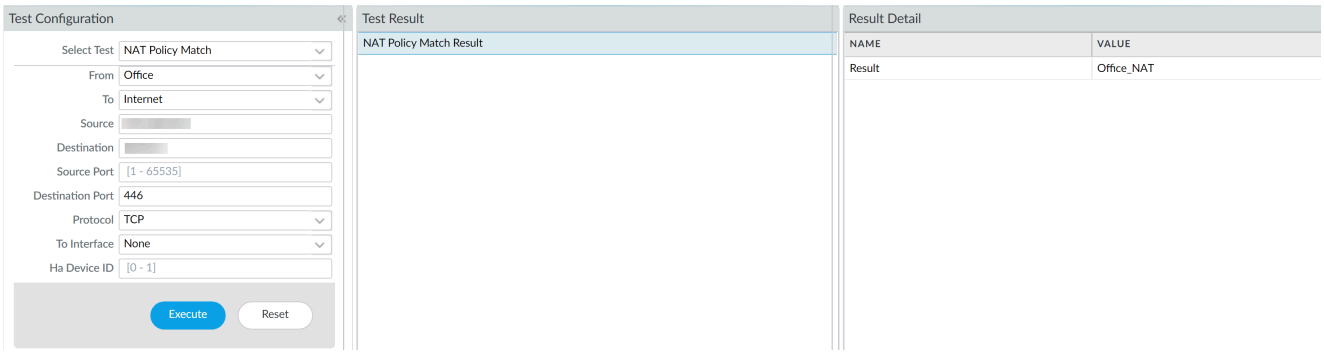
STEP 2 | Select **Device** > **Troubleshooting** to perform a policy match or connectivity test.

STEP 3 | Enter the required information to perform the policy match test. In this example, we run a NAT policy match test.

1. **Select Test**—Select **NAT Policy Match**.
2. **From**—Select the zone traffic is originating from.
3. **To**—Select the target zone of the traffic.
4. **Source**—Enter the IP address from which traffic originated.
5. **Destination**—Enter the IP address of the target device for the traffic.
6. **Destination Port**—Enter the port used for the traffic. This port varies depending on the IP protocol used in the following step.
7. **Protocol**—Enter the IP protocol used for the traffic.
8. If necessary, enter any additional information relevant for your NAT policy rule testing.

STEP 4 | **Execute** the NAT policy match test.

STEP 5 | Review the **NAT Policy Match Result** to see the policy rules that match the test criteria.



 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Virtual Systems

This topic describes virtual systems, their benefits, typical use cases, and how to configure them. It also provides links to other topics where virtual systems are documented as they function with other features.

> Virtual Systems Overview
> Communication Between Virtual Systems
> Shared Gateway
> Configure Virtual Systems
> Configure Inter-Virtual System Communication within the Firewall
> Configure a Shared Gateway
> Customize Service Routes for a Virtual System
> Virtual System Functionality with Other Features

https://t.me/learningnets

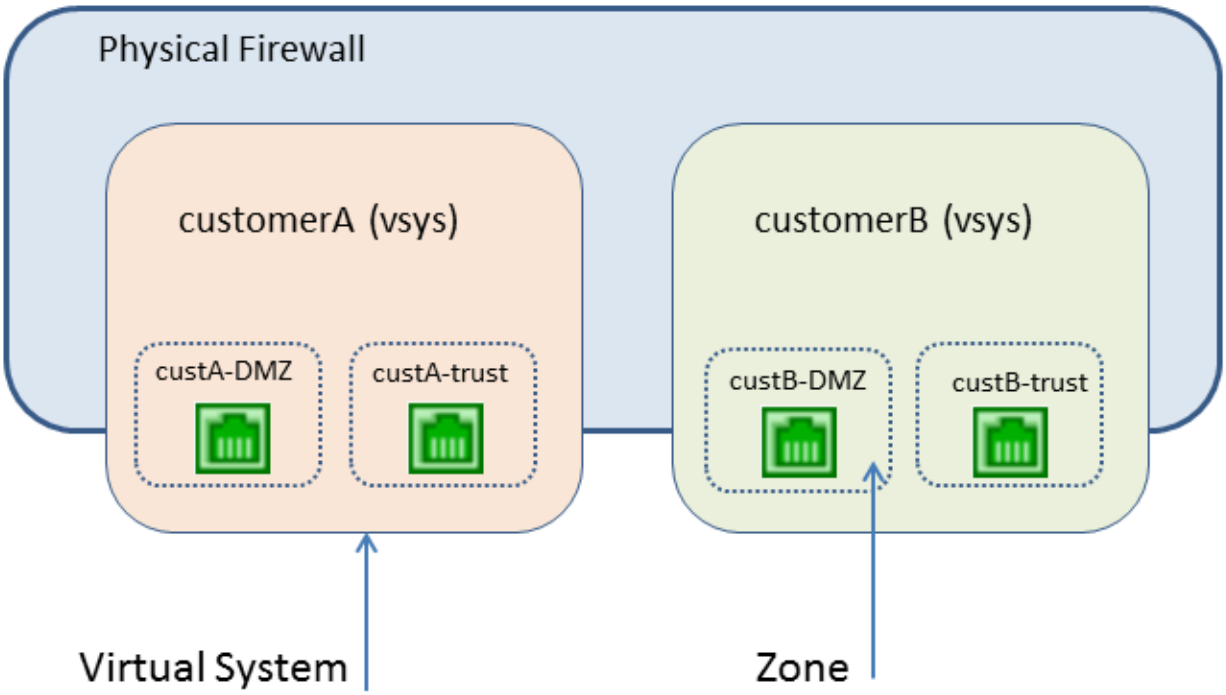
Virtual Systems Overview

Virtual systems are separate, logical firewall instances within a single physical Palo Alto Networks firewall. Rather than using multiple firewalls, managed service providers and enterprises can use a single pair of firewalls (for high availability) and enable virtual systems on them. Each virtual system (vsys) is an independent, separately-managed firewall with its traffic kept separate from the traffic of other virtual systems.

- Virtual System Components and Segmentation
- Benefits of Virtual Systems
- Use Cases for Virtual Systems
- Platform Support and Licensing for Virtual Systems
- Administrative Roles for Virtual Systems
- Shared Objects for Virtual Systems

Virtual System Components and Segmentation

A virtual system is an object that creates an administrative boundary, as shown in the following figure.



A virtual system consists of a set of physical and logical interfaces and subinterfaces (including VLANs and virtual wires), virtual routers, and security zones. You choose the deployment mode(s) (any combination of virtual wire, Layer 2, or Layer 3) of each virtual system. By using virtual systems, you can segment any of the following:

- Administrative access

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

- The management of all policies (Security, NAT, QoS, Policy-based Forwarding, Decryption, Application Override, Tunnel Inspection, Authentication, and DoS protection)
- All objects (such as address objects, application groups and filters, external dynamic lists, security profiles, decryption profiles, custom objects, etc.)
- User-ID
- Certificate management
- Server profiles
- Logging, reporting, and visibility functions

Virtual systems affect the security functions of the firewall, but virtual systems alone do not affect networking functions such as static and dynamic routing. You can segment routing for each virtual system by creating one or more virtual routers for each virtual system, as in the following use cases:

- If you have virtual systems for departments of one organization, and the network traffic for all of the departments is within a common network, you can create a single virtual router for multiple virtual systems.
- If you want routing segmentation and each virtual system's traffic must be isolated from other virtual systems, you can create one or more virtual routers for each virtual system.
- If you want to segment the user mappings so that not all mappings are shared across virtual systems, you can configure the User-ID sources on a virtual system that is not a User-ID hub. See Share User-ID Mappings Across Virtual Systems.

Benefits of Virtual Systems

Virtual systems provide the same basic functions as a physical firewall, along with additional benefits:

- **Segmented administration**—Different organizations (or customers or business units) can control (and monitor) a separate firewall instance, so that they have control over their own traffic without interfering with the traffic or policies of another firewall instance on the same physical firewall.
- **Scalability**—After the physical firewall is configured, adding or removing customers or business units can be done efficiently. An ISP, managed security service provider, or enterprise can provide different security services to each customer.
- **Reduced capital and operational expenses**—Virtual systems eliminate the need to have multiple physical firewalls at one location because virtual systems co-exist on one firewall. By not having to purchase multiple firewalls, an organization can save on the hardware expense, electric bills, and rack space, and can reduce maintenance and management expenses.
- **Ability to share IP-address-to-username mappings**—By assigning a virtual system as a User-ID hub, you can share the IP-address-to-username mappings across virtual systems to leverage the full User-ID capacity of the firewall and reduce operational complexity.

Use Cases for Virtual Systems

There are many ways to use virtual systems in a network. One common use case is for an ISP or a managed security service provider (MSSP) to deliver services to multiple customers with a single firewall. Customers can choose from a wide array of services that can be enabled or

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

disabled easily. The firewall's role-based administration allows the ISP or MSSP to control each customer's access to functionality (such as logging and reporting) while hiding or offering read-only capabilities for other functions.

Another common use case is within a large enterprise that requires different firewall instances because of different technical or confidentiality requirements among multiple departments. Like the above case, different groups can have different levels of access while IT manages the firewall itself. Services can be tracked and/or billed back to departments to thereby make separate financial accountability possible within an organization.

Platform Support and Licensing for Virtual Systems

Virtual systems are supported on PA-3200 Series, PA-5200 Series, and PA-7000 Series firewalls. Each firewall series supports a base number of virtual systems; the number varies by platform. A Virtual Systems license is required to support multiple virtual systems on PA-3200 Series firewalls, and to create more than the base number of virtual systems supported on a platform.

For license information, see Subscriptions. For the base and maximum number of virtual systems supported, see Compare Firewalls tool.

Multiple virtual systems are not supported on the PA-220, PA-800 Series, or VM-Series firewalls.

*The default is **vsys1**. You cannot delete vsys1 because it is relevant to the internal hierarchy on the firewall; vsys1 appears even on firewall models that don't support multiple virtual systems.*

You can limit the resource allocations for sessions, rules and VPN tunnels allowed for a virtual system, and thereby control firewall resources. Each resource setting displays the valid range of values, which varies per firewall model. The default setting is 0, which means the limit for the virtual system is the limit for the firewall model. However, the limit for a specific setting isn't replicated for each virtual system. For example, if a firewall has four virtual systems, each virtual system can't have the total number of Decryption Rules allowed per firewall. After the total number of Decryption Rules for all of the virtual systems reaches the firewall limit, you cannot add more.

Administrative Roles for Virtual Systems

A **Superuser** administrator can create virtual systems and add a **Device administrator**, **vsysadmin**, or **vsysreader**. A **Device administrator** can access all virtual systems, but cannot add administrators. When you create an Admin Role profile and select the role to be **Virtual System**, the role applies to specific virtual systems on the firewall. From the **Command Line** tab, the two types of virtual system administrative roles are:

- **vsysadmin**—Has access to specific virtual systems on the firewall to create and manage specific aspects of virtual systems. A vsysadmin doesn't have access to network interfaces, VLANs, virtual wires, virtual routers, IPSec tunnels, GRE tunnels, DHCP, DNS Proxy, QoS, LLDP, or network profiles. Persons with vsysadmin permission can commit configurations for only the virtual systems assigned to them.
- **vsysreader**—Has read-only access to specific virtual systems on the firewall and specific aspects of virtual systems. A vsysreader doesn't have access to network interfaces, VLANs, virtual wires, virtual routers, IPSec tunnels, GRE tunnels, DHCP, DNS Proxy, QoS, LLDP, or network profiles.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

A virtual system administrator can view logs of only the virtual systems assigned to that administrator. A **Superuser** or **Device administrator** can view all of the logs, select a virtual system to view, or configure a virtual system as a User-ID hub.

Shared Objects for Virtual Systems

If your administrator account extends to multiple virtual systems, you can choose to configure objects (such as an address object) and policies for a specific virtual system or as shared objects, which apply to all of the virtual systems on the firewall. If you try to create a shared object with the same name and type as an existing object in a virtual system, the virtual system object is used.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Communication Between Virtual Systems

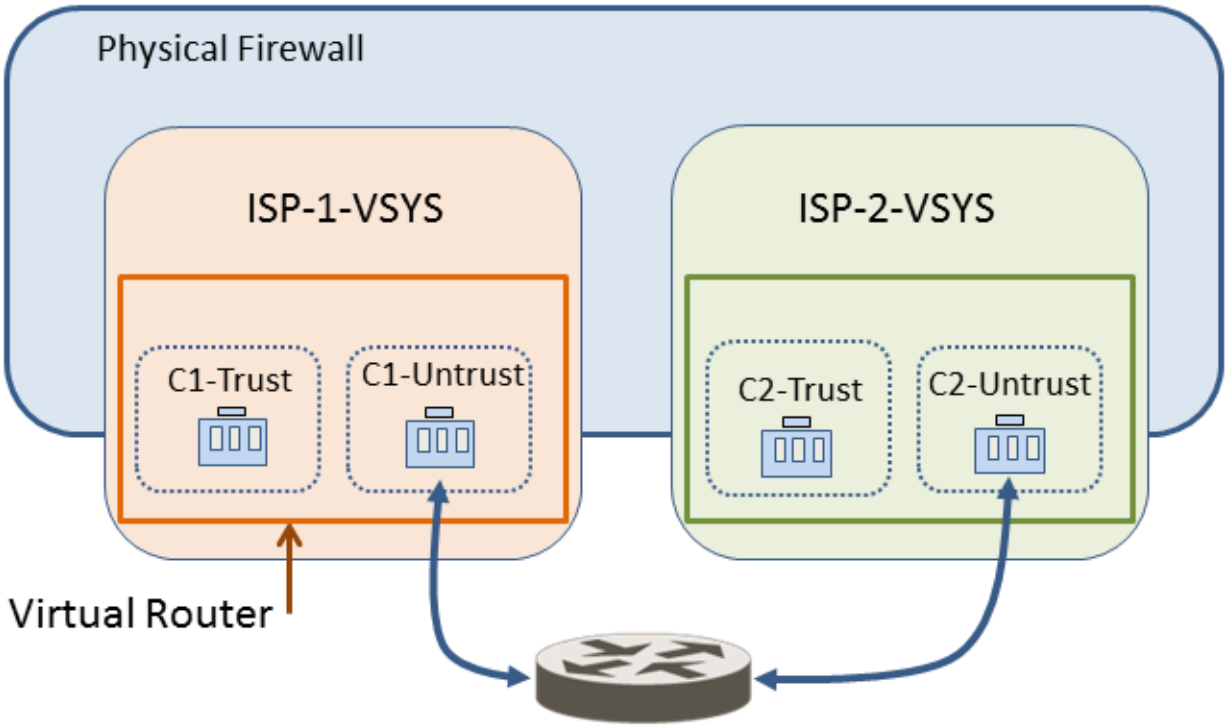
There are two typical scenarios where communication between virtual systems (inter-vsys traffic) is desirable. In a multi-tenancy environment, communication between virtual systems can occur by having traffic leave the firewall, go through the Internet, and re-enter the firewall. In a single organization environment, communication between virtual systems can remain within the firewall. This section discusses both scenarios.

- Inter-VSYS Traffic That Must Leave the Firewall
- Inter-VSYS Traffic That Remains Within the Firewall
- Inter-VSYS Communication Uses Two Sessions

Inter-VSYS Traffic That Must Leave the Firewall

An ISP that has multiple customers on a firewall (known as multi-tenancy) can use a virtual system for each customer, and thereby give each customer control over its virtual system configuration. The ISP grants **vsysadmin** permission to customers. Each customer's traffic and management are isolated from the others. Each virtual system must be configured with its own IP address and one or more virtual routers in order to manage traffic and its own connection to the Internet.

If the virtual systems need to communicate with each other, that traffic goes out the firewall to another Layer 3 routing device and back to the firewall, even though the virtual systems exist on the same physical firewall, as shown in the following figure.



 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Inter-VSYS Traffic That Remains Within the Firewall

Unlike the preceding multi-tenancy scenario, virtual systems on a firewall can be under the control of a single organization. The organization wants to both isolate traffic between virtual systems and allow communications between virtual systems. This common use case arises when the organization wants to provide departmental separation and still have the departments be able to communicate with each other or connect to the same network(s). In this scenario, the inter-vsys traffic remains within the firewall, as described in the following topics:

- External Zone
- External Zones and Security Policies For Traffic Within a Firewall

External Zone

The communication desired in the use case above is achieved by configuring security policies that point to or from an *external* zone. An external zone is a security object that is associated with a specific virtual system that it can reach; the zone is external to the virtual system. A virtual system can have only one external zone, regardless of how many security zones the virtual system has within it. External zones are required to allow traffic between zones in different virtual systems, without the traffic leaving the firewall.

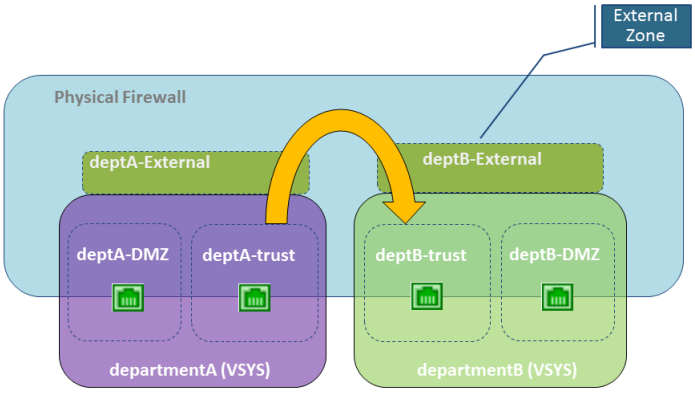
The virtual system administrator configures the security policies needed to allow traffic between two virtual systems. Unlike security zones, an external zone is not associated with an interface; it is associated with a virtual system. The security policy allows or denies traffic between the security (internal) zone and the external zone.

Because external zones do not have interfaces or IP addresses associated with them, some zone protection profiles are not supported on external zones.

Remember that each virtual system is a separate instance of a firewall, which means that each packet moving between virtual systems is inspected for security policy and App-ID evaluation.

External Zones and Security Policies For Traffic Within a Firewall

In the following example, an enterprise has two separate administrative groups: the departmentA and departmentB virtual systems. The following figure shows the external zone associated with each virtual system, and traffic flowing from one trust zone, out an external zone, into an external zone of another virtual system, and into its trust zone.



 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

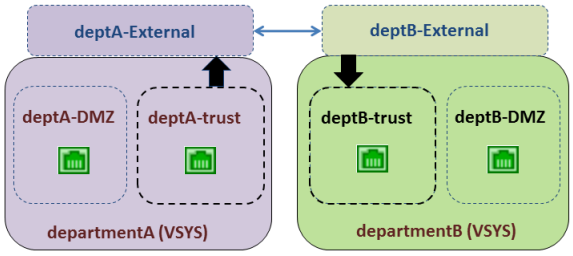
To create external zones, the firewall administrator must configure the virtual systems so that they are *visible* to each other. External zones do not have security policies between them because their virtual systems are visible to each other.

To communicate between virtual systems, the ingress and egress interfaces on the firewall are either assigned to a single virtual router or else they are connected using inter-virtual router static routes. The simpler of these two approaches is to assign all virtual systems that must communicate with each other to a single virtual router.

There might be a reason that the virtual systems need to have their own virtual router, for example, if the virtual systems use overlapping IP address ranges. Traffic can be routed between the virtual systems, but each virtual router must have static routes that point to the other virtual router(s) as the next hop.

Referring to the scenario in the figure above, we have an enterprise with two administrative groups: departmentA and departmentB. The departmentA group manages the local network and the DMZ resources. The departmentB group manages traffic in and out of the sales segment of the network. All traffic is on a local network, so a single virtual router is used. There are two external zones configured for communication between the two virtual systems. The departmentA virtual system has three zones used in security policies: deptA-DMZ, deptA-trust, and deptA-External. The departmentB virtual system also has three zones: deptB-DMZ, deptB-trust, and deptB-External. Both groups can control the traffic passing through their virtual systems.

In order to allow traffic from deptA-trust to deptB-trust, two security policies are required. In the following figure, the two vertical arrows indicate where the security policies (described below the figure) are controlling traffic.



- Security Policy 1: In the preceding figure, traffic is destined for the deptB-trust zone. Traffic leaves the deptA-trust zone and goes to the deptA-External zone. A security policy must allow traffic from the source zone (deptA-trust) to the destination zone (deptA-External). A virtual system allows any policy type to be used for this traffic, including NAT.

  No policy is needed between external zones because traffic sent to an external zone appears in and has automatic access to the other external zones that are visible to the original external zone.
- Security Policy 2: In the preceding figure, the traffic from deptB-External is still destined to the deptB-trust zone, and a security policy must be configured to allow it. The policy must allow traffic from the source zone (deptB-External) to the destination zone (deptB-trust).

The departmentB virtual system could be configured to block traffic from the departmentA virtual system, and vice versa. Like traffic from any other zone, traffic from external zones must be explicitly allowed by policy to reach other zones in a virtual system.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

In addition to external zones being required for inter-virtual system traffic that does not leave the firewall, external zones are also required if you configure a Shared Gateway, *in which case the traffic is intended to leave the firewall.*

Inter-VSYS Communication Uses Two Sessions

It is helpful to understand that communication between two virtual systems uses two sessions, unlike the one session used for a single virtual system. Let's compare the scenarios.

Scenario 1—Vsys1 has two zones: trust1 and untrust1. A host in the trust1 zone initiates traffic when it needs to communicate with a device in the untrust1 zone. The host sends traffic to the firewall, and the firewall creates a new session for source zone trust1 to destination zone untrust1. Only one session is needed for this traffic.

Scenario 2—A host from vsys1 needs to access a server on vsys2. A host in the trust1 zone initiates traffic to the firewall, and the firewall creates the first session: source zone trust1 to destination zone untrust1. Traffic is routed to vsys2, either internally or externally. Then the firewall creates a second session: source zone untrust2 to destination zone trust2. Two sessions are needed for this inter-vsys traffic.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Shared Gateway

This topic includes the following information about shared gateways:

- External Zones and Shared Gateway
- Networking Considerations for a Shared Gateway

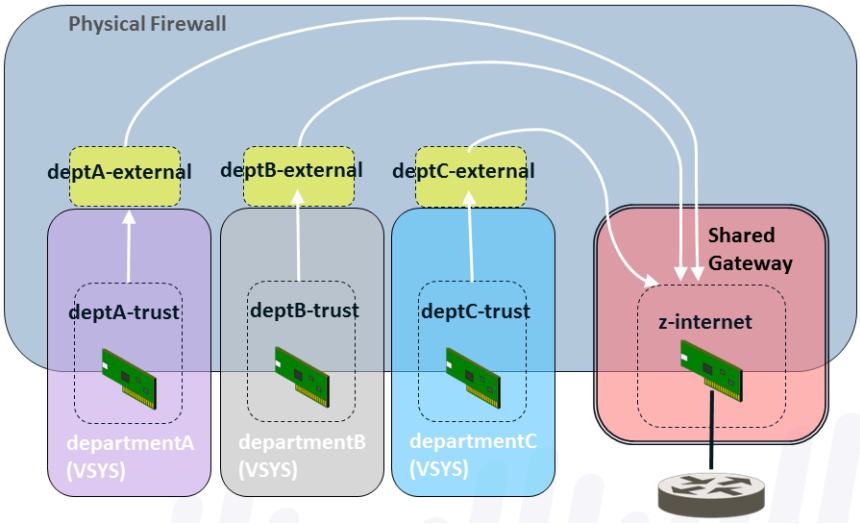
External Zones and Shared Gateway

A shared gateway is an interface that multiple virtual systems share in order to communicate over the Internet. Each virtual system requires an External Zone, which acts as an intermediary, for configuring security policies that allow or deny traffic from the virtual system's internal zone to the shared gateway.

The shared gateway uses a single virtual router to route traffic for all virtual systems. A shared gateway is used in cases when an interface does not need a full administrative boundary around it, or when multiple virtual systems must share a single Internet connection. This second case arises if an ISP provides an organization with only one IP address (interface), but multiple virtual systems need external communication.

Unlike the behavior between virtual systems, security policy and App-ID evaluations are not performed between a virtual system and a shared gateway. That is why using a shared gateway to access the Internet involves less overhead than creating another virtual system to do so.

In the following figure, three customers share a firewall, but there is only one interface accessible to the Internet. Creating another virtual system would add the overhead of App-ID and security policy evaluation for traffic being sent to the interface through the added virtual system. To avoid adding another virtual system, the solution is to configure a shared gateway, as shown in the following diagram.



The shared gateway has one globally-routable IP address used to communicate with the outside world. Interfaces in the virtual systems have IP addresses too, but they can be private, non-routable IP addresses.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

You will recall that an administrator must specify whether a virtual system is visible to other virtual systems. Unlike a virtual system, a shared gateway is always visible to all of the virtual systems on the firewall.

A shared gateway ID number appears as **sg<ID>** on the web interface. It is recommended that you name your shared gateway with a name that includes its ID number.

When you add objects such as zones or interfaces to a shared gateway, the shared gateway appears as an available virtual system in the vsys menu.

A shared gateway is a limited version of a virtual system; it supports NAT and policy-based forwarding (PBF), but does not support Security, DoS policies, QoS, Decryption, Application Override, or Authentication policies.

Networking Considerations for a Shared Gateway

Keep the following in mind while you are configuring a shared gateway.

- The virtual systems in a shared gateway scenario access the Internet through the shared gateway's physical interface, using a single IP address. If the IP addresses of the virtual systems are not globally routable, configure source NAT to translate those addresses to globally-routable IP addresses.
- A virtual router routes the traffic for all of the virtual systems through the shared gateway.
- The default route for the virtual systems should point to the shared gateway.
- Security policies must be configured for each virtual system to allow the traffic between the internal zone and external zone, which is visible to the shared gateway.
- A firewall administrator should control the virtual router, so that no member of a virtual system can affect the traffic of other virtual systems.
- Within a Palo Alto Networks firewall, a packet may hop from one virtual system to another virtual system or a shared gateway. A packet may not traverse more than two virtual systems or shared gateways. For example, a packet cannot go from vsys1 to vsys2 to vsys3, or similarly from vsys1 to vsys2 to shared gateway1. Both examples involve more than two virtual systems, which is not permitted.

To save configuration time and effort, consider the following advantages of a shared gateway:

- Rather than configure NAT for multiple virtual systems associated with a shared gateway, you can configure NAT for the shared gateway.
- Rather than configure policy-based routing (PBR) for multiple virtual systems associated with a shared gateway, you can configure PBR for the shared gateway.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Configure Virtual Systems

Creating a virtual system requires that you have the following:

- A **superuser** administrative role.
- An interface configured.
- A Virtual Systems license if you are creating more than the base number of virtual systems supported on the platform. See Platform Support and Licensing for Virtual Systems.

STEP 1 | Enable virtual systems.

1. Select **Device** > **Setup** > **Management** and edit the **General Settings**.
2. Select the **Multi Virtual System Capability** check box and click **OK**. This action triggers a commit if you approve it.

   Only after enabling virtual systems will the **Device** tab display the **Virtual Systems** and **Shared Gateways** options.

STEP 2 | Create a virtual system.

1. Select **Device** > **Virtual Systems**, click **Add** and enter a virtual system **ID**, which is appended to “vsys” (range is 1-255).

   *The default is **vsys1**. You cannot delete vsys1 because it is relevant to the internal hierarchy on the firewall; vsys1 appears even on firewall models that don’t support multiple virtual systems.*
2. Select **Allow forwarding of decrypted content** if you want to allow the firewall to forward decrypted content to an outside service. For example, you must enable this option for the firewall to be able to send decrypted content to WildFire for analysis.
3. Enter a descriptive **Name** for the virtual system. A maximum of 31 alphanumeric, space, and underscore characters is allowed.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 3 | Assign interfaces to the virtual system.

The virtual routers, virtual wires, or VLANs can either be configured already or you can configure them later, at which point you specify the virtual system associated with each.

1. On the **General** tab, select a **DNS Proxy** object if you want to apply DNS proxy rules to the interface.
2. In the **Interfaces** field, click **Add** to enter the interfaces or subinterfaces to assign to the virtual system. An interface can belong to only one virtual system.
3. Do any of the following, based on the deployment type(s) you need in the virtual system:
   - In the **VLANs** field, click **Add** to enter the VLAN(s) to assign to the vsys.
   - In the **Virtual Wires** field, click **Add** to enter the virtual wire(s) to assign to the vsys.
   - In the **Virtual Routers** field, click **Add** to enter the virtual router(s) to assign to the vsys.
4. In the **Visible Virtual System** field, check all virtual systems that should be made visible to the virtual system being configured. This is required for virtual systems that need to communicate with each other.

   In a multi-tenancy scenario where strict administrative boundaries are required, no virtual systems would be checked.
5. Click **OK**.

STEP 4 | (Optional) Limit the resource allocations for sessions, rules, and VPN tunnels allowed for the virtual system. The flexibility of being able to allocate limits per virtual system allows you to effectively control firewall resources.

1. On the **Resource** tab, optionally set limits for a virtual system. Each field displays the valid range of values, which varies per firewall model. The default setting is 0, which means the limit for the virtual system is the limit for the firewall model. However, the limit for a specific setting isn't replicated for each virtual system. For example, if a firewall has four virtual systems, each virtual system can't have the total number of Decryption

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Rules allowed per firewall. After the total number of Decryption Rules for all of the virtual systems reaches the firewall limit, you cannot add more.

- **Sessions Limit**

  If you use the show session meter CLI command, it displays the Maximum number of sessions allowed per dataplane, the Current number of sessions being used by the virtual system, and the Throttled number of sessions per virtual system. On a PA-5200 or PA-7000 Series firewall, the Current number of sessions being used can be greater than the Maximum configured for Sessions Limit because there are multiple dataplanes per virtual system. The Sessions Limit you configure on a PA-5200 Series or PA-7000 Series firewall is per dataplane, and will result in a higher maximum per virtual system.

- **Security Rules**
- **NAT Rules**
- **Decryption Rules**
- **QoS Rules**
- **Application Override Rules**
- **Policy Based Forwarding Rules**
- **Authentication Rules**
- **DoS Protection Rules**
- **Site to Site VPN Tunnels**
- **Concurrent SSL VPN Tunnels**

2. Click **OK**.

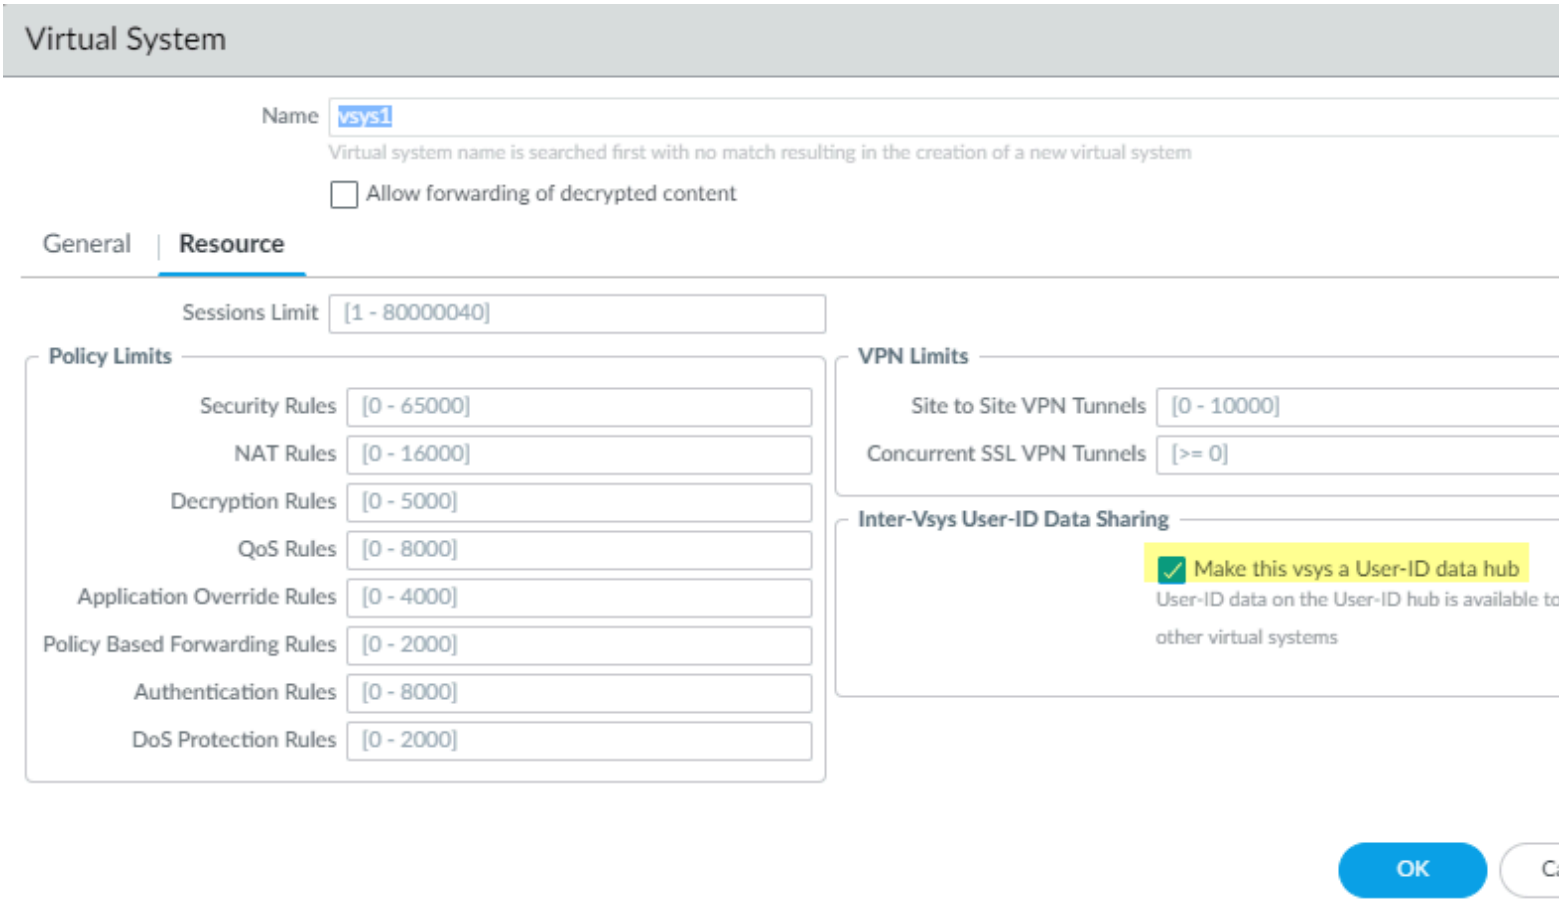
 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 5 | (Optional) Configure a virtual system as a User-ID hub to Share User-ID Mappings Across Virtual Systems.

IP-address-and-port-to-username mapping information from Terminal Server agents and group mapping data is not shared between the virtual system hub and the connected virtual systems.

1. For any existing virtual systems, transfer the configuration for the User-ID sources you want to share (such as monitored servers and User-ID agents) to the virtual system you will use as a hub.
2. On the **Resource** tab, select **Make this vsys a User-ID data hub**.

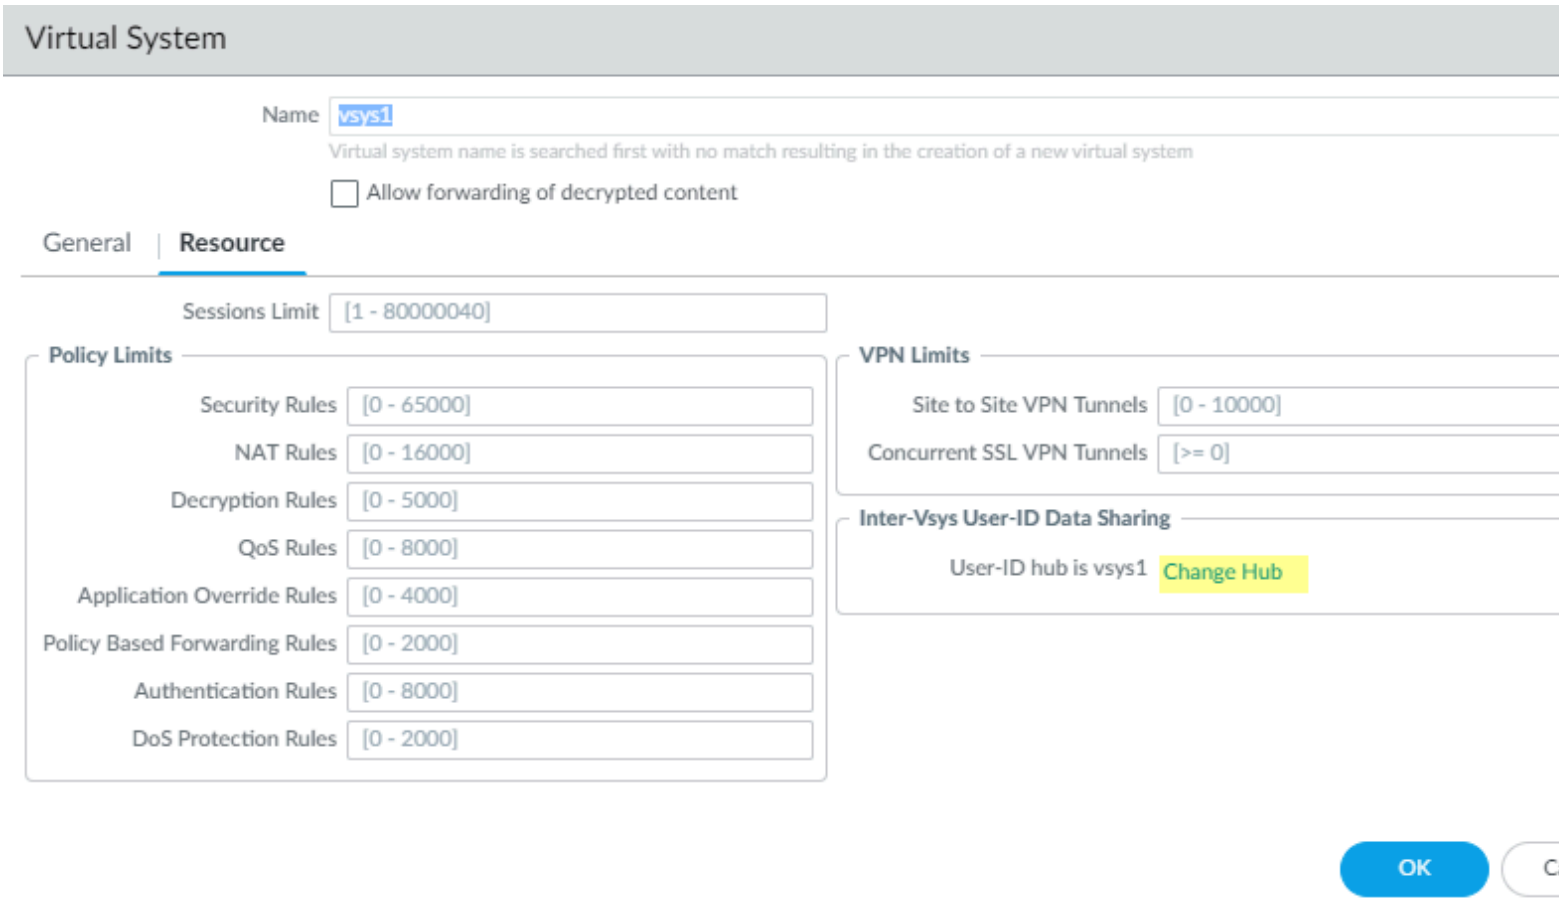


3. Click **Yes** to confirm, then click **OK**.

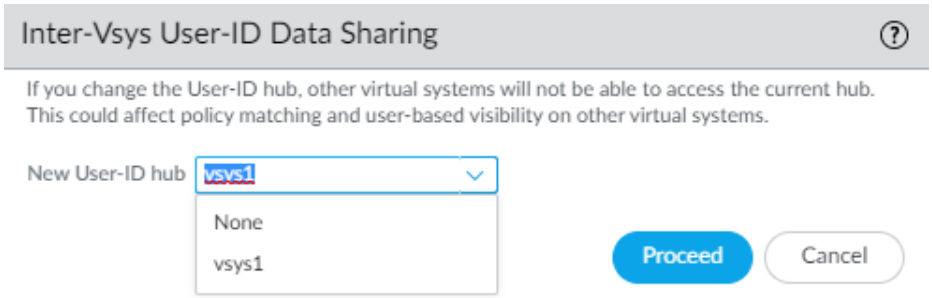
   If you want to change the User-ID hub to a different virtual system or disable it, select the virtual system currently configured as a User-ID hub, then select **Resource** > **Change Hub**.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets



Select the **New User-ID hub** from the list, or select **none** to disable the User-ID hub and stop sharing mappings across virtual systems.



Click **Proceed** to confirm and commit your changes.

STEP 6 | Commit the configuration.

Click **Commit**. The virtual system is now an object accessible from the **Objects** tab.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 7 | Create at least one virtual router for the virtual system in order to make the virtual system capable of networking functions, such as static and dynamic routing.

Alternatively, your virtual system might use a VLAN or a virtual wire, depending on your deployment.

1. Select **Network** > **Virtual Routers** and **Add** a virtual router by **Name**.
2. For **Interfaces**, click **Add** and select the interfaces that belong to the virtual router.
3. Click **OK**.

STEP 8 | Configure a security zone for each interface in the virtual system.

For at least one interface, create a Layer 3 security zone. See Configure Interfaces and Zones.

STEP 9 | Configure the security policy rules that allow or deny traffic to and from the zones in the virtual system.

See Create a Security Policy Rule.

STEP 10 | Commit the configuration.

Click **Commit**.



After creating a virtual system, you can use the CLI to commit a configuration for only a specific virtual system:

```
commit partial vsys <vsys-id>
```

STEP 11 | (Optional) View the security policies configured for a virtual system.

Open an SSH session to use the CLI. To view the security policies for a virtual system, in operational mode, use the following commands:

```
set system setting target-vsys <vsys-id>
show running security-policy
```

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Configure Inter-Virtual System Communication within the Firewall

Perform this task if you have a use case, perhaps within a single enterprise, where you want the virtual systems to be able to communicate with each other within the firewall. Such a scenario is described in Inter-VSYS Traffic That Remains Within the Firewall. This task presumes:

- You completed the task, Configure Virtual Systems.
- When configuring the virtual systems, in the **Visible Virtual System** field, you checked the boxes of all virtual systems that must communicate with each other to be visible to each other.

STEP 1 | Configure an external zone for each virtual system.

1. Select **Network** > **Zones** and **Add** a new zone by **Name**.
2. For **Location**, select the virtual system for which you are creating an external zone.
3. For **Type**, select **External**.
4. For **Virtual Systems**, click **Add** and enter the virtual system that the external zone can reach.
5. (Optional) Select a **Zone Protection Profile** (or configure one later) that provides flood, reconnaissance, or packet-based attack protection.
6. (Optional) In **Log Setting**, select a log forwarding profile for forwarding zone protection logs to an external system.
7. (Optional) Select **Enable User Identification** to enable User-ID for the external zone.
8. Click **OK**.

STEP 2 | Configure the Security policy rules to allow or deny traffic from the internal zones to the external zone of the virtual system, and vice versa.

- See Create a Security Policy Rule.
- See Inter-VSYS Traffic That Remains Within the Firewall.

STEP 3 | Commit your changes.

Click **Commit**.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Configure a Shared Gateway

Perform this task if you need multiple virtual systems to share an interface (a Shared Gateway) to the Internet. This task presumes:

- You configured an interface with a globally-routable IP address, which will be the shared gateway.
- You completed the prior task, Configure Virtual Systems. For the interface, you chose the external-facing interface with the globally-routable IP address.
- When configuring the virtual systems, in the **Visible Virtual System** field, you checked the boxes of all virtual systems that must communicate to be visible to each other.

STEP 1 | Configure a Shared Gateway.

1. Select **Device** > **Shared Gateway**, click **Add** and enter an **ID**.
2. Enter a helpful **Name**, preferably including the **ID** of the gateway.
3. In the **DNS Proxy** field, select a DNS proxy object if you want to apply DNS proxy rules to the interface.
4. **Add** an **Interface** that connects to the outside world.
5. Click **OK**.

STEP 2 | Configure the zone for the shared gateway.

*When adding objects such as zones or interfaces to a shared gateway, the shared gateway itself will be listed as an available vsys in the **VSYS** menu.*

1. Select **Network** > **Zones** and **Add** a new zone by **Name**.
2. For **Location**, select the shared gateway for which you are creating a zone.
3. For **Type**, select **Layer3**.
4. (Optional) Select a **Zone Protection Profile** (or configure one later) that provides flood, reconnaissance, or packet-based attack protection.
5. (Optional) In **Log Setting**, select a log forwarding profile for forwarding zone protection logs to an external system.
6. (Optional) Select **Enable User Identification** to enable User-ID for the shared gateway.
7. Click **OK**.

STEP 3 | Commit your changes.

Click **Commit**.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Customize Service Routes for a Virtual System

When a firewall is enabled for multiple virtual systems, the virtual systems inherit the global service and service route settings. For example, the firewall can use a shared email server to originate email alerts to all virtual systems. In some scenarios, you'd want to create different service routes for each virtual system.

One use case for configuring service routes at the virtual system level is if you are an ISP who needs to support multiple individual tenants on a single Palo Alto Networks firewall. Each tenant requires custom service routes to access service such as DNS, Kerberos, LDAP, NetFlow, RADIUS, TACACS+, Multi-Factor Authentication, email, SNMP trap, syslog, HTTP, User-ID Agent, VM Monitor, and Panorama (deployment of content and software updates). Another use case is an IT organization that wants to provide full autonomy to groups that set servers for services. Each group can have a virtual system and define its own service routes.

You can select a virtual router for a service route in a virtual system; you cannot select the egress interface. After you select the virtual router and the firewall sends the packet from the virtual router, the firewall selects the egress interface based on the destination IP address. Therefore, if a virtual system has multiple virtual routers, packets to all of the servers for a service must egress out of only one virtual router. A packet with an interface source address may egress a different interface, but the return traffic would be on the interface that has the source IP address, creating asymmetric traffic.

- Customize Service Routes to Services for Virtual Systems
- Configure a PA-7000 Series Firewall for Logging Per Virtual System
- Configure Administrative Access Per Virtual System or Firewall

Customize Service Routes to Services for Virtual Systems

When you enable Multi Virtual System Capability, any virtual system that does not have specific service routes configured inherits the global service and service route settings for the firewall. You can instead configure a virtual system to use a different service route, as described in the following workflow.

A firewall with multiple virtual systems must have interfaces and subinterfaces with non-overlapping IP addresses. A per-virtual system service route for SNMP traps or for Kerberos is for IPv4 only.

The service route for a service strictly follows how you configured the server profile for the service:

- If you define a server profile (**Device** > **Server Profiles**) for the Shared location, the firewall uses the global service route for that service.
- If you define a server profile for a specific virtual system, the firewall uses the virtual system-specific service route for that service.
- If you define a server profile for a specific virtual system but the virtual system-specific service route for that service is not configured, the firewall uses the global service route for that service.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

The firewall supports syslog forwarding on a virtual system basis. When multiple virtual systems on a firewall are connecting to a syslog server using SSL transport, the firewall can generate only one certificate for secure communication. The firewall does not support each virtual system having its own certificate.

STEP 1 | Customize service routes for a virtual system.

1. Select **Device** > **Setup** > **Services** > **Virtual Systems**, and select the virtual system you want to configure.
2. Click the **Service Route Configuration** link.
3. Select one:
   - **Inherit Global Service Route Configuration**—Causes the virtual system to inherit the global service route settings relevant to a virtual system. If you choose this option, skip the step to customize.
   - **Customize**—Allows you to specify a source address for each service.
4. If you chose **Customize**, select the **IPv4** or **IPv6** tab, depending on what type of addressing the server offering the service uses. You can specify both IPv4 and IPv6 addresses for a service. Click on a service. (Only services that are relevant to a virtual system are available.)

   *To easily use the same source address for multiple services, select the checkbox for the services, click **Set Selected Routes**, and continue.*

   - To limit the list for Source Address, select a **Source Interface**, then select a Source Address (from that interface) as the service route. Selecting **Any** Source Interface makes all IP addresses on all interfaces for the virtual system available in the Source Address list from which you select an address. You can select **Inherit Global Setting**.
   - **Source Address** will indicate **Inherited** if you selected **Inherit Global Setting** for the **Source Interface** or it will indicate the source address you selected. If you selected **Any** for **Source Interface**, select an IP address or enter an IP address (using the IPv4 or IPv6 format that matches the tab you chose) to specify the source address that will be used in packets sent to the external service.
   - If you modify an address object and the IP family type (IPv4/IPv6) changes, a **Commit** is required to update the service route family to use.
5. Click **OK**.
6. Repeat the prior steps to configure source addresses for other external services.
7. Click **OK**.

STEP 2 | Commit your changes.

Click **Commit** and **OK**.

If you are configuring per-virtual system service routes for logging services for a PA-7000 Series firewall, continue to the task Configure a PA-7000 Series Firewall for Logging Per Virtual System.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Configure a PA-7000 Series Firewall for Logging Per Virtual System

For Traffic, HIP Match, Threat, and WildFire log types, the PA-7000 Series firewall does not use service routes for SNMP Trap, Syslog, and email services. Instead, the PA-7000 Series firewall supports using a logging card.

Depending on your firewall configuration, you might have one of the following card types:

- **Log Processing Card (LPC)**—Supports virtual system-specific paths from LPC subinterfaces to an on-premise switch to the respective service on a server. For System and Config logs, the PA-7000 Series firewall uses global service routes, and not the LPC. If your firewall has an LPC installed, you need to configure a log card port.
- **Log Forwarding Card (LFC)**—Supports high-speed log forwarding of all dataplane logs to an external log collector (for example, Panorama and syslog servers). If your firewall has an LFC installed, you do not need to configure a log card port.



Log forwarding to an external server is not yet supported on LFC subinterfaces.

In other Palo Alto Networks models, the dataplane sends logging service route traffic to the management plane, which sends the traffic to logging servers. In a PA-7000 Series firewall, the LPC or LFC have only one interface, and dataplanes for multiple virtual systems send logging server traffic (types mentioned above) to the PA-7000 Series firewall logging card. The logging card is configured with multiple subinterfaces, over which the platform sends the logging service traffic out to a customer's switch, which can be connected to multiple logging servers.

Each subinterface can be configured with a subinterface name and a dotted subinterface number. The subinterface is assigned to a virtual system, which is configured for logging services. The other service routes on a PA-7000 Series firewall function similarly to service routes on other Palo Alto Networks platforms. For information about the LPC or LFC, see the PA-7000 Series Hardware Reference Guide.

- Configure a PA-7000 Series LPC for Logging per Virtual System
- Configure a PA-7000 Series LFC for Logging per Virtual System

Configure a PA-7000 Series LPC for Logging per Virtual System

If you have enabled multi-vsys capability on a PA-7000 Series firewall with a Log Processing Card (LPC) installed, you can configure logging for different virtual systems as described in the following workflow.

STEP 1 | Create a Log Card subinterface.

1. Select **Network** > **Interfaces** > **Ethernet** and select the interface to be the Log Card interface.
2. Enter the **Interface Name**.
3. For **Interface Type**, select **Log Card**.
4. Click **OK**.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 2 | Add a subinterface for each tenant on the LPCs physical interface.

1. Highlight the Ethernet interface that is a Log Card interface type and click **Add Subinterface**.
2. For **Interface Name**, after the period, enter the subinterface assigned to the tenant's virtual system.
3. For **Tag**, enter a VLAN tag value.

   Make the tag the same as the subinterface number for ease of use, but it could be a different number.
4. (Optional) Enter a **Comment**.
5. On the **Config** tab, in the **Assign Interface to Virtual System** field, select the virtual system to which the LPC subinterface is assigned. Alternatively, you can click **Virtual Systems** to add a new virtual system.
6. Click **OK**.

STEP 3 | Enter the addresses assigned to the subinterface, and configure the default gateway.

1. Select the **Log Card Forwarding** tab, and do one or both of the following:
   - For the IPv4 section, enter the **IP Address** and **Netmask** assigned to the subinterface. Enter the **Default Gateway** (the next hop where packets will be sent that have no known next hop address in the Routing Information Base [RIB]).
   - For the IPv6 section, enter the **IPv6 Address** assigned to the subinterface. Enter the **IPv6 Default Gateway**.
2. Click **OK**.

STEP 4 | Commit your changes.

Click **OK** and **Commit**.

STEP 5 | If you haven't already done so, configure the remaining service routes for the virtual system.

Customize Service Routes for a Virtual System.

Configure a PA-7000 Series LFC for Logging per Virtual System

If you have enabled multiple virtual system (multi-vsys) capability on a PA-7000 Series firewall with a Log Forwarding Card (LFC) installed, you can configure logging for different virtual systems. The LFC can then forward logs to a Panorama Log Collector or syslog server.

You can choose to configure only the physical interface. Because syslog forwarding via subinterfaces is not yet supported on LFCs, each virtual system uses the single untagged physical interface.

If you configure an LFC subinterface to forward logs externally, the interfaces will no longer work as expected.

To configure a separate subinterface for each virtual system, add subinterfaces to the physical interface and assign the necessary tag to segment the subinterface traffic.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

For a PA-7000 Series firewall managed by a Panorama management server, you cannot override or revert the LFC configuration locally on the firewall if the LFC configuration is pushed from Panorama. To override the LFC configuration pushed from Panorama, you must log in to the firewall CLI and delete the Panorama pushed configuration.

```
admin> configure
```

```
admin# delete deviceconfig log-fwd-card
```

```
admin# commit
```

Configure Administrative Access Per Virtual System or Firewall

If you have a superuser administrative account, you can create and configure granular permissions for a vsysadmin or device admin role.

STEP 1 | Create an Admin Role Profile that grants or disables permission to an Administrator to configure or read-only various areas of the web interface.

1. Select **Device** > **Admin Roles** and **Add** an **Admin Role Profile**.
2. Enter a **Name** and optional **Description** of the profile.
3. For **Role**, specify which level of control the profile affects:
   - **Device**—The profile allows the management of the global settings and any virtual systems.
   - **Virtual System**—The profile allows the management of only the virtual system(s) assigned to the administrator(s) who have this profile. (The administrator will be able to access **Device** > **Setup** > **Services** > **Virtual Systems**, but not the **Global** tab.)
4. On the **Web UI** tab for the Admin Role Profile, scroll down to **Device**, and leave the green check mark (Enable).
   - Under **Device**, enable **Setup**. Under **Setup**, enable the areas to which this profile will grant configuration permission to the administrator, as shown below. (The Read Only lock icon appears in the Enable/Disable rotation if Read Only is allowed for that setting.)
     - **Management**—Allows an admin with this profile to configure settings on the **Management** tab.
     - **Operations**—Allows an admin with this profile to configure settings on the **Operations** tab.
     - **Services**—Allows an admin with this profile to configure settings on the **Services** tab. An admin must have **Services** enabled in order to access the **Device** > **Setup**

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Services > **Virtual Systems** tab. If the **Role** was specified as **Virtual System** in the prior step, **Services** is the only setting that can be enabled under **Device** > **Setup**.

- **Content-ID**—Allows an admin with this profile to configure settings on the **Content-ID** tab.
- **WildFire**—Allows an admin with this profile to configure settings on the **WildFire** tab.
- **Session**—Allows an admin with this profile to configure settings on the **Session** tab.
- **HSM**—Allows an admin with this profile to configure settings on the **HSM** tab.

5. Click **OK**.
6. (Optional) Repeat the entire step to create another Admin Role profile with different permissions, as necessary.

STEP 2 | Apply the Admin role profile to an administrator.

1. Select **Device** > **Administrators**, click **Add** and enter the **Name** to add an Administrator.
2. (Optional) Select an **Authentication Profile**.
3. (Optional) Select **Use only client certificate authentication (Web)** to have bi-directional authentication; to get the server to authenticate the client.
4. Enter a **Password** and **Confirm Password**.
5. (Optional) Select **Use Public Key Authentication (SSH)** if you want to use a much stronger, key-based authentication method using an SSH public key rather than just a password.
6. For **Administrator Type**, select **Role Based**.
7. For **Profile**, select the profile that you just created.
8. (Optional) Select a **Password Profile**.
9. Click **OK**.

STEP 3 | Commit the configuration.

Click **Commit**.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Virtual System Functionality with Other Features

Many firewall features and functionality are capable of being configured, viewed, logged, or reported per virtual system. Therefore, virtual systems are mentioned in other relevant locations in the documentation and that information is not repeated here. Some of the specific chapters are the following:

- If you are configuring Active/Passive HA, the two firewalls must have the same virtual system capability (single or multiple virtual system capability). See High Availability.
- To configure QoS for virtual systems, see Configure QoS for a Virtual System.
- For information about configuring a firewall with virtual systems in a virtual wire deployment that uses subinterfaces (and VLAN tags), see Virtual Wire Interfaces.
- If you have configured User-ID and multiple virtual systems, you can share user mappings across virtual systems. See Share User-ID Mappings Across Virtual Systems.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Zone Protection and DoS Protection

Segmenting the network into functional and organizational zones reduces the network's attack surface—the portion of the network exposed to potential attackers. Zone protection defends network zones against flood attacks, reconnaissance attempts, packet-based attacks, and attacks that use non-IP protocols. Tailor a Zone Protection profile to protect each zone (you can apply the same profile to similar zones). Denial-of-service (DoS) protection defends specific critical systems against flood attacks, especially devices that user access from the internet such as web servers and database servers, and protects resources from session floods. Tailor DoS Protection profiles and policy rules to protect each set of critical devices. Visit the Best Practices documentation portal to get a checklist of Zone Protection and DoS Protection best practices.



*Check and monitor firewall dataplane CPU consumption to ensure that each firewall is properly sized to support DoS and Zone Protection along with any other features that consume CPU cycles, such as decryption. If you use Panorama to manage your firewalls, use Device Monitor (**Panorama** > **Managed Devices** > **Health**) to check and monitor the CPU consumption of all managed firewalls at one time.*

- Network Segmentation Using Zones
- How Do Zones Protect the Network?
- Zone Defense
- Configure Zone Protection to Increase Network Security
- DoS Protection Against Flooding of New Sessions

https://t.me/learningnets

Network Segmentation Using Zones

The larger the network, the more difficult it is to protect. A large, unsegmented network presents a large attack surface that can be difficult to manage and protect. Because traffic and applications have access to the entire network, once an attacker gains entry to a network, the attacker can move laterally through the network to access critical data. A large network is also more difficult to monitor and control. Segmenting the network limits an attacker's ability to move through the network by preventing lateral movement between zones.

A security zone is a group of one or more physical or virtual firewall interfaces and the network segments connected to the zone's interfaces. You control protection for each zone individually so that each zone receives the specific protections it needs. For example, a zone for the finance department may not need to allow all of the applications that a zone for IT allows.

To fully protect your network, all traffic must flow through the firewall. Configure Interfaces and Zones to create separate zones for different functional areas such as the internet gateway, sensitive data storage, and business applications, and for different organizational groups such as finance, IT, marketing, and engineering. Wherever there is a logical division of functionality, application usage, or user access privileges, you can create a separate zone to isolate and protect the area and apply the appropriate security policy rules to prevent unnecessary access to data and applications that only one or some groups need to access. The more granular the zones, the greater the visibility and control you have over network traffic. Dividing your network into zones helps to create a Zero Trust architecture that executes a security philosophy of trusting no users, devices, applications, or packets, and verifying everything. The end goal is to create a network that allows access only to the users, devices, and applications that have legitimate business needs, and to deny all other traffic.

How to appropriately restrict and permit access to zones depends on the network environment. For example, environments such as semiconductor manufacturing floors or robotic assembly plants, where the workstations control sensitive manufacturing equipment, or highly restricted access areas, may require physical segmentation that permits no access from outside devices (no mobile device access).

In environments where users can access the network with mobile devices, enabling User-ID and App-ID in conjunction with segmenting the network into zones ensures that users receive the appropriate access privileges regardless of where they access the network, because access privileges are tied to a user or a user group instead of to a device in one particular zone.

The protection requirements for different functional areas and groups may also differ. For example, a zone that handles a large amount of traffic may require different flood protection thresholds than a zone that normally handles less traffic. The ability to define the appropriate protection for each zone is another reason to segment the network. What appropriate protection is depends on your network architecture, what you want to protect, and what traffic you want to permit and deny.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

How Do Zones Protect the Network?

Zones not only protect your network by segmenting it into smaller, more easily managed areas, zones also protect the network because you can control access to zones and traffic movement between zones.

Zones prevent uncontrolled traffic from flowing through the firewall interfaces into your network because firewall interfaces can't process traffic until you assign them to zones. The firewall applies zone protection on ingress interfaces, where traffic enters the firewall in the direction of flow from the originating client to the responding server (c2s), to filter traffic before it enters a zone.

The firewall interface type and the zone type (Tap, virtual wire, L2, L3, Tunnel, or External) must match, which helps to protect the network against admitting traffic that doesn't belong in a zone. For example, you can assign an L2 interface to an L2 zone or an L3 interface to an L3 zone, but you can't assign an L2 interface to an L3 zone.

In addition, a firewall interface can belong to one zone only. Traffic destined for different zones can't use the same interface, which helps to prevent inappropriate traffic from entering a zone and enables you to configure the protection appropriate for each individual zone. You can connect more than one firewall interface to a zone to increase bandwidth, but each interface can connect to only one zone.

After the firewall admits traffic to a zone, traffic flows freely within that zone and is not logged. The more granular you make each zone, the greater the control you have over the traffic that accesses each zone, and the more difficult it is for malware to move laterally across the network between zones. Traffic can't flow between zones unless a security policy rule allows it and the zones are of the same zone type (Tap, virtual wire, L2, L3, Tunnel, or External). For example, a security policy rule can allow traffic between two L3 zones, but not between an L3 zone and an L2 zone. The firewall logs traffic that flows between zones when a security policy rule permits interzone traffic.

By default, security policy rules prevent lateral movement of traffic between zones, so malware can't gain access to one zone and then move freely through the network to other targets.



Tunnel zones are for non-encrypted tunnels. You can apply different security policy rules to the tunnel content and to the zone of the outer tunnel, as described in the Tunnel Content Inspection Overview.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Zone Defense

Zone Protection profiles defend zones against flood, reconnaissance, packet-based, and non-IP-protocol-based attacks. DoS Protection profiles used in DoS Protection policy rules defend specific, critical devices against targeted flood and resource-based attacks. A DoS attack overloads the network or targeted critical systems with large amounts of unwanted traffic an attempt to disrupt network services.

Plan to defend your network against different types of DoS attacks:

- **Application-Based Attacks**—Target weaknesses in a particular application and try to exhaust its resources so legitimate users can't use it. An example of this is the Slowloris attack.
- **Protocol-Based Attacks**—Also known as state-exhaustion attacks, these attacks target protocol weaknesses. A common example is a SYN flood attack.
- **Volumetric Attacks**—High-volume attacks that attempt to overwhelm the available network resources, especially bandwidth, and bring down the target to prevent legitimate users from accessing those resources. An example of this is a UDP flood attack.

There are no default Zone Protection profiles or DoS Protection profiles and DoS Protection policy rules. Configure and apply zone protection based on each zone's traffic characteristics and configure DoS protection based on the individual critical systems you want to protect in each zone.

- Zone Defense Tools
- How Do the Zone Defense Tools Work?
- Firewall Placement for DoS Protection
- Zone Protection Profiles
- Packet Buffer Protection
- DoS Protection Profiles and Policy Rules

Zone Defense Tools

Effective defense against DoS attacks requires a layered approach. The first layer of defense should be a dedicated, high-volume DDoS protection device at the internet-facing network perimeter and a perimeter router, switch, or other hardware-based packet drop device with appropriate access control lists (ACLs) to defend against volumetric attacks that the session-based firewall isn't designed to handle. The firewall adds more granular layers of DoS attack defense and also visibility into application traffic that dedicated DDoS devices don't provide.

Palo Alto Networks firewalls provide four complementary tools to layer in DoS protection for your network zones and critical devices:

- **Zone Protection profiles** defend the ingress zone edge against IP flood attacks, reconnaissance port scans and host sweeps, IP packet-based attacks, and non-IP protocol attacks. The ingress zone is where traffic enters the firewall in the direction of flow from the client to the server (c2s), where the client is the originator of the flow and the server is the responder. Zone Protection profiles provide a second layer of broad defense against DoS attacks, based on the aggregate traffic entering the zone, by limiting the new connections-per-second (CPS) to the

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

zone. Zone Protection profiles don't take individual devices (IP addresses) into account because the profiles apply to the aggregate traffic entering the zone.

Zone protection profiles defend the network as a session is formed, before the firewall performs DoS Protection policy and Security policy rule lookups, and consume fewer CPU cycles than a DoS Protection policy or Security policy rule lookup. If a Zone Protection profile denies traffic, the firewall doesn't spend CPU cycles on policy rule lookups.

Apply Zone Protection profiles to every zone, both internet-facing and internal.

- **DoS Protection profiles and policy rules** defend specific individual endpoints and resources against flood attacks, especially high-value targets that users access from the internet. While a Zone Protection profile defends the zone from flood attacks, a DoS Protection policy rule with an appropriate DoS Protection profile defends critical individual systems in a zone from targeted flood attacks, providing a granular third layer of defense against DoS attacks.

  Because the intent of DoS protection is to defend critical devices and because it consumes resources, DoS protection defends only the devices you specify in a DoS Protection policy rule. No other devices are protected.

  DoS Protection profiles set flood protection thresholds (new CPS limits) for individual devices or groups of devices, resource protection thresholds (session limits for specified endpoints and resources), and whether the profile applies to aggregate or classified traffic. DoS Protection policy rules specify match criteria (source, destination, service ports), the action to take when traffic matches the rule, and the aggregate and classified DoS Protection profiles associated with each rule.

  Aggregate DoS Protection policy rules apply the CPS thresholds defined in an aggregate DoS Protection profile to the combined traffic of all the devices that meet the DoS Protection policy rule match criteria. For example, if you configure the aggregate DoS Protection profile to limit the CPS rate to 20,000, the 20,000 CPS limit applies to the aggregate number of connections for the entire group. In this case, one device could receive the majority of the allowed connections.

  Classified DoS Protection policy rules apply the CPS thresholds defined in a classified DoS Protection profile to each individual device that matches the policy rule. For example, if you configure the classified DoS Protection profile to limit the CPS rate to 4,000, then no device in the group can accept more than 4,000 CPS. A DoS Protection policy can have one aggregate profile and one classified profile.

  Classified profiles can classify connections by source IP, destination IP, or both. For internet-facing zones, classify by destination IP only because the firewall can't scale to hold the internet routing table.

  Apply DoS Protection only to critical devices, especially popular attack targets that users access from the internet, such as web servers and database servers.

- For existing sessions, **Packet Buffer Protection** protects the firewall (and therefore the zone) against single-session DoS attacks that attempt to overwhelm the firewall's packet buffer, using thresholds and timers to mitigate abusive sessions. You configure Packet Buffer Protection settings globally and apply them per zone.
- **Security Policy rules** affect both the ingress and egress flows of a session. To establish a session, incoming traffic must match an existing Security policy rule. If there is no match, the

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

firewall discards the packet. A Security policy allows or denies traffic between zones (interzone) and within zones (intrazone) using criteria including zones, IP addresses, users, applications, services, and URL categories.

Apply the best practice Vulnerability Protection profile to each Security policy rule to help defend against DoS attacks.

The default Security policy rules don't permit traffic to travel between zones, so you need to configure a Security policy rule if you want to allow interzone traffic. All intrazone traffic is allowed by default. You can configure Security policy rules to match and control intrazone, interzone, or universal (intrazone and interzone) traffic.

Zone Protection profiles, DoS Protection profiles and policy rules, and Security policy rules only affect dataplane traffic on the firewall. Traffic originating on the firewall management interface does not cross the dataplane, so the firewall does not match management traffic against these profiles or policy rules.

- You can also search the Palo Alto Networks Threat Vault (requires a valid support account and login) for threats by hash, CVE, signature ID, domain name, URL, or IP address.

How Do the Zone Defense Tools Work?

When a packet arrives at the firewall, the firewall attempts to match the packet to an existing session, based on the ingress zone, egress zone, source IP address, destination IP address, protocol, and application derived from the packet header. If the firewall finds a match, then the packet uses the Security policy rules that already control the session. If the packet doesn't match an existing session, the firewall uses Zone Protection profiles, DoS Protection profiles and policy rules, and Security policy rules to determine whether to establish a session or discard the packet, and the level of access the packet receives.

After traffic passes through your dedicated DDoS device at the internet-facing network edge, the first protection the firewall applies is the broad defense of the Zone Protection profile, if one is attached to the zone. The firewall determines the zone from the interface on which the packet arrives (each interface is assigned to only one zone and all interfaces that carry traffic must belong to a zone). If the Zone Protection profile denies the packet, the firewall discards the packet and saves resources by not needing to look up the DoS Protection policy or Security policy. The firewall applies Zone Protection profiles only to new sessions (packets that do not match an existing session). After the firewall establishes a session, the firewall bypasses the Zone Protection profile lookup for succeeding packets in that session.

If the Zone Protection profile doesn't drop the packet, the second protection the firewall applies is a DoS Protection policy rule. If a Zone Protection profile allows a packet based on the total aggregate amount of traffic going to the zone, a DoS Protection policy rule may deny the packet if it is going to a particular destination or coming from a particular source that has exceeded the flood protection or resource protection settings in the rule's DoS Protection profile. If the packet matches a DoS Protection policy rule, the firewall applies the rule to the packet. If the rule denies access, the firewall discards the packet and doesn't perform a Security policy lookup. If the rule allows access, the firewall performs a Security policy lookup. Like the Zone Protection profile, the firewall enforces DoS Protection policy only on new sessions.

The third protection the firewall applies is a Security policy lookup, which happens only if the Zone Protection profile and DoS Protection policy rules allow the packet. If the firewall finds no

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Security policy rule match for the packet, the firewall discards the packet. If the firewall finds a matching Security policy rule, the firewall applies the rule to the packet. The firewall enforces the Security policy rule on traffic in both directions (c2s and s2c) for the life of the session. Apply the best practice Vulnerability Protection profile to all Security policy rules to help defend against DoS attacks.

The fourth protection the firewall applies is packet buffer protection, which you apply globally to protect the device and can also apply individually to zones to prevent single-session DoS attacks that attempt to overwhelm the firewall's packet buffer. For global protection, the firewall used Random Early Drop (RED) to drop packets (not sessions) when the level of traffic crosses protection thresholds. For per-zone protection, the firewall blocks the source IP address if it violates the packet buffer thresholds. Unlike zone and DoS protection, packet buffer protection applies to existing sessions.

Firewall Placement for DoS Protection

The firewall is a session-based device that isn't designed to scale to millions of connections-per-second (CPS) to defend against large volumetric DoS attacks. The firewall treats each unique flow (based on ingress and egress zone, source and destination IP, protocol, and application) as a session, spends CPU cycles on packet inspection at the port and the IP level to provide visibility into application traffic, and must count each session for the flood threshold counters, so firewall placement is critical to avoid flooding the firewall.

For the best DoS protection, *place firewalls as close to the resources you're protecting as possible.* This reduces the number of sessions the firewall needs to handle and therefore the amount of firewall resources required to provide DoS protection.

At the internet-facing perimeter, do *not* place firewalls you use for DoS protection or zone protection in front of dedicated DDoS devices and perimeter routers and switches. Make those high-volume devices your first line of DoS defense to mitigate volumetric flood attacks. For zone and DoS protection at the perimeter, use high-capacity firewalls and place them *behind* the high-volume devices. As a rule, the closer a firewall is to the perimeter, the higher capacity it must be to handle the volume of traffic.

The way you segment your network into zones can help mitigate internal DoS attacks. Smaller zones provide greater visibility into traffic and prevent lateral movement of malware better because more traffic must cross zones, and to allow interzonal traffic requires you to create a specific Security policy rule (all intrazonal traffic is allowed by default). Consider revisiting your segmentation approach if your network is relatively unsegmented.

Baseline CPS Measurements for Setting Flood Thresholds

Flood protection thresholds determine the number of new connections-per-second (CPS) to allow for a zone (Zone Protection profile), for a group of devices in a zone (aggregate DoS Protection policy), or for individual devices in a zone (classified DoS Protection policy), when to throttle new connections to begin mitigating a potential flood attack, and when to drop all new connections. The default Zone Protection profile and DoS Protection profile flood protection thresholds aren't appropriate for most networks because each network is unique. You need to understand the aggregate normal and peak CPS for each zone to set effective Zone Protection profile thresholds, and for the individual critical systems you want to defend to set effective DoS Protection profile thresholds that don't inadvertently set thresholds too high and allow flood attacks or set thresholds too low and throttle traffic.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

- CPS Measurements to Take
- How to Measure CPS

CPS Measurements to Take

Measure average and peak CPS traffic over the course of at least five business days or until you're confident that the measurements reflect the network's typical traffic patterns; the longer measurement period, the more accurate the measurements. Take into account special events, quarterly events, and annual events that may spike the number of CPS you need to support. You may need to adjust Zone Protection profiles and schedule adjusted DoS Protection policy rules to accommodate these types of events if your firewalls have the capacity to handle extra traffic. Take the following baseline measurements:

- For Zone Protection profiles, measure the average and peak CPS ingressing each zone.
- For aggregate DoS Protection profiles, measure the combined average and peak CPS for each group of devices you want to protect.
- For classified DoS Protection profiles, measure the average and peak CPS of the individual devices you want to protect.

Also understand the capacity of your firewalls and how other resource-consuming features such as decryption affect the number of connections each firewall can control. As a general rule, the closer a firewall is to the perimeter, the greater its capacity needs to be because it handles more traffic. The datasheet for each firewall model includes the total new sessions per second (CPS) the firewall supports and the Firewall Comparison Tool enables you to compare the CPS (and other metrics) of different firewall models.

How to Measure CPS

There are many ways to measure CPS:

- If you use Panorama to manage your firewalls, use Device Monitoring to measure CPS coming into a firewall (**Panorama** > **Managed Devices** > **Health** > **All Devices**). Device Monitoring can also show you a 90-day trend line of CPU average and peak use to help you understand the typical available capacity of each firewall.
- Run the operational CLI command **`show session info`**.

  *The operational CLI command **`show counter interface`** displays two times the actual CPS value. If you use this command, divide the CPS value by two to derive the real CPS value.*

- For setting appropriate DoS Protection profile thresholds, work with application teams to understand the normal and peak CPS to their servers and the maximum CPS those servers can support.

  In addition, you can filter firewall Traffic logs and Threat logs for the destination IP addresses of the critical devices you want to protect to obtain normal and peak session activity information.
- Use third-party tools such as Wireshark or NetFlow to collect and analyze network traffic.
- Use scripts to automate CPS information collection and continuous monitoring, and to mine information from the logs.
- Configure every Security policy rule on the firewall to **Log at Session End**. If you have no monitoring tools such as NetFlow or Wireshark, and cannot obtain or develop automated

https://t.me/learningnets

scripts, **Log at Session End** captures the number of connections at the session end. While this doesn't provide CPS information, it does show you the number of sessions ending in the selected time duration and you can make an approximate calculation of the sessions per second from that information.

To conserve resources, the firewall measures the aggregate CPS at ten-second intervals. For this reason, measurements you see on the firewall may not catch bursts within the ten-second interval. Although the average CPS measurements aren't affected, the peak CPS measurements may not be precise. For example, if the firewall logs report a 5,000 CPS average in a ten-second interval, it's possible that 4,000 CPS came in a one-second burst and the other 1,000 CPS were spread out over the remaining nine seconds.

To gather historical CPS data over time, if you use an SNMP server, you can use your own management tools to poll SNMP MIBs. However, it is important to understand that the CPS measurements in the MIBs show twice the actual CPS value (for example, if the true CPS measurement is 10,000, the MIBs show 20,000 as the value). You can still see trends from the MIBs and you can divide the CPS values by two to derive the true values. The SNMP MIB OIDs are: PanZoneActiveTcpCps, PanZoneActiveUdpCps, and PanZoneOtherIpCps. Because the firewall only takes measurements and updates the SNMP server every 10 seconds, poll every 10 seconds.

In addition, create separate log forwarding profiles for flood events so the appropriate administrator receives emails that contain only flood (potential DoS attack) events. Set Log Forwarding for both zone protection and DoS protection threshold events.

After you implement Zone and DoS protection, use these methods to monitor the deployment, so as your network evolves and traffic patterns change, you adjust flood protection thresholds.

Zone Protection Profiles

Apply a Zone Protection profile to each zone to defend it based on the aggregate traffic entering the ingress zone.

In addition to configuring zone protection and DoS protection, apply the best practice Vulnerability Protection profile to each Security policy rule to help defend against DoS attacks.

- Flood Protection
- Reconnaissance Protection
- Packet-Based Attack Protection
- Protocol Protection
- Ethernet SGT Protection

Flood Protection

A Zone Protection profile with flood protection configured defends an entire ingress zone against SYN, ICMP, ICMPv6, UDP, and other IP flood attacks. The firewall measures the aggregate amount of each flood type entering the zone in new connections-per-second (CPS) and compares the

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

totals to the thresholds you configure in the Zone Protection profile. (You protect critical individual devices within a zone with DoS Protection profiles and policy rules.)

Measure and monitor firewall dataplane CPU consumption to ensure that each firewall is properly sized to support DoS and Zone Protection and any other features that consume CPU cycles, such as decryption. If you use Panorama to manage your firewalls, Device Monitoring *(**Panorama** > **Managed Devices** > **Health** > **All Devices**) shows you the CPU and memory consumption of each managed firewall. It can also show you a 90-day trend line of CPU average and peak use to help you understand the typical available capacity of each firewall.*

For each flood type, you set three thresholds for new CPS entering the zone, and you can set a drop **Action** for SYN floods. If you know the baseline CPS rates for the zone, use these guidelines to set the initial thresholds, and then monitor and adjust the thresholds as necessary.

- **Alarm Rate**—The new CPS threshold to trigger an alarm. Target setting the **Alarm Rate** to 15-20% above the average CPS rate for the zone so that normal fluctuations don't cause alerts.
- **Activate**—The new CPS threshold to activate the flood protection mechanism and begin dropping new connections. For ICMP, ICMPv6, UDP, and other IP floods, the protection mechanism is Random Early Drop (RED, also known as Random Early Detection). For SYN floods only, you can set the drop **Action** to SYN Cookies or RED. Target setting the **Activate** rate to just above the peak CPS rate for the zone to begin mitigating potential floods.
- **Maximum**—The number of connections-per-second to drop incoming packets when RED is the protection mechanism. Target setting the **Maximum** rate to approximately 80-90% of firewall capacity, taking into account other features that consume firewall resources.

If you don't know the baseline CPS rates for the zone, start by setting the **Maximum** CPS rate to approximately 80-90% of firewall capacity and use it to derive reasonable flood mitigation alarm and activation rates. Set the **Alarm Rate** and **Activate** rate based on the Maximum rate. For example, you could set the **Alarm Rate** to half the **Maximum** rate and adjust it depending on how many alarms you receive and the firewall resources being consumed. Be careful setting the **Activate Rate** since it begins to drop connections. Because normal traffic loads experience some fluctuation, it's best not to drop connections too aggressively. Err on the high side and adjust the rate if firewall resources are impacted.

*SYN Flood Protection is the only type for which you set the drop **Action**. Start by setting the **Action** to **SYN Cookies**. SYN Cookies treats legitimate traffic fairly and only drops traffic that fails the SYN handshake, while using Random Early Drop drops traffic randomly, so RED may affect legitimate traffic. However, SYN Cookies is more resource-intensive because the firewall acts as a proxy for the target server and handles the three-way handshake for the server. The tradeoff is not dropping legitimate traffic (SYN Cookies) versus preserving firewall resources (RED). Monitor the firewall, and if SYN Cookies consumes too many resources, switch to RED. If you don't have a dedicated DDoS prevention device in front of the firewall, always use RED as the drop mechanism.*

*When **SYN Cookies** is activated, the firewall does not honor the TCP options that the server sends because it does not know these values at the time that it proxies the SYN/ACK. Therefore, values such as the TCP server's window size and MSS values cannot be negotiated during the TCP handshake and the firewall will use its own default values. In the scenario where the MSS of the path to the server is smaller than the firewall's default MSS value, the packet will need to be fragmented.*

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

The default threshold values are high so that activating a Zone Protection profile doesn't unexpectedly drop legitimate traffic. Adjust the thresholds to values appropriate for your network's traffic. The best method for understanding how to set reasonable flood thresholds is to take baseline measurements of average and peak CPS for each flood type to determine the normal traffic conditions for each zone and to understand the capacity of the firewall, including the impact of other resource-consuming features such as decryption. Monitor and adjust the flood thresholds as needed and as your network evolves.



*Firewalls with multiple dataplane processors (DPs) distribute connections across DPs. In general, the firewall divides the CPS threshold settings equally across its DPs. For example, if a firewall has five DPs and you set the **Alarm Rate** to 20,000 CPS, each DP has an **Alarm Rate** of 4,000 CPS (20,000 / 5 = 4,000), so if the new sessions on a DP exceeds 4,000, it triggers the **Alarm Rate** threshold for that DP.*

Reconnaissance Protection

Similar to the military definition of reconnaissance, the network security definition of reconnaissance is when attackers attempt to gain information about your network's vulnerabilities by secretly probing the network to find weaknesses. Reconnaissance activities are often preludes to a network attack. *Enable Reconnaissance Protection on all zones* to defend against port scans and host sweeps:

- **Port scans** discover open ports on a network. A port scanning tool sends client requests to a range of port numbers on a host, with the goal of locating an active port to exploit in an attack. Zone Protection profiles defend against TCP and UDP port scans.
- **Host sweeps** examine multiple hosts to determine if a specific port is open and vulnerable.

You can use reconnaissance tools for legitimate purposes such as pen testing of network security or the strength of a firewall. You can specify up to 20 IP addresses or netmask address objects to exclude from Reconnaissance Protection so that your internal IT department can conduct pen tests to find and fix network vulnerabilities.

You can set the action to take when reconnaissance traffic (excluding pen testing traffic) exceeds the configured threshold when you Configure Reconnaissance Protection. Retain the default **Interval** and **Threshold** to log a few packets for analysis before blocking the reconnaissance operation.

Packet-Based Attack Protection

Packet-based attacks take many forms. Zone Protection profiles check IP, TCP, ICMP, IPv6, and ICMPv6 packet headers and protect a zone by:

- Dropping packets with undesirable characteristics.
- Stripping undesirable options from packets before admitting them to the zone.

Select the drop characteristics for each packet type when you Configure Packet Based Attack Protection. The best practices for each IP protocol are:

- **IP Drop**—Drop **Unknown** and **Malformed** packets. Also drop **Strict Source Routing** and **Loose Source Routing** because allowing these options allows adversaries to bypass Security policy rules that use the Destination IP address as the matching criteria. For internal zones only, check **Spoofed IP Address** so only traffic with a source address that matches the firewall routing table can access the zone.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

- **TCP Drop**—Retain the default **TCP SYN with Data** and **TCP SYNACK with Data** drops, drop **Mismatched overlapping TCP segment** and **Split Handshake** packets, and strip the **TCP Timestamp** from packets.

  *Enabling **Rematch Sessions** (**Device** > **Setup** > **Session** > **Session Settings**) is a best practice that applies committed newly configured or edited Security Policy rules to existing sessions. However, if you configure Tunnel Content Inspection on a zone and **Rematch Sessions** is enabled, you must also disable **Reject Non-SYN TCP** (change the selection from **Global** to **No**), or else when you enable or edit a Tunnel Content Inspection policy, the firewall drops all existing tunnel sessions. Create a separate Zone Protection profile to disable **Reject Non-SYN TCP** only on zones that have Tunnel Content Inspection policies and only when you enable **Rematch Sessions**.*

- **ICMP Drop**—There are no standard best practice settings because dropping ICMP packets depends on how you use ICMP (or if you use ICMP). For example, if you want to block ping activity, you can block **ICMP Ping ID 0**.
- **IPv6 Drop**—If compliance matters, ensure that the firewall drops packets with non-compliant routing headers, extensions, etc.
- **ICMPv6 Drop**—If compliance matters, ensure that the firewall drops certain packets if the packets don't match a Security policy rule.

Protocol Protection

In a Zone Protection profile, Protocol Protection defends against non-IP protocol based attacks. Enable Protocol Protection to block or allow non-IP protocols between security zones on a Layer 2 VLAN or on a virtual wire, or between interfaces within a single zone on a Layer 2 VLAN (Layer 3 interfaces and zones drop non-IP protocols so non-IP Protocol Protection doesn't apply). Configure Protocol Protection to reduce security risks and facilitate regulatory compliance by preventing less secure protocols from entering a zone, or an interface in a zone.

If you don't configure a Zone Protection profile that prevents non-IP protocols in the same zone from going from one Layer 2 interface to another, the firewall allows the traffic because of the default intrazone allow Security policy rule. You can create a Zone Protection profile that blocks protocols such as LLDP within a zone to prevent discovery of networks reachable through other zone interfaces.

If you need to discover which non-IP protocols are running on your network, use monitoring tools such as NetFlow, Wireshark, or other third-party tools discover non-IP protocols on your network. Examples of non-IP protocols you can block or allow are LLDP, NetBEUI, Spanning Tree, and Supervisory Control and Data Acquisition (SCADA) systems such as Generic Object Oriented Substation Event (GOOSE), among many others.

Create an **Exclude List** or an **Include List** to configure Protocol Protection for a zone. The **Exclude List** is a block list—the firewall blocks all of the protocols you place in the **Exclude List** and allows all other protocols. The **Include List** is an allow list—the firewall allows only the protocols you specify in the list and blocks all other protocols.

Use include lists for Protocol Protection instead of exclude lists. Include lists specifically sanction only the protocols you want to allow and block the protocols you don't need or didn't know were on your network, which reduces the attack surface and blocks unknown traffic.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

A list supports up to 64 Ethertype entries, each identified by its IEEE hexadecimal Ethertype code. Other sources of Ethertype codes are standards.ieee.org/develop/regauth/ethertype/eth.txt and http://www.cavebear.com/archive/cavebear/Ethernet/type.html. When you configure zone protection for non-IP protocols on zones that have Aggregated Ethernet (AE) interfaces, you can't block or allow a non-IP protocol on only one AE interface member because AE interface members are treated as a group.

*Protocol Protection doesn't allow blocking IPv4 (Ethertype 0x0800), IPv6 (0x86DD), ARP (0x0806), or VLAN-tagged frames (0x8100). The firewall always implicitly allows these four Ethertypes in an **Include List** even if you don't explicitly list them and doesn't permit you to add them to an **Exclude List**.*

Ethernet SGT Protection

In a Cisco TrustSec network, a Cisco Identity Services Engine (ISE) assigns a Layer 2 Security Group Tag (SGT) of 16 bits to a user's or endpoint's session. You can create a Zone Protection profile with Ethernet SGT protection when your firewall is part of a Cisco TrustSec network. The firewall can inspect headers with 802.1Q (Ethertype 0x8909) for specific Layer 2 security group tag (SGT) values and drop the packet if the SGT matches the list you configure for the Zone Protection profile attached to the interface. Determine which SGT values you want to deny access to a zone.

Packet Buffer Protection

Packet Buffer Protection defends your firewall and network from single session DoS attacks that can overwhelm the firewall's packet buffer and cause legitimate traffic to drop. Although you don't configure Packet Buffer Protection in a Zone Protection profile or in a DoS Protection profile or policy rule, Packet Buffer Protection defends ingress zones. While zone and DoS protection apply to new sessions (connections) and are granular, Packet Buffer Protection applies to existing sessions and is global.

You Configure Packet Buffer Protection globally to protect the entire firewall and you also enable Packet Buffer Protection on each zone to protect zones:

- **Global Packet Buffer Protection**—The firewall monitors sessions from all zones (regardless of whether Packet Buffer Protection is enabled in a zone) and how those sessions utilize the packet buffer. You must configure Packet Buffer Protection globally (**Device** > **Setup** > **Session Settings**) to protect the firewall and to enable it on individual zones. When packet buffer consumption reaches the configured **Activate** percentage, the firewall used Random Early Drop (RED) to drop packets from the offending sessions (the firewall doesn't drop complete sessions at the global level).
- **Per-Zone Packet Buffer Protection**—Enable Packet Buffer Protection on each zone (**Network** > **Zones**) to layer in a second level of protection. When packet buffer consumption crosses the **Activate** threshold and global protection begins to apply RED to session traffic, that starts the **Block Hold Time** timer. The **Block Hold Time** is the amount of time in seconds that the offending session can continue before the firewall blocks the entire session. The offending session remains blocked until the **Block Duration** time expires.

You must enable Packet Buffer Protection globally in order for it to be active in zones.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

There are two types of packet buffer protection:

- Packet Buffer Protection Based on Buffer Utilization
- Packet Buffer Protection Based on Latency

Packet Buffer Protection Based on Buffer Utilization

Packet Buffer Protection based on buffer utilization is enabled by default. Take baseline measurements of firewall packet buffer utilization over a period of time—at least one business week, but a longer measurement period provides a better baseline—to understand typical usage.

To see packet buffer utilization for a specified period of time (or to see the top five sessions that use as least 2 percent of the packet buffer), use the operational CLI command:

```
admin1138@thxvm1>show running resource-monitor [day | hour | ingress-
backlogs | minute | second | week]
```

The CLI command provides a snapshot of buffer utilization for the specified period of time, but is neither automated nor continuous. To automate continuous packet buffer utilization measurements so you can monitor changes in behavior and anomalous events, use a script. Your Palo Alto Networks account team can provide a sample script that you can modify to develop your own script; however, the script is not officially supported and there is no technical support available for script usage or modification.

If baseline measurements consistently show abnormally high packet buffer utilization, then the firewall's capacity may be undersized for typical traffic loads. In this case, consider resizing the firewall deployment. Otherwise, you need to tune the Packet Buffer Protection thresholds carefully to prevent impacted buffers from overflowing (and to prevent dropping legitimate traffic). When firewall sizing is correct for the deployment, only an attack should cause a large spike in buffer usage.



Overrunning the firewall packet buffer negatively impacts the firewall's packet forwarding capabilities. When the buffers are full, no packets can enter the firewall on any interface, not just the interface that experienced the attack.

The best practices for setting the thresholds are:

- **Alert** and **Activate**—Start with the default threshold values, monitor packet buffer utilization, and adjust the thresholds as necessary. The **Alert** threshold defaults to 50%; when packet buffer utilization exceeds the threshold for more than 10 seconds, the firewall creates an alert entry in the System log every minute. The **Activate** threshold defaults to 80%; when the threshold is reached, the firewall begins to mitigate the most abusive sessions. If the firewall is sized correctly, buffer utilization should be well below 50%.
- **Block Hold Time**—When packet buffer utilization triggers the **Activate** threshold, the **Block Hold Time** sets the amount of time the offending session can continue before the firewall blocks the session. During the **Block Hold Time**, the firewall continues to apply RED to the packets of offending sessions. Start with the default **Block Hold Time** threshold value (60 seconds), monitor packet buffer utilization, and adjust the threshold as necessary. If the packet buffer utilization percentage falls below the **Activate** threshold before the **Block Hold Time** expires, the timer resets and doesn't start until the **Activate** threshold is crossed again.

https://t.me/learningnets

Increasing the **Block Hold Time** imposes a greater penalty on offending sessions and reducing it imposes a lesser penalty on offending sessions.

- **Block Duration**—When the **Block Hold Time** expires, the firewall blocks the offending session for the period of time defined by the **Block Duration**. Start with the default threshold value (3600 seconds), monitor packet buffer utilization, and adjust the threshold as necessary. When you enable Packet Buffer Protection on a zone, **Block Duration** affects every session from the IP address even if only one session from an IP address overutilizes the packet buffer. If you believe that blocking an IP address for one hour (3600 seconds) is too great a penalty, reduce the **Block Duration** to an acceptable value.

In addition to monitoring the buffer utilization of individual sessions, Packet Buffer Protection can also block an IP address if certain criteria are met. While the firewall monitors the packet buffers, if it detects a source IP address rapidly creating sessions that would not individually be seen as an attack, it blocks that IP address for the configured **Block Duration**.

Network Address Translation (NAT) *(an external source that has translated its internet-bound traffic using source NAT) can give the appearance of greater packet buffer utilization because of IP address translation activity. If this occurs, adjust the thresholds in a way that penalizes individual sessions but doesn't penalize the underlying IP addresses (so other sessions from the same IP address aren't affected). To do this, reduce the* ***Block Hold Time*** *so the firewall blocks individual sessions that overutilize the buffers faster, and reduce the* ***Block Duration*** *so that the underlying IP address is not unduly penalized.*

Packet Buffer Protection Based on Latency

As an alternative to packet buffer protection based on utilization, you can trigger packet buffer protection based on packet latency caused by dataplane packet buffering, which indicates congestion on the firewall. Such packet buffer protection mitigates head-of-line blocking by alerting you to the congestion and performing random early drop (RED) on packets. Packet buffer protection based on latency can trigger the protection before latency-sensitive protocols or applications are affected.

If your traffic includes protocols or applications that are latency-sensitive, then packet buffer protection based on latency will be more helpful than packet buffer protection based on buffer utilization.

Packet buffer protection based on latency includes setting a **Latency Alert** threshold (in milliseconds), above which the firewall starts generating an Alert log event. The **Latency Activate** threshold indicates when the firewall activates RED on incoming packets and starts generating an Activate log. The **Latency Max Tolerate** threshold indicates when the firewall uses with RED with almost 100% drop probability.

The **Block Hold Time** and **Block Duration** settings function for packet buffer protection based on latency in the same way they do for packet buffer protection based on utilization.

DoS Protection Profiles and Policy Rules

DoS Protection profiles and DoS Protection policy rules combine to protect specific groups of critical resources and individual critical resources against session floods. Compared to Zone Protection profiles, which protect entire zones from flood attacks, DoS protection provides granular defense for specific systems, especially critical systems that users access from the internet and are often attack targets, such as web servers and database servers. Apply both types

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

of protection because if you only apply a Zone Protection profile, then a DoS attack that targets a particular system in the zone can succeed if the total connections-per-second (CPS) doesn't exceed the zone's **Activate** and **Maximum** rates.

DoS Protection is resource-intensive, so use it only for critical systems. Similar to Zone Protection profiles, DoS Protection profiles specify flood thresholds. DoS Protection policy rules determine the devices, users, zones, and services to which DoS Profiles apply.



In addition to configuring DoS protection and zone protection, apply the best practice Vulnerability Protection profile to each Security policy rule to help defend against DoS attacks.

- Classified Versus Aggregate DoS Protection
- DoS Protection Profiles
- DoS Protection Policy Rules

Classified Versus Aggregate DoS Protection

You can configure *aggregate* and *classified* DoS Protection Profiles, and apply one profile or one of each type of profile to DoS Protection Policy Rules when you configure DoS Protection.

- **Aggregate**—Sets thresholds that apply to the entire group of devices specified in a DoS Protection policy rule instead of to each individual device, so one device could receive the majority of the allowed connection traffic. For example, a **Max Rate** of 20,000 CPS means the total CPS for the group is 20,000, and an individual device can receive up to 20,000 CPS if other devices don't have connections. Aggregate DoS Protection policies provide another layer of broad protection (after your dedicated DDoS device at the internet perimeter and Zone Protection profiles) for a particular group of critical devices when you want to apply extra constraints on specific subnets, users, or services.
- **Classified**—Sets flood thresholds that apply to each individual device specified in a DoS Protection policy rule. For example, if you set an **Max Rate** of 5,000 CPS, each device specified in the rule can accept up to 5,000 CPS before it drops new connections. If you apply a classified DoS Protection policy rule to more than one device, the devices governed by the rule should be similar in terms of capacity and how you want to control their CPS rates because classified thresholds apply to each individual device. Classified profiles protect individual critical resources.

  When you configure a DoS Protection policy rule with a classified DoS Protection profile (**Option/Protection** > **Classified** > **Address**), use the **Address** field to specify whether incoming connections count toward the profile thresholds based on matching the **source-ip-only**, **destination-ip-only**, or **scr-dest-ip-both** (the firewall counts both the source and the destination IP addresses matches toward the thresholds). Counters consume resources, so the way you count address matches affects firewall resource consumption. You can use classified DoS protection to:

  - Protect critical individual devices, especially servers that users access from the internet and are often attack targets, such as web servers, database servers, and DNS servers. Set appropriate flood and resource protection thresholds in a classified DoS Protection profile. Create a DoS Protection policy rule that applies the profile to each server's IP address by

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

adding the IP addresses as the rule's destination criteria, and set the **Address** to **destination-ip-only**.

*Do not use **source-IP-only** or **src-dest-ip-both** classification for internet-facing zones in classified DoS Protection policy rules because the firewall doesn't have the capacity to store counters for every possible IP address on the internet. Increment the threshold counter for source IPs only for internal zone or same-zone rules. In perimeter zones, use **destination-ip-only**.*

- Monitor the CPS rate for a suspect host or group of hosts (the zone that contains the hosts cannot be internet-facing). Set an appropriate alarm threshold in a classified DoS Protection profile to notify you if a host initiates an unusually large number of connections. Create a DoS Protection policy rule that applies the profile to the individual source or source address group and set the **Address** to **source-ip-only**. Investigate hosts that initiate enough new connections to set off the alarm.

How you configure the **Address** (**source-ip-only**, **destination-ip-only**, or **src-dest-ip-both**) for classified profiles depends on your DoS protection goals, what you are protecting, and whether the protected device(s) are in internet-facing zones.

*The firewall uses more resources to track **src-dest-ip-both** as the **Address** than to track **source-IP-only** or **destination-ip-only** because the counters consume resources for both the source and destination IP addresses instead of just one of the two.*

If you apply both an aggregate and a classified DoS Protection profile to the same DoS Protection policy rule, the firewall applies the aggregate profile first and then applies the classified profile if needed. For example, we protect a group of five web servers with both types of profiles in a DoS Protection policy rule. The aggregate profile configuration drops new connections when the combined total for the group reaches a **Max Rate** of 25,000 CPS. The classified profile configuration drops new connections to any individual web server in the group when it reaches a **Max Rate** of 6,000 CPS. There are three scenarios where new connection traffic crosses **Max Rate** thresholds:

- The new CPS rate exceeds the aggregate **Max Rate** but doesn't exceed the classified **Max Rate**. In this scenario, the firewall applies the aggregate profile and blocks all new connections for the configured Block Duration.
- The new CPS rate doesn't exceed the aggregate **Max Rate**, but the CPS to one of the web servers exceeds the classified **Max Rate**. In this scenario, the firewall checks the aggregate profile and finds that the rate for the group is less than 25,000 CPS, so the firewall doesn't block new connections based on that. Next, the firewall checks the classified profile and finds that the rate for a particular server exceeds 6,000 CPS. The firewall applies the classified profile and blocks new connections to that particular server for the configured Block Duration. Because the other servers in the group are within the classified profile's **Max Rate**, their traffic is not affected.
- The new CPS rate exceeds the aggregate **Max Rate** and also exceeds the classified **Max Rate** for one of the web servers. In this scenario, the firewall checks the aggregate profile and finds that the rate for the group exceeds 25,000 CPS, so the firewall blocks new connections to limit the group's total CPS. The firewall then checks the classified profile and finds that the rate for a particular server exceeds 6,000 CPS (so the aggregate profile enforced the group's combined limit, but that wasn't enough to protect this particular server). The firewall applies the classified profile and blocks new connections to that particular server for the configured Block Duration.

https://t.me/learningnets

Because the other servers in the group are within the classified profile's **Max Rate**, their traffic is not affected.

*If you want both an aggregate and a classified DoS Protection profile to apply to the same traffic, you must apply both profiles to the same DoS Protection policy rule. If you apply the aggregate profile to one rule and the classified profile to a different rule, even if they specify exactly the same traffic, the firewall can apply only one profile because when the traffic matches the first DoS Protection policy rule, the firewall executes the **Action** specified in that rule and doesn't compare to the traffic to any subsequent rules, so the traffic never matches the second rule and the firewall can't apply its action. (This is the same way that Security policy rules work.)*

DoS Protection Profiles

DoS Protection profiles set thresholds that protect against new session IP flood attacks and provide resource protection (maximum concurrent session limits for specified endpoints and resources). DoS Protection profiles protect specific devices (classified profiles) and groups of devices (aggregate profiles) against SYN, UDP, ICMP, ICMPv6, and Other IP flood attacks. Configuring Flood Protection thresholds in a DoS Protection profile is similar to configuring Flood Protection in a Zone Protection profile, but Zone Protection profiles protect entire ingress zones, while DoS protection profiles and policy rules are granular and targeted, and can even be classified to a single device (IP address). The firewall measures the aggregate number of connections-per-second (CPS) to a group of devices (aggregate profile) or measures the CPS to individual devices (classified profile).

*Measure and monitor firewall dataplane CPU consumption to ensure that each firewall is properly sized to support DoS and Zone Protection and any other features that consume CPU cycles, such as decryption. If you use Panorama to manage your firewalls, Device Monitoring (**Panorama** > **Managed Devices** > **Health** > **All Devices**) shows you the CPU and memory consumption of each managed firewall. It can also show you a 90-day trend line of CPU average and peak use to help you understand the typical available capacity of each firewall.*

For each flood type, you set three thresholds for new CPS to a group of devices (aggregate) or to individual devices (classified) and a **Block Duration**, and you can set a drop **Action** for SYN floods:

- **Alarm Rate**—When new CPS exceeds this threshold, the firewall generates a DoS alarm. For classified profiles, set the rate to 15-20% above the device's average CPS rate so that normal fluctuations don't cause alerts. For aggregate profiles, set the rate to 15-20% above the group's average CPS rate.
- **Activate Rate**—When new CPS exceeds this threshold, the firewall begins to drop new connections to mitigate the flood until the CPS rate drops below the threshold. For classified profiles, the **Max Rate** should be an acceptable CPS rate for the device(s) you're protecting (the **Max Rate** won't flood the critical device(s)). You can set the **Activate Rate** to the same threshold as the **Max Rate** so that the firewall doesn't use RED or SYN Cookies to begin dropping traffic before it reaches the **Max Rate**. Set the **Activate Rate** lower than the **Max Rate** only if you want to drop traffic before it reaches the **Max Rate**. For aggregate profiles, set the threshold just above the average peak CPS rate for the group to begin mitigating floods using RED (or SYN Cookies for SYN floods).

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

- **Max Rate**—When new CPS exceeds this threshold, the firewall blocks (drops) all new connections from the offending IP address for the specified **Block Duration** time period. For classified profiles, base the **Max Rate** threshold on the capacity of the device(s) you're protecting so that the CPS rate can't flood them. For aggregate profiles, set to 80-90% of the group's capacity.
- **Block Duration**—When new CPS exceeds the **Max Rate**, the firewall blocks new connections from the offending IP address. The **Block Duration** specifies the amount of time the firewall continues to block the IP address's new connections. While the firewall blocks new connections, it doesn't count incoming connections and doesn't increment the threshold counters. For classified and aggregate profiles, use the default value (300 seconds) to block the attacking session without penalizing legitimate sessions from the source for too long a period of time.



*SYN Flood Protection is the only type for which you set the drop **Action**. Start by setting the **Action** to **SYN Cookies**. SYN Cookies treats legitimate traffic fairly and only drops traffic that fails the SYN handshake, while using Random Early Drop drops traffic randomly, so RED may affect legitimate traffic. However, SYN Cookies is more resource-intensive because the firewall acts as a proxy for the target server and handles the three-way handshake for the server. The tradeoff is not dropping legitimate traffic (SYN Cookies) versus preserving firewall resources (RED). Monitor the firewall, and if SYN Cookies consumes too many resources, switch to RED. If you don't have a dedicated DDoS prevention device in front of the firewall, always use RED as the drop mechanism.*

The default threshold values are high so that DoS Protection profiles don't unexpectedly drop legitimate traffic. Monitor connection traffic and adjust the thresholds to values appropriate for your network. Start by taking baseline measurements of average and peak CPS for each flood type to determine the normal traffic conditions for the critical devices you want to protect. Because normal traffic loads experience some fluctuation, it's best not to drop connections too aggressively. Monitor and adjust the flood thresholds as needed and as your network evolves.

Another method of setting flood thresholds is to use the baseline measurements to set the maximum CPS you want to allow and work back from there to derive reasonable flood mitigation alarm and activation rates.



*Firewalls with multiple dataplane processors (DPs) distribute connections across DPs. In general, the firewall divides the CPS threshold settings equally across its DPs. For example, if a firewall has five DPs and you set the **Alarm Rate** to 20,000 CPS, each DP has an **Alarm Rate** of 4,000 CPS (20,000 / 5 = 4,000), so if the new sessions on a DP exceeds 4,000, it triggers the **Alarm Rate** threshold for that DP.*

In addition to setting IP flood thresholds, you can also use DoS Protection profiles to detect and prevent session exhaustion attacks in which a large number of hosts (bots) establish as many sessions as possible to consume a target's resources. On the profile's **Resources Protection** tab, you can set the maximum number of concurrent sessions that the device(s) defined in the DoS Protection policy rule to which you apply the profile can receive. When the number of concurrent sessions reaches its maximum limit, new sessions are dropped.

The maximum number of concurrent sessions to set depends on your network context. Understand the number of concurrent sessions that the resources you are protecting (defined in

https://t.me/learningnets

the DoS Protection policy rule to which you attach the profile) can handle. Set the threshold to approximately 80% of the resources' capacity, then monitor and adjust the threshold as needed.

For aggregate profiles, the **Resources Protection** threshold applies to all traffic of the devices defined in the policy rule (source and destination). For classified profiles, the **Resources Protection** threshold applies to the traffic based on whether the classified policy rule applies to the source IP only, to the destination IP only, or to both the source and destination IPs.

DoS Protection Policy Rules

DoS Protection policy rules control the systems to which the firewall applies DoS protection (the flood thresholds configured in DoS Protection profiles that you attach to DoS Protection policy rules), what action to take when traffic matches the criteria defined in the rule, and how to log DoS traffic. Because DoS protection consumes firewall resources, use it only to defend specific critical resources against session floods, especially common targets that users access from the internet, such as web servers and database servers. Use Zone Protection profiles to protect entire zones against floods and other attacks. DoS Protection policy rules provide granular matching criteria so that you have the flexibility to define exactly what you want to protect:

- Source zone, interface, IP address (including whole regions), and user.
- Destination zone, interface, and IP address (including whole regions).
- Services (by port and protocol). DoS protection applies only to the services you specify. However, specifying services doesn't allow the services and implicitly block all other services. Specifying services limits DoS protection to those services, but doesn't block other services.



In addition to protecting service ports in use on critical servers, you can also protect against DoS attacks on the unused service ports of critical servers. For critical systems, you can do this by creating one DoS Protection policy rule and profile to protect ports with services running, and a different DoS Protection policy rule and profile to protect ports with no services running. For example, you can protect a web server's normal service ports, such as 80 and 443, with one policy/profile, and protect all of the other service ports with the other policy/profile. Be aware of the firewall's capacity so that servicing the DoS counters doesn't impact performance.

When traffic matches a DoS Protection policy rule, the firewall takes one of three actions:

- **Deny**—The firewall denies access and doesn't apply a DoS Protection profile. Traffic that matches the rule is blocked.
- **Allow**—The firewall permits access and doesn't apply a DoS Protection profile. Traffic that matches the rule is allowed.
- **Protect**—The firewall protects the devices defined in the DoS Protection policy rule by applying the specified DoS Protection profile or profiles thresholds to traffic that matches the rule. A rule can have one aggregate DoS Protection profile and one classified DoS Protection profile, and for classified profiles, you can use the source IP, destination IP, or both to increment the flood threshold counters, as described in Classified Versus Aggregate DoS Protection. Incoming packets count against both DoS Protection profile thresholds if the they match the rule.

The firewall applies DoS Protection profiles only if the **Action** is **Protect**. If the DoS Protection policy rule's **Action** is **Protect**, specify the appropriate aggregate and/or classified DoS Protection profiles in the rule so that the firewall applies the DoS Protection profile's thresholds to traffic that matches the rule. Most rules are **Protect** rules.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

The **Allow** and **Deny** actions enable you to make exceptions within larger groups but do not apply DoS protection to the traffic. For example, you can deny the traffic from most of a group but allow a subset of that traffic. Conversely, you can allow the traffic from most of a group and deny a subset of that traffic.

You can **Schedule** when a DoS Protection policy rule is active (start and end time, recurrence period). One use case for scheduling is to apply different flood thresholds at different times of the day or week. For example, if your business experiences significantly less traffic at night than during the day, you may want to apply higher flood thresholds during the day than at night. Another use case is to schedule special thresholds for special events, providing that the firewall supports the CPS rates.

For easier management and granular reporting, configure **Log Forwarding** to separate DoS protection logs from other threat logs. Forward DoS threshold violation events directly to the administrators via email in addition to forwarding the logs to a server such an SNMP or syslog server. Providing that the firewalls are appropriately sized, threshold breaches should not be frequent and will be strong indicators of an attack attempt.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Configure Zone Protection to Increase Network Security

The following topics provide zone protection configuration examples:

- Configure Reconnaissance Protection
- Configure Packet Based Attack Protection
- Configure Protocol Protection
- Configure Packet Buffer Protection
- Configure Packet Buffer Protection Based on Latency
- Configure Ethernet SGT Protection

Configure Reconnaissance Protection

Configure one of the following Reconnaissance Protection actions for the firewall to take in response to the corresponding reconnaissance attempt:

- **Allow**—The firewall allows the port scan or host sweep reconnaissance to continue.
- **Alert**—The firewall generates an alert for each port scan or host sweep that matches the configured threshold within the specified time interval. Alert is the default action.
- **Block**—The firewall drops all subsequent packets from the source to the destination for the remainder of the specified time interval.
- **Block IP**—The firewall drops all subsequent packets for the specified **Duration**, in seconds (the range is 1-3,600). **Track By** determines whether the firewall blocks source or source-and-destination traffic.

STEP 1 | Configure Reconnaissance Protection.

1. Select **Network** > **Network Profiles** > **Zone Protection**.
2. Select a Zone Protection profile or **Add** a new profile and enter a **Name** for it.
3. On the Reconnaissance Protection tab, select the scan types to protect against.
4. Select an **Action** for each scan. If you select Block IP, you must also configure **Track By** (source or source-and-destination) and **Duration**.
5. Set the **Interval** in seconds. This options defines the time interval for port scan and host sweep detection.
6. Set the **Threshold**. The threshold defines the number of port scan events or host sweeps that occurs within the interval configured above that triggers an action.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 2 | (Optional) Configure a Source Address Exclusion.

1. On the Reconnaissance Protection tab, **Add** a Source Address Exclusion.
   1. Enter a descriptive **Name** for the address you want to exclude.
   2. Set the Address Type to **IPv4** or **IPv6** and then select an address object or enter an IP address.
   3. Click **OK**.
2. Click **OK** to save the Zone Protection profile.
3. **Commit** your changes.

Configure Packet Based Attack Protection

To enhance security for a zone, Packet-Based Attack Protection allows you to specify whether the firewall drops IP, IPv6, TCP, ICMP, or ICMPv6 packets that have certain characteristics or strips certain options from the packets.

For example, you can drop TCP SYN and SYN-ACK packets that contain data in the payload during a TCP three-way handshake. A Zone Protection profile by default is set to drop SYN and SYN-ACK packets with data (you must apply the profile to the zone).

The TCP Fast Open option (RFC 7413) preserves the speed of a connection setup by including data in the payload of SYN and SYN-ACK packets. A Zone Protection profile treats handshakes that use the TCP Fast Open option separately from other SYN and SYN-ACK packets; the profile by default is set to allow the handshake packets if they contain a valid Fast Open cookie.

If you have existing Zone Protection profiles in place when you upgrade to PAN-OS 8.0, the three default settings will apply to each profile and the firewall will act accordingly.

Beginning with PAN-OS 8.1.2 and later releases, you can use a CLI command (Step 4 in this task) to enable the firewall to generate a Threat log when the firewall receives and drops the following types of packets, so that you can more easily analyze these occurrences and also fulfill audit and compliance requirements:

- Teardrop attack
- DoS attack using ping of death

Furthermore, the same CLI command also enables the firewall to generate Threat logs for the following types of packets if you enable the corresponding Packet Based Attack Protection:

- Fragmented IP packets
- IP address spoofing
- ICMP packets larger than 1024 bytes
- Packets containing ICMP fragments
- ICMP packets embedded with an error message
- First packets for a TCP session that are not SYN packets

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 1 | Create a Zone Protection profile and configure Packet-Based Attack Protection settings.

1. Select **Network** > **Network Profiles** > **Zone Protection** and **Add** a new profile.
2. Enter a **Name** for the profile and an optional **Description**.
3. Select **Packet Based Attack Protection**.
4. On each tab (**IP Drop**, **TCP Drop**, **ICMP Drop**, **IPv6 Drop**, and **ICMPv6 Drop**), select the Packet-Based Attack Protection settings you want to enforce to protect a zone.
5. Click **OK**.

STEP 2 | Apply the Zone Protection profile to a security zone that is assigned to interfaces you want to protect.

1. Select **Network** > **Zones** and select the zone where you want to assign the Zone Protection profile.
2. **Add** the **Interfaces** belonging to the zone.
3. For **Zone Protection Profile**, select the profile you just created.
4. Click **OK**.

STEP 3 | **Commit** your changes.

STEP 4 | (PAN-OS 8.1.2 and later releases) Enable the firewall to generate Threat logs for a teardrop attack and a DoS attack using ping of death, and also generate Threat logs for the types of packets listed above if you enable the corresponding packet-based attack protection (in Step 1). For example, if you enable packet-based attack protection for **Spoofed IP address**, using the following CLI causes the firewall to generate a Threat log when the firewall receives and drops a packet with a spoofed IP address.

1. Access the CLI.
2. Use the operational CLI command **`set system setting additional-threat-log on`**. Default is **`off`**.

Configure Protocol Protection

Protect virtual wire or Layer 2 security zones from non-IP protocol packets by using Protocol Protection.

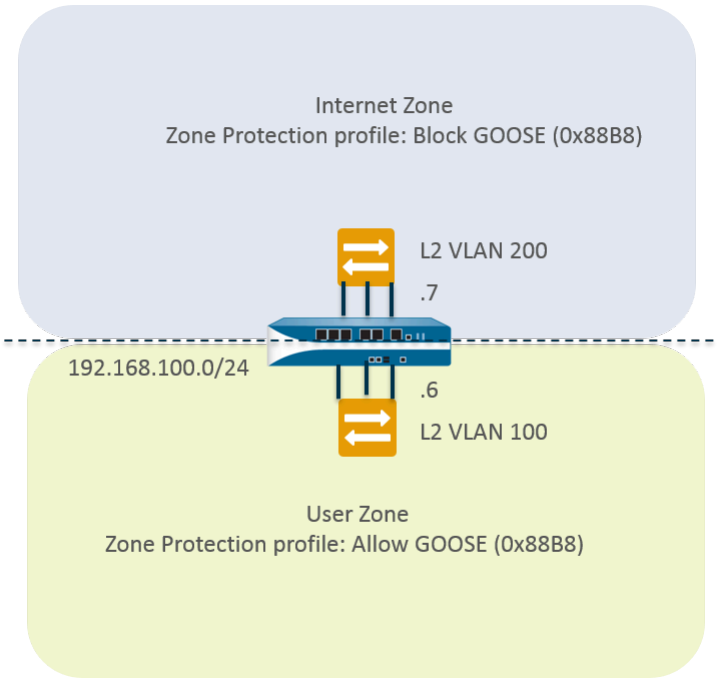
- Use Case: Non-IP Protocol Protection Between Security Zones on Layer 2 Interfaces
- Use Case: Non-IP Protocol Protection Within a Security Zone on Layer 2 Interfaces

Use Case: Non-IP Protocol Protection Between Security Zones on Layer 2 Interfaces

In this use case, the firewall is in a Layer 2 VLAN divided into two subinterfaces. VLAN 100 is 192.168.100.1/24, subinterface .6. VLAN 200 is 192.168.100.1/24, subinterface .7. Non-IP protocol protection applies to ingress zones. In this use case, if the Internet zone is the ingress zone, the firewall blocks the Generic Object Oriented Substation Event (GOOSE) protocol. If the User zone is the ingress zone, the firewall allows the GOOSE protocol. The firewall implicitly allows IPv4, IPv6, ARP, and VLAN-tagged frames in both zones.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets



STEP 1 | Configure two VLAN subinterfaces.

1. Select **Network** > **Interfaces** > **VLAN** and **Add** an interface.
2. **Interface Name** defaults to vlan. After the period, enter 7.
3. On the **Config** tab, **Assign Interface To** the **VLAN** 200.
4. Click **OK**.
5. Select **Network** > **Interfaces** > **VLAN** and **Add** an interface.
6. **Interface Name** defaults to vlan. After the period, enter 6.
7. On the **Config** tab, **Assign Interface To** the **VLAN** 100.
8. Click **OK**.

STEP 2 | Configure protocol protection in a Zone Protection profile to block GOOSE protocol packets.

1. Select **Network** > **Network Profiles** > **Zone Protection** and **Add** a profile.
2. Enter the **Name** Block GOOSE.
3. Select **Protocol Protection**.
4. Choose **Rule Type** of **Exclude List**.
5. Enter the **Protocol Name**, GOOSE, to easily identify the Ethertype on the list. The firewall doesn't verify that the name you enter matches the Ethertype code; it uses only the Ethertype code to filter.
6. Enter **Ethertype** code 0x88B8. The Ethertype must be preceded by 0x to indicate a hexadecimal value. Range is 0x0000 to 0xFFFF.
7. Select **Enable** to enforce the protocol protection. You can disable a protocol on the list, for example, for testing.
8. Click **OK**.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 3 | Apply the Zone Protection profile to the Internet zone.

1. Select **Network** > **Zones** and **Add** a zone.
2. Enter the **Name** of the zone, Internet.
3. For **Location**, select the virtual system where the zone applies.
4. For **Type**, select **Layer2**.
5. **Add** the **Interface** that belongs to the zone, vlan.7.
6. For **Zone Protection Profile**, select the profile Block GOOSE.
7. Click **OK**.

STEP 4 | Configure protocol protection to allow GOOSE protocol packets.

Create another Zone protection profile named Allow GOOSE, and choose **Rule Type** of **Include List**.

When configuring an Include list, include all required non-IP protocols; an incomplete list can result in legitimate non-IP traffic being blocked.

STEP 5 | Apply the Zone Protection profile to the User zone.

1. Select **Network** > **Zones** and **Add** a zone.
2. Enter the **Name** of the zone, User.
3. For **Location**, select the virtual system where the zone applies.
4. For **Type**, select **Layer2**.
5. **Add** the **Interface** that belongs to the zone, vlan.6.
6. For **Zone Protection Profile**, select the profile Allow GOOSE.
7. Click **OK**.

STEP 6 | Commit.

Click **Commit**.

STEP 7 | View the number of non-IP packets the firewall has dropped based on protocol protection.

Access the CLI.

```
> show counter global name pkt_nonip_pkt_drop
> show counter global name pkt_nonip_pkt_drop delta yes
```

Use Case: Non-IP Protocol Protection Within a Security Zone on Layer 2 Interfaces

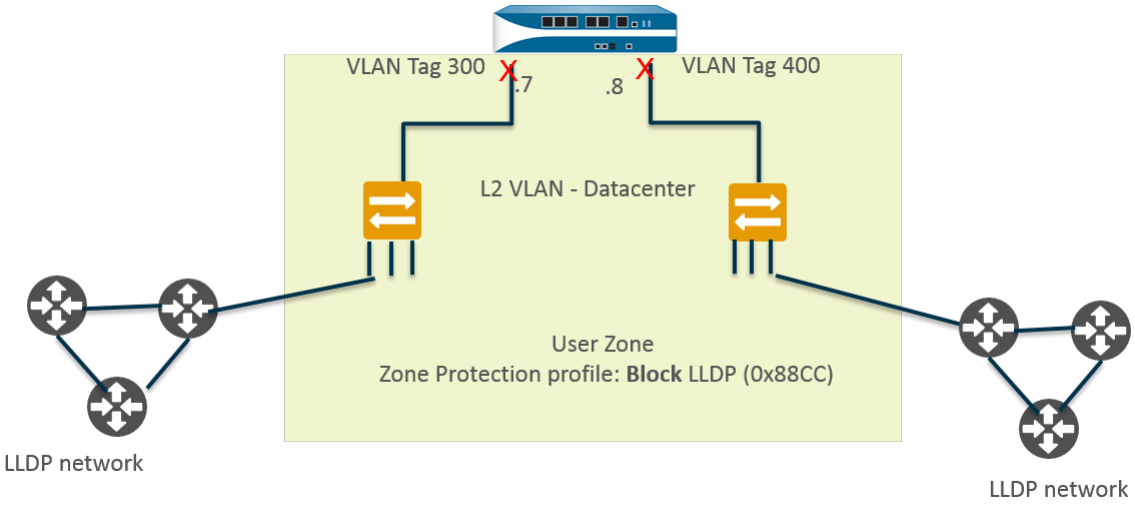
If you don't implement a Zone Protection profile with non-IP protocol protection, the firewall allows non-IP protocols in a single zone to go from one Layer 2 interface to another. In this use case, blocking LLDP packets ensures that LLDP for one network doesn't discover a network reachable through another interface in the zone.

In the following figure, the Layer 2 VLAN named Datacenter is divided into two subinterfaces: 192.168.1.1/24, subinterface .7 and 192.168.1.2/24, subinterface .8. The VLAN belongs to the User zone. By applying a Zone Protection profile that blocks LLDP to the User zone:

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

- Subinterface .7 blocks LLDP from its switch to the firewall at the red X on the left, preventing that traffic from reaching subinterface .8.
- Subinterface .8 blocks LLDP from its switch to the firewall at the red X on the right, preventing that traffic from reaching subinterface .7.



STEP 1 | Create a subinterface for an Ethernet interface.

1. Select **Network** > **Interfaces** > **Ethernet** and select a Layer 2 interface, in this example, ethernet1/1.
2. Select **Add Subinterfaces**.
3. The **Interface Name** defaults to the interface (ethernet 1/1). After the period, enter 7.
4. For **Tag**, enter 300.
5. For **Security Zone**, select User.
6. Click **OK**.

STEP 2 | Create a second subinterface for the Ethernet interface.

1. Select **Network** > **Interfaces** > **Ethernet** and select the Layer 2 interface: ethernet1/1.
2. Select **Add Subinterfaces**.
3. The **Interface Name** defaults to the interface (ethernet 1/1). After the period, enter 8.
4. For **Tag**, enter 400.
5. For **Security Zone**, select User.
6. Click **OK**.

STEP 3 | Create a VLAN for the Layer2 interface and two subinterfaces.

1. Select **Network** > **VLANs** and **Add** a VLAN.
2. Enter the **Name** of the VLAN; for this example, enter Datacenter.
3. For **VLAN Interface**, select **None**.
4. For **Interfaces**, click **Add** and select the Layer 2 interface: ethernet1/1, and two subinterfaces: ethernet1/1.7 and ethernet1/1.8.
5. Click **OK**.

https://t.me/learningnets

STEP 4 | Block non-IP protocol packets in a Zone Protection profile.

1. Select **Network** > **Network Profiles** > **Zone Protection** and **Add** a profile.
2. Enter the **Name**, in this example, Block LLDP.
3. Enter a profile **Description**—Block LLDP packets from an LLDP network to other interfaces in the zone (intrazone).
4. Select **Protocol Protection**.
5. Choose **Rule Type** of **Exclude List**.
6. Enter **Protocol Name** LLDP.
7. Enter **Ethertype** code 0x88cc. The Ethertype must be preceded by 0x to indicate a hexadecimal value.
8. Select **Enable**.
9. Click **OK**.

STEP 5 | Apply the Zone Protection profile to the security zone to which Layer 2 VLAN belongs.

1. Select **Network** > **Zones**.
2. **Add** a zone.
3. Enter the **Name** of the zone, User.
4. For **Location**, select the virtual system where the zone applies.
5. For **Type**, select **Layer2**.
6. **Add** an **Interface** that belongs to the zone, ethernet1/1.7
7. **Add** an **Interface** that belongs to the zone, ethernet1/1.8.
8. For **Zone Protection Profile**, select the profile Block LLDP.
9. Click **OK**.

STEP 6 | Commit.

Click **Commit**.

STEP 7 | View the number of non-IP packets the firewall has dropped based on protocol protection.

Access the CLI.

```
> show counter global name pkt_nonip_pkt_drop
> show counter global name pkt_nonip_pkt_drop delta yes
```

Configure Packet Buffer Protection

You can configure Packet Buffer Protection at two levels: the device level (global) and if enabled globally, you can also enable it at the zone level. Global packet buffer protection (**Device** > **Setup** > **Session**) is to protect firewall resources and ensure that malicious traffic does not cause the firewall to become non-responsive.

Packet buffer protection per ingress zone (**Network** > **Zones**) is a second layer of protection that starts blocking the offending IP address if it continues to exceed the packet buffer protection thresholds. The firewall can block all traffic from the offending source IP address. Keep in mind that if the source IP address is a translated NAT IP address, many users can be using the same

https://t.me/learningnets

IP address. If one abusive user triggers packet buffer protection and the ingress zone has packet buffer protection enabled, all traffic from that offending source IP address (even from non-abusive users) can be blocked when the firewall puts the IP address on its block list.

The most effective way to block DoS attacks against a service behind the firewall is to configure packet buffer protection globally and per ingress zone.

You can **Enable Packet Buffer Protection** for a zone, but it is not active until you enable packet buffer protection globally and specify the settings.

STEP 1 | Enable packet buffer protection globally.

1. Select **Device** > **Setup** > **Session** and edit the Session Settings.
2. Select **Packet Buffer Protection**.
3. Define the packet buffer protection behavior:
   - **Alert (%)**—When packet buffer utilization exceeds this threshold for more than 10 seconds, the firewall creates a log event every minute. Range s 0% to 99%; default is 50%. If the value is 0%, the firewall does not create a log event.
   - **Activate (%)**—When packet buffer utilization reaches this threshold, the firewall begins to mitigate the most abusive sessions by applying random early drop (RED). Range is 0% to 99%; default is 50%. If the value is 0%, the firewall does not apply RED. If the abuser is ingressing a zone that has Packet Buffer Protection enabled, the firewall can also discard the abusive session or block the offending source IP address. Start with the default threshold and adjust it if necessary.

     The firewall records alert events in the System log, and records events for dropped traffic, discarded sessions, and blocked IP address in the Threat log.
   - **Block Hold Time (sec)**—Number of seconds a RED-mitigated session is allowed to continue before the firewall discards it. Range is 0 to 65,535; default is 60. If the value is 0, the firewall does not discard sessions based on packet buffer protection.
   - **Block Duration (sec)**—Number of seconds a session remains discarded or an IP address remains blocked. Range is 1 to 15,999,999; default is 3,600.
4. Click **OK**.
5. **Commit** your changes.

STEP 2 | Enable additional packet buffer protection on an ingress zone.

1. Select **Network** > **Zones**.
2. Choose an ingress zone and click on its name.
3. **Enable Packet Buffer Protection** in the Zone Protection section.
4. Click **OK**.
5. **Commit** your changes.

Configure Packet Buffer Protection Based on Latency

Configure packet buffer protection based on latency and apply it to zones that have traffic consisting of protocols and applications that are latency-sensitive.

STEP 1 | Select **Device** > **Setup** > **Session**.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 2 | Edit the Session Settings section and enable **Packet Buffer Protection**.

STEP 3 | Enable **Buffering Latency Based**.

STEP 4 | Enter the **Latency Alert (milliseconds)** threshold above which the firewall starts generating an Alert log event every minute; range is 1 to 20,000; default is 50.

STEP 5 | Enter the **Latency Activate (milliseconds)** threshold above which the firewall activates random early drop (RED) on incoming packets and starts generating an Activate log every 10 seconds; range is 1 to 20,000ms; default is 200ms.

STEP 6 | Enter the **Latency Max Tolerate (milliseconds)** threshold above which the firewall uses RED with close to 100% drop probability; range is 1 to 20,000ms; default is 500ms.

If the current latency is a value between the **Latency Activate** threshold and the **Latency Max Tolerate** threshold, the firewall calculates the RED drop probability as follows: (current latency - **Latency Activate** threshold) / (**Latency Max Tolerate** threshold - **Latency Activate** threshold). For example, if the current latency is 300, **Latency Activate** is 200, and **Latency Max Tolerate** is 500, then (300-200)/(500-200) = 1/3, meaning the firewall uses approximately 33% RED drop probability.

STEP 7 | Configure the **Block Hold Time** and **Block Duration** as for Packet Buffer Protection based on utilization.

STEP 8 | Click **OK**.

STEP 9 | Enable the second layer of protection for each zone where you want packet buffer protection based on latency.

1. Select **Network** > **Zones** and select a zone.
2. Enable **Packet Buffer Protection**.

STEP 10 | **Commit**.

Configure Ethernet SGT Protection

Use the following task to configure an Ethernet SGT Protection profile.

STEP 1 | Create a Zone Protection profile to provide Ethernet SGT Protection.

1. Select **Network** > **Network Profiles** > **Zone Protection**.
2. **Add** a Zone Protection profile by **Name**.
3. Select **Ethernet SGT Protection**.
4. **Add** a **Layer 2 SGT Exclude List** by name.
5. Enter one or more **Tag** values for the list; range is 0 to 65,535. You can enter individual entries that are a contiguous range of tag values (for example, 100-500). You can add up to 100 (individual or range) tag entries in an Exclude List.
6. **Enable** the Layer 2 SGT Exclude List. You can disable the list at any time.
7. Click **OK**.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 2 | Apply the Zone Protection profile to the security zone to which the Layer 2, virtual wire, or tap interfaces belong.

1. Select **Network** > **Zones**.
2. **Add** a zone.
3. Enter the **Name** of the zone.
4. For **Location**, select the virtual system where the zone applies.
5. For **Type**, select **Layer2**, **Virtual Wire**, or **Tap**.
6. **Add** an **Interface** that belongs to the zone.
7. For **Zone Protection Profile**, select the profile you created.
8. Click **OK**.

STEP 3 | **Commit**.

STEP 4 | View the global counter of packets that the firewall dropped as a result of all Zone Protection profiles that employ Ethernet SGT Protection.

1. Access the CLI.
2. `> show counter global name pan_flow_dos_l2_sec_tag_drop`

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

DoS Protection Against Flooding of New Sessions

DoS protection against flooding of new sessions is beneficial against high-volume single-session and multiple-session attacks. In a single-session attack, an attacker uses a single session to target a device behind the firewall. If a Security rule allows the traffic, the session is established and the attacker initiates an attack by sending packets at a very high rate with the same source IP address and port number, destination IP address and port number, and protocol, trying to overwhelm the target. In a multiple-session attack, an attacker uses multiple sessions (or connections per second [cps]) from a single host to launch a DoS attack.

This feature defends against DoS attacks of new sessions only, that is, traffic that has not been offloaded to hardware. An offloaded attack is not protected by this feature. However, this topic describes how you can create a Security policy rule to reset the client; the attacker reinitiates the attack with numerous connections per second and is blocked by the defenses illustrated in this topic.

DoS Protection Profiles and Policy Rules work together to provide protection against flooding of many incoming SYN, UDP, ICMP, and ICMPv6 packets, and other types of IP packets. You determine what thresholds constitute flooding. In general, the DoS Protection profile sets the thresholds at which the firewall generates a DoS alarm, takes action such as Random Early Drop, and drops additional incoming connections. A DoS Protection policy rule configured to protect (rather than to allow or deny packets) determines the criteria for packets to match (such as source address) in order to be counted toward the thresholds. This flexibility allows you to block certain traffic, or allow certain traffic and treat other traffic as DoS traffic. When the incoming rate exceeds your maximum threshold, the firewall blocks incoming traffic from the source address.

- Multiple-Session DoS Attack
- Single-Session DoS Attack
- Configure DoS Protection Against Flooding of New Sessions
- End a Single Session DoS Attack
- Identify Sessions That Use Too Much of the On-Chip Packet Descriptor
- Discard a Session Without a Commit

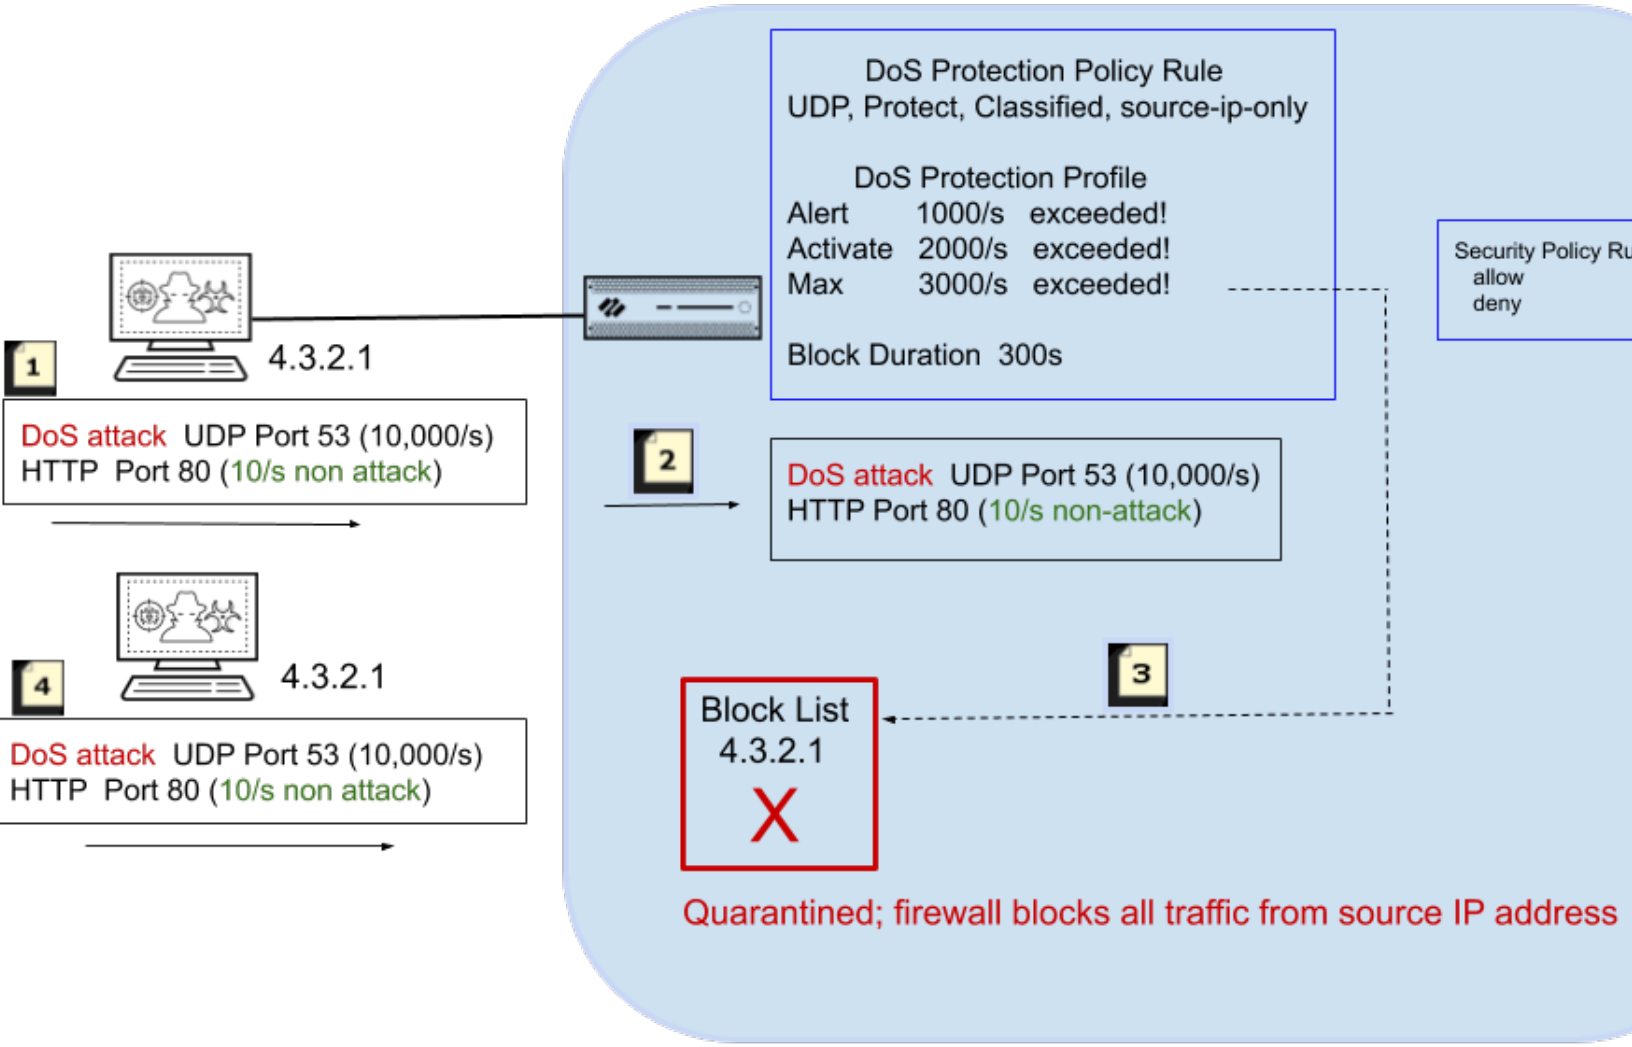
Multiple-Session DoS Attack

Configure DoS Protection Against Flooding of New Sessions by configuring a DoS Protection policy rule, which determines the criteria that, when matched by incoming packets, trigger the **Protect** action. The DoS Protection profile counts each new connection toward the Alarm Rate, Activate Rate, and Max Rate thresholds. When the incoming new connections per second exceed the Activate Rate, the firewall takes the action specified in the DoS Protection profile.

The following figure and table describe how the Security policy rules, DoS Protection policy rules and profile work together in an example.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets



Sequence of Events as Firewall Quarantines an IP Address	
1	In this example, an attacker launches a DoS attack at a rate of 10,000 new connections per second to UDP port 53. The attacker also sends 10 new connections per second to HTTP port 80.
2	The new connections match criteria in the DoS Protection policy rule, such as a source zone or interface, source IP address, destination zone or interface, destination IP address, or a service, among other settings. In this example, the policy rule specifies UDP. The DoS Protection policy rule also specifies the **Protect** action and **Classified**, two settings that dynamically put the DoS Protection profile settings into effect. The DoS Protection profile specifies that a Max Rate of 3000 packets per second is allowed. When incoming packets match the DoS Protection policy rule,

 ©2022 Palo Alto Networks, Inc.

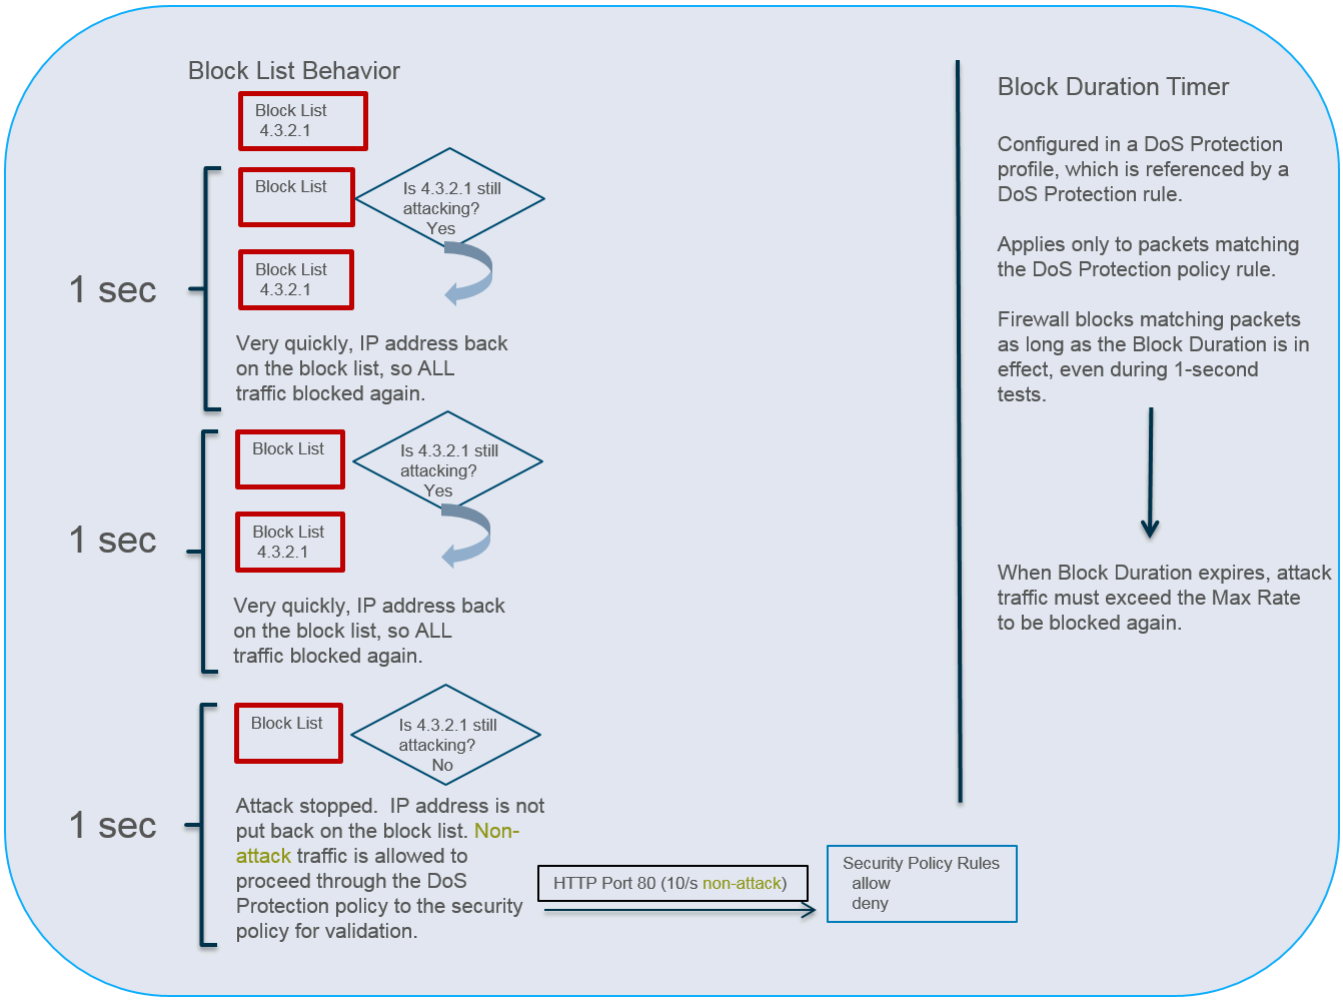
https://t.me/learningnets

Sequence of Events as Firewall Quarantines an IP Address	
	new connections per second are counted toward the **Alert**, **Activate**, and **Max Rate** thresholds. *You can also use a Security policy rule to block all traffic from the source IP address if you deem that address to be malicious all the time.*
3	The 10,000 new connections per second exceed the **Max Rate** threshold. When all of the following occur: • the threshold is exceeded, • a **Block Duration** is specified, and • **Classified** is set to include source IP address, the firewall puts the offending source IP address on the block list.
4	An IP address on the block list is in quarantine, meaning all traffic from that IP address is blocked. The firewall blocks the offending source IP address before additional attack packets reach the Security policy.

The following figure describes in more detail what happens after an IP address that matches the DoS Protection policy rule is put on the block list. It also describes the Block Duration timer.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets



Every one second, the firewall allows the IP address to come off the block list so that the firewall can test the traffic patterns and determine if the attack is ongoing. The firewall takes the following action:

- During this one-second test period, the firewall allows packets that don't match the DoS Protection policy criteria (HTTP traffic in this example) through the DoS Protection policy rules to the Security policy for validation. Very few packets, if any, have time to get through because the first attack packet that the firewall receives after the IP address is let off the block list will match the DoS Protection policy criteria, quickly causing the IP address to be placed back on the block list for another second. The firewall repeats this test each second until the attack stops.
- The firewall blocks all attack traffic from going past the DoS Protection policy rules (the address remains on the block list) until the Block Duration expires.

The 1-second checks illustrated in the preceding figure occur on firewall models that have multiple dataplane CPUs and a hardware network processor. All single dataplane systems or systems without a hardware network processor perform this mitigation in software and use a 5-second interval.

When the attack stops, the firewall does not put the IP address back on the block list. The firewall allows non-attack traffic to proceed through the DoS Protection policy rules to the Security policy

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

rules for evaluation. You must configure a Security policy rule to allow or deny traffic because without one, an implicit Deny rule denies all traffic.

The block list is based on a source zone and source address combination. This behavior allows duplicate IP addresses to exist as long as they are in different zones belonging to separate virtual routers.

The Block Duration setting in a DoS Protection profile specifies how long the firewall blocks the [offending] packets that match a DoS Protection policy rule. The attack traffic remains blocked until the Block Duration expires, after which the attack traffic must again exceed the Max Rate threshold to be blocked again.

If the attacker uses multiple sessions or bots that initiate multiple attack sessions, the sessions count toward the thresholds in the DoS Protection profile without a Security policy deny or drop rule in place. Hence, a single-session attack requires a Security policy deny or drop rule in order for each packet to count toward the thresholds; a multiple-session attack does not.

Therefore, the DoS protection against flooding of new sessions allows the firewall to efficiently defend against a source IP address while attack traffic is ongoing and to permit non-attack traffic to pass as soon as the attack stops. Putting the offending IP address on the block list allows the DoS protection functionality to take advantage of the block list, which is designed to quarantine all activity from that source IP address, such as packets with a different application. Quarantining the IP address from all activity protects against a modern attacker who attempts a rotating application attack, in which the attacker simply changes applications to start a new attack or uses a combination of different attacks in a hybrid DoS attack. You can Monitor Blocked IP Addresses to view the block list, remove entries from it, and get additional information about an IP address on the block list.

Beginning with PAN-OS 7.0.2, it is a change in behavior that the firewall places the attacking source IP address on the block list. When the attack stops, non-attack traffic is allowed to proceed to Security policy enforcement. The attack traffic that matched the DoS Protection profile and DoS Protection policy rules remains blocked until the Block Duration expires.

Single-Session DoS Attack

A single-session DoS attack typically will not trigger Zone or DoS Protection profiles because they are attacks that are formed after the session is created. These attacks are allowed by the Security policy because a session is allowed to be created, and after the session is created, the attack drives up the packet volume and takes down the target device.

Configure DoS Protection Against Flooding of New Sessions to protect against flooding of new sessions (single-session and multiple-session flooding). In the event of a single-session attack that is underway, additionally End a Single Session DoS Attack.

Configure DoS Protection Against Flooding of New Sessions

STEP 1 | Configure Security policy rules to deny traffic from the attacker's IP address and allow other traffic based on your network needs. You can specify any of the match criteria in a Security policy rule, such as source IP address. (Required for single-session attack mitigation or attacks

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

that have not triggered the DoS Protection policy threshold; optional for multiple-session attack mitigation).

This step is one of the steps typically performed to stop an existing attack. See End a Single Session DoS Attack.

- Create a Security Policy Rule

STEP 2 | Configure a DoS Protection profile for flood protection.

Because flood attacks can occur over multiple protocols, as a best practice, activate protection for all of the flood types in the DoS Protection profile.

1. Select **Objects** > **Security Profiles** > **DoS Protection** and **Add** a profile **Name**.
2. Select **Classified** as the **Type**.
3. For **Flood Protection**, select all types of flood protection:
   - **SYN Flood**
   - **UDP Flood**
   - **ICMP Flood**
   - **ICMPv6 Flood**
   - **Other IP Flood**
4. When you enable **SYN Flood**, select the **Action** that occurs when connections per second (cps) exceed the **Activate Rate** threshold:
   1. **Random Early Drop**—The firewall uses an algorithm to progressively start dropping that type of packet. If the attack continues, the higher the incoming cps rate (above the **Activate Rate**) gets, the more packets the firewall drops. The firewall drops packets until the incoming cps rate reaches the **Max Rate**, at which point the firewall drops all incoming connections. **Random Early Drop** (RED) is the default action for **SYN Flood**, and the only action for **UDP Flood**, **ICMP Flood**, **ICMPv6 Flood**, and **Other IP Flood**. RED is more efficient than SYN Cookies and can handles larger attacks, but doesn't discern between good and bad traffic.
   2. **SYN Cookies**—Rather than immediately sending the SYN to the server, the firewall generates a cookie (on behalf of the server) to send in the SYN-ACK to the client. The client responds with its ACK and the cookie; upon this validation the firewall then sends the SYN to the server. The **SYN Cookies** action requires more firewall resources than **Random Early Drop**; it's more discerning because it affects bad traffic.
5. (Optional) On each of the flood tabs, change the following thresholds to suit your environment:
   - **Alarm Rate (connections/s)**—Specify the threshold rate (cps) above which a DoS alarm is generated. (Range is 0-2,000,000; default is 10,000.)
   - **Activate Rate (connections/s)**—Specify the threshold rate (cps) above which a DoS response is activated. When the **Activate Rate** threshold is reached, **Random Early**

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Drop occurs. Range is 0-2,000,000; default is 10,000. (For SYN Flood, you can select the action that occurs.)

- **Max Rate (connections/s)**—Specify the threshold rate of incoming connections per second that the firewall allows. When the threshold is exceeded, new connections that arrive are dropped. (Range is 2-2,000,000; default is 40,000.)

The default threshold values in this step are only starting points and might not be appropriate for your network. You must analyze the behavior of your network to properly set initial threshold values.

6. On each of the flood tabs, specify the **Block Duration** (in seconds), which is the length of time the firewall blocks packets that match the DoS Protection policy rule that references this profile. Specify a value greater than zero. (Range is 1-21,600; default is 300.)

   *Set a low **Block Duration** value if you are concerned that packets you incorrectly identify as attack traffic will be blocked unnecessarily.*

   Set a high **Block Duration** value if you are more concerned about blocking volumetric attacks than you are about incorrectly blocking packets that aren't part of an attack.
7. Click **OK**.

STEP 3 | Configure a DoS Protection policy rule that specifies the criteria for matching the incoming traffic.

The firewall resources are finite, so you wouldn't want to classify using source address on an internet-facing zone because there can be an enormous number of unique IP addresses that match the DoS Protection policy rule. That would require many counters and the firewall would run out of tracking resources. Instead, define a DoS Protection policy rule that classifies using the destination address (of the server you are protecting).

1. Select **Policies** > **DoS Protection** and **Add** a **Name** on the **General** tab. The name is case-sensitive and can be a maximum of 31 characters, including letters, numbers, spaces, hyphens, and underscores.
2. On the **Source** tab, choose the **Type** to be a **Zone** or **Interface**, and then **Add** the zone(s) or interface(s). Choose zone or interface depending on your deployment and what you

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

want to protect. For example, if you have only one interface coming into the firewall, choose Interface.

3. (Optional) For **Source Address**, select **Any** for any incoming IP address to match the rule or **Add** an address object such as a geographical region.
4. (Optional) For **Source User**, select **any** or specify a user.
5. (Optional) Select **Negate** to match any sources except those you specify.
6. (Optional) On the **Destination** tab, choose the **Type** to be a **Zone** or **Interface**, and then **Add** the destination zone(s) or interface(s). For example, enter the security zone you want to protect.
7. (Optional) For **Destination Address**, select **Any** or enter the IP address of the device you want to protect.
8. (Optional) On the **Option/Protection** tab, **Add** a **Service**. Select a service or click **Service** and enter a **Name**. Select **TCP** or **UDP**. Enter a **Destination Port**. Not specifying a particular service allows the rule to match a flood of any protocol type without regard to an application-specific port.
9. On the **Option/Protection** tab, for **Action**, select **Protect**.
10. Select **Classified**.
11. For **Profile**, select the name of the **DoS Protection** profile you created.
12. For **Address**, select **source-ip-only** or **src-dest-ip-both**, which determines the type of IP address to which the rule applies. Choose the setting based on how you want the firewall to identify offending traffic:
    - Specify **source-ip-only** if you want the firewall to classify only on the source IP address. Because attackers often test the entire network for hosts to attack, **source-ip-only** is the typical setting for a wider examination.
    - Specify **src-dest-ip-both** if you want to protect against DoS attacks only on the server that has a specific destination address, and you also want to ensure that every source IP address won't surpass a specific cps threshold to that server.
13. Click **OK**.

STEP 4 | Commit.

Click **Commit**.

End a Single Session DoS Attack

To mitigate a single-session DoS attack, you would still Configure DoS Protection Against Flooding of New Sessions in advance. At some point after you configure the feature, a session might be established before you realize a DoS attack (from the IP address of that session) is underway. When you see a single-session DoS attack, perform the following task to end the session, so that subsequent connection attempts from that IP address trigger the DoS protection against flooding of new sessions.

STEP 1 | Identify the source IP address that is causing the attack.

For example, use the firewall Packet Capture feature with a destination filter to collect a sample of the traffic going to the destination IP address. Alternatively, use the ACC to filter on destination address to view the activity to the target host being attacked.

https://t.me/learningnets

STEP 2 | Create a DoS Protection policy rule that will block the attacker's IP address after the attack thresholds are exceeded.

STEP 3 | Create a Security policy rule to deny the source IP address and its attack traffic.

STEP 4 | End any existing attacks from the attacking source IP address by executing the **clear session all filter source *<ip-address>*** operational command.

Alternatively, if you know the session ID, you can execute the **clear session id *<value>*** command to end that session only.

*If you use the **clear session all filter source *<ip-address>*** command, all sessions matching the source IP address are discarded, which can include both good and bad sessions.*

After you end the existing attack session, any subsequent attempts to form an attack session are blocked by the Security policy. The DoS Protection policy counts all connection attempts toward the thresholds. When the Max Rate threshold is exceeded, the source IP address is blocked for the Block Duration, as described in Multiple-Session DoS Attack.

Identify Sessions That Use Too Much of the On-Chip Packet Descriptor

When a firewall exhibits signs of resource depletion, it might be experiencing an attack that is sending an overwhelming number of packets. In such events, the firewall starts buffering inbound packets. You can quickly identify the sessions that are using an excessive percentage of the on-chip packet descriptor and mitigate their impact by discarding them.

Perform the following task on any hardware-based firewall model (not a VM-Series firewall) to identify, for each slot and dataplane, the on-chip packet descriptor percentage used, the top five sessions using more than two percent of the on-chip packet descriptor, and the source IP addresses associated with those sessions. Having that information allows you to take appropriate action.

STEP 1 | View firewall resource usage, top sessions, and session details. Execute the following operational command in the CLI (sample output from the command follows):

```
admin@PA-7050> show running resource-monitor ingress-backlogs
-- SLOT:s1, DP:dp1 -- USAGE - ATOMIC: 92%  TOTAL: 93%
TOP SESSIONS:SESS-ID      PCT   GRP-ID   COUNT
6            92%   1        156                   7        1732
SESSION DETAILS SESS-
ID PROTO SZONESRC       SPORT  DST       DPORT  IGR-IF    EGR-
IF       APP
```

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

```
6    6     trust 192.168.2.35 55653  10.1.8.89 80  ethernet1/21
 ethernet1/22 undecided
```

The command displays a maximum of the top five sessions that each use 2% or more of the on-chip packet descriptor.

The sample output above indicates that Session 6 is using 92% of the on-chip packet descriptor with TCP packets (protocol 6) coming from source IP address 192.168.2.35.

- **SESS-ID**—Indicates the global session ID that is used in all other **`show session`** commands. The global session ID is unique within the firewall.
- **GRP-ID**—Indicates an internal stage of processing packets.
- **COUNT**—Indicates how many packets are in that GRP-ID for that session.
- **APP**—Indicates the App-ID extracted from the Session information, which can help you determine whether the traffic is legitimate. For example, if packets use a common TCP or UDP port but the CLI output indicates an APP of `undecided`, the packets are possibly attack traffic. The APP is `undecided` when Application IP Decoders cannot get enough information to determine the application. An APP of `unknown` indicates that Application IP Decoders cannot determine the application; a session of `unknown` APP that uses a high percentage of the on-chip packet descriptor is also suspicious.

To restrict the display output:

On a PA-7000 Series model only, you can limit output to a slot, a dataplane, or both. For example:

```
admin@PA-7050> show running resource-monitor ingress-backlogs slot
 s1
admin@PA-7050> show running resource-monitor ingress-backlogs slot
 s1 dp dp1
```

On PA-5200 Series and PA-7000 Series models only, you can limit output to a dataplane. For example:

```
admin@PA-5260> show running resource-monitor ingress-backlogs dp
 dp1
```

STEP 2 | Use the command output to determine whether the source at the source IP address using a high percentage of the on-chip packet descriptor is sending legitimate or attack traffic.

In the sample output above, a single-session attack is likely occurring. A single session (Session ID 6) is using 92% of the on-chip packet descriptor for Slot 1, DP 1, and the application at that point is `undecided`.

- If you determine a single user is sending an attack and the traffic is not offloaded, you can End a Single Session DoS Attack. At a minimum, you can Configure DoS Protection Against Flooding of New Sessions.
- On a hardware model that has a field-programmable gate array (FPGA), the firewall offloads traffic to the FPGA when possible to increase performance. If the traffic is offloaded to

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

hardware, clearing the session does not help because then it is the software that must handle the barrage of packets. You should instead Discard a Session Without a Commit.

To see whether a session is offloaded or not, use the **show session id <session-id>** operational command in the CLI as shown in the following example. The `layer7processing` value indicates `completed` for sessions offloaded or `enabled` for sessions not offloaded.

```
admin@PA-5060> show session id 68088184

Session        68088184

        c2s flow:
                source:      1.1.42.15 [trust]
                dst:         1.2.27.99
                proto:       6
                sport:       55993           dport:      6881
                state:       ACTIVE          type:       FLOW
                src user:    unknown
                dst user:    unknown
                offload:     Yes

        s2c flow:
                source:      1.2.27.99 [untrust]
                dst:         1.1.42.15
                proto:       6
                sport:       6881            dport:      55993
                state:       ACTIVE          type:       FLOW
                src user:    unknown
                dst user:    unknown
                offload:     Yes

        DP                                   : 2
        index(local):                        : 979320
        start time                           : Tue Oct 27 14:20:09 2015
        timeout                              : 1200 sec
        time to live                         : 1167 sec
        total byte count(c2s)                : 270
        total byte count(s2c)                : 270
        layer7 packet count(c2s)             : 3
        layer7 packet count(s2c)             : 3
        vsys                                 : vsys1
        application                          : bittorrent
        rule                                 : rule1
        session to be logged at end          : True
        session in session ager              : True
        session updated by HA peer           : False
        layer7 processing                    : completed
        URL filtering enabled                : False
        session via syn-cookies              : False
        session terminated on host           : False
        session traverses tunnel             : False
        captive portal session               : False
        ingress interface                    : ethernet1/21
        egress interface                     : ethernet1/22
        session QoS rule                     : N/A (class 4)
        tracker stage l7proc                 : ctd decoder bypass
        end-reason                           : unknown
```

If the **show session id <session-id>** command output shows information similar to the following, the output implies that the session has not yet been installed on the PAN-OS

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

firewall. One reason why this can occur is because the traffic is denied due to a configured Security policy rule.

```
> show session id xxxxxxxxxx
Session xxxxxxxxxx
Bad Key: c2s: 'c2s'
Bad Key: s2c: 's2c'
index(local): : yyyyyyy
```

Discard a Session Without a Commit

Perform this task to permanently discard a session, such as a session that is overloading the packet buffer or on-chip packet descriptor. No commit is required; the session is discarded immediately after executing the command. The commands apply to both offloaded and non-offloaded sessions.

STEP 1 | In the CLI, execute the following operational command on any hardware model:

```
admin@PA-7050> request session-discard [timeout <seconds>]
 [reason <reason-string>] id <session-id>
```

The default timeout is 3,600 seconds.

STEP 2 | Verify that sessions have been discarded.

```
admin@PA-7050> show session all filter state discard
```

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Certifications

The following topics describe how to configure Palo Alto Networks® firewalls and appliances to support the Common Criteria and the Federal Information Processing Standard 140-2 (FIPS 140-2), which are security certifications that ensure a standard set of security assurances and functionalities. These certifications are often required by civilian U.S. government agencies and government contractors.

For details about product certifications and third-party validation, refer to the Certifications page.

> Enable FIPS and Common Criteria Support
> FIPS-CC Security Functions
> Scrub Swap Memory on a Firewall or Appliances in FIPS-CC Mode

https://t.me/learningnets

Enable FIPS and Common Criteria Support

Use the following procedures to enable FIPS-CC mode on a software version that supports Common Criteria and the Federal Information Processing Standards 140-2 (FIPS 140-2). When you enable FIPS-CC mode, all FIPS and CC functionality is included.

FIPS-CC mode is supported on all Palo Alto Networks next-generation firewalls and appliances —including VM-Series firewalls. To enable FIPS-CC mode, first boot the firewall into the Maintenance Recovery Tool (MRT) and then change the operational mode from normal mode to FIPS-CC mode. The procedure to change the operational mode is the same for all firewalls and appliances but the procedure to access the MRT varies.



When you enable FIPS-CC mode, the firewall will reset to the factory default settings; all configuration will be removed.

- Access the Maintenance Recovery Tool (MRT)
- Change the Operational Mode to FIPS-CC Mode

Access the Maintenance Recovery Tool (MRT)

The Maintenance Recovery Tool (MRT) enables you to perform several tasks on Palo Alto Networks firewalls and appliances. For example, you can revert the firewall or appliance to factory default settings, revert PAN-OS or a content update to a previous version, run diagnostics on the file system, gather system information, and extract logs. Additionally, you can use the MRT to Change the Operational Mode to FIPS-CC Mode or from FIPS-CC mode to normal mode.

The following procedures describe how to access the Maintenance Recovery Tool (MRT) on various Palo Alto Networks products.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

- Access the MRT on hardware firewalls and appliances (such as PA-220 firewalls, PA-7000 Series firewalls, or M-Series appliances).
  1. Establish a serial console session to the firewall or appliance.
     1. Connect a serial cable from the serial port on your computer to the console port on the firewall or appliance.

        If your computer does not have a 9-pin serial port but does have a USB port, use a serial-to-USB converter to establish the connection. If the firewall has a micro USB console port, connect to the port using a standard Type-A USB to micro USB cable.
     2. Open terminal emulation software on your computer and set to 9600-8-N-1 and then connect to the appropriate COM port.

        On a Windows system, you can go to the Control Panel to view the COM port settings for Device and Printers to determine which COM port is assigned to the console.
     3. Log in using an administrator account. (The default username/password is admin/admin.)
  2. Enter the following CLI command and press **y** to confirm:

     ```
     debug system maintenance-mode
     ```
  3. After the firewall or appliance boots to the MRT welcome screen (in approximately 2 to 3 minutes), press Enter on **Continue** to access the MRT main menu.

     *You can also access the MRT by rebooting the firewall or appliance and entering **maint** at the maintenance mode prompt. A direct serial console connection is required.*

     After the firewall or appliance boots into the MRT, you can access the MRT remotely by establishing an SSH connection to the management (MGT) interface IP address. At the login prompt, enter **maint** as the username and the firewall or appliance serial number as the password.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

- Access the MRT on VM-Series firewalls deployed in a private cloud (such as on a VMware ESXi or KVM hypervisor).
  1. Establish an SSH session to the management IP address of the firewall and log in using an administrator account.
  2. Enter the following CLI command and press **y** to confirm:

     ```
     debug system maintenance-mode
     ```

     It will take approximately 2 to 3 minutes for the firewall to boot to the MRT. During this time, your SSH session will disconnect.
  3. After the firewall boots to the MRT welcome screen, log in based on the operational mode:
     - **Normal mode**—Establish an SSH session to the management IP address of the firewall and log in using **maint** as the username and the firewall or appliance serial number as the password.
     - **FIPS-CC mode**—Access the virtual machine management utility (such as the vSphere client) and connect to the virtual machine console.
  4. From the MRT welcome screen, press Enter on **Continue** to access the MRT main menu.

- Access the MRT on VM-Series firewalls deployed in the public cloud (such as AWS or Azure).
  1. Establish an SSH session to the management IP address of the firewall and log in using an administrator account.
  2. Enter the following CLI command and press **y** to confirm:

     ```
     debug system maintenance-mode
     ```

     It will take approximately 2 to 3 minutes for the firewall to boot to the MRT. During this time, your SSH session will disconnect.
  3. After the firewall boots to the MRT welcome screen, log in based on the virtual machine type:
     - **AWS**—Log in as **ec2-user** and select the SSH public key associated with the virtual machine when you deployed it.
     - **Azure**—Enter the credentials you created when you deployed the VM-Series firewall.
     - **GCP**—Log in as **gcp-user** and select the SSH public key associated with the virtual machine when you deployed it.
  4. From the MRT welcome screen, press Enter on **Continue** to access the MRT main menu.

Change the Operational Mode to FIPS-CC Mode

The following procedure describes how to change the operational mode of a Palo Alto Networks product from normal mode to FIPS-CC mode.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

When the appliance is in FIPS-CC mode, you will not be able to configure any settings via the console, including the management interface settings. Before enabling FIPS-CC mode, make sure that your network is set up to allow access to the management interface via SSH or the web interface. The management interface will default to a static address of 192.168.1.1 if using a PA-Series firewall or to an address retrieved via DHCP if it is a VM-Series firewall. The WildFire, virtual Panorama, and M-series Panorama appliances will default to a static address of 192.168.1.1.

Once FIPS-CC mode is enabled, all configurations and settings are erased. If an administrator has configurations or settings they would like to reuse after FIPS-CC mode is enabled, the administrator can save and export the configuration before changing to FIPS-CC mode. The configuration can then be imported once the operational mode change is complete. The imported configuration must be edited per the FIPS-CC Security Functions or else the import process will fail.

Keys, passwords, and other critical security parameters cannot be shared across modes.

If you change the operational mode of a firewall or Dedicated Log Collector managed by a Panorama management server to FIPS-CC mode, you must also change the operational mode of Panorama to FIPS-CC mode. This is required to secure password hashes for local admin passwords pushed from Panorama.

STEP 1 | (Public Cloud VM-Series firewalls or Public Cloud Panorama Virtual Appliances only) Create an SSH key and log in to the firewall or Panorama.

On some public cloud platforms, such as Microsoft Azure, you must have an SSH key to prevent an authentication failure after changing to FIPS-CC mode. Verify that you have deployed the firewall to authenticate using the SSH key. Although on Azure you can deploy the VM-Series firewall or Panorama and log in using a username and password, you will be unable to authenticate using the username and password after changing the operational mode to FIPS-CC. After resetting to FIPS-CC mode, you must use the SSH key to log in and can then configure a username and password that you can use for subsequently logging in to the firewall web interface.

STEP 2 | Connect to the firewall or appliance and Access the Maintenance Recovery Tool (MRT).

STEP 3 | Select **`Set FIPS-CC Mode`** from the menu.

STEP 4 | Select **`Enable FIPS-CC Mode`**. The mode change operation begins a full factory reset and a status indicator shows the progress. After the mode change is complete, the status shows `Success`.

All configurations and settings are erased and cannot be retrieved once the mode change is complete.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

STEP 5 | When prompted, select **Reboot**.

*If you change the operational mode on a VM-Series firewall deployed in the public cloud and you lose your SSH connection to the MRT before you are able to **Reboot**, you must wait 10-15 minutes for the mode change to complete, log back into the MRT, and then reboot the firewall to complete the operation. After resetting to FIPS-CC mode, on some virtual form factors (Panorama or VM-Series) you can only log in using the SSH key, and if you have not set up authentication using an SSH key, you can no longer log in to the firewall on reboot.*

After you switch to FIPS-CC mode, you see the following status: `FIPS-CC mode enabled successfully`.

In addition, the following changes are in effect:

- FIPS-CC displays at all times in the status bar at the bottom of the web interface.
- The default administrator login credentials change to admin/paloalto.

See FIPS-CC Security Functions for details on the security functions that are enforced in FIPS-CC mode.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

FIPS-CC Security Functions

When FIPS-CC mode is enabled, the following security functions are enforced on all firewalls and appliances:

- To log in, the browser must be TLS 1.1 (or later) compatible; on a WF-500 appliance, you manage the appliance only through the CLI and you must connect using an SSHv2-compatible client application.
- All passwords must be at least eight characters.
- You must ensure that **Failed Attempts** and **Lockout Time (min)** are greater than 0 in authentication settings. If an administrator reaches the **Failed Attempts** threshold, the administrator is locked out for the duration defined in the **Lockout Time (min)** field.
- You must ensure that the **Idle Timeout** is greater than 0 in authentication settings. If a login session is idle for more than the specified time, the administrator is automatically logged out.
- You can configure the **Absolute Session Length** to set the maximum length of time in minutes that a user can be logged in. The minimum length that can be set is 60 minutes. You will receive a session termination warning 5 minutes before timeout. This feature cannot be disabled in FIPS-CC mode and defaults at a session of 30 days.
- You can configure the **Max No. of Sessions** to set how many users can be concurrently logged in to the same administrator account.
- The firewall or appliance automatically determines the appropriate level of self-testing and enforces the appropriate level of strength in encryption algorithms and cipher suites.
- Unapproved FIPS-CC algorithms are not decrypted—they are ignored during decryption.
- MS-CHAPv2 is not compatible with FIPS-CC mode. It is recommended to use RADIUS with TLS.
- When configuring an IPSec VPN, the administrator must select a cipher suite option presented to them during the IPSec setup.
- (For Panorama and WildFire only) IPSec can be enabled on the management interface to protect protocols such as NTP, RADIUS, TACACS, and DNS.
- Self-generated and imported certificates must contain public keys that are either RSA 2,048 bits (or more) or ECDSA 256 bits (or more); you must also use a digest of SHA256 or greater.

  You cannot use a hardware security module (HSM) to store the private ECDSA keys used for SSL Forward Proxy or SSL Inbound Inspection.
- Telnet, TFTP, and HTTP management connections are not available.
- You must enable encryption for the HA1 control link. You must set automatic rekeying parameters; you must set the data parameter to a value no greater than 1000 MB (you cannot let it default) and you must set a time interval (you cannot leave it disabled).
- The serial console port in FIPS-CC mode functions as a limited status output port only; CLI access is not available.
- The serial console port on hardware and private-cloud VM-Series firewalls booted into the MRT provides interactive access to the MRT.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

- Interactive console access is not supported in the hypervisor environment private-cloud VM-Series firewalls booted into the MRT; you can access the MRT only using SSH.
- You must manually configure a new master key before the old master key expires; **Auto Renew Master Key** is not supported in FIPS-CC mode.

  If the master key expires, the firewall or Panorama automatically reboots in Maintenance mode. You must then Reset the Firewall to Factory Default Settings.
- Zero Touch Provisioning (ZTP) mode is disabled on the PA-5450 Firewall and the PA-400 Series Firewalls if FIPS-CC mode is enabled.
- (Panorama managed devices) Review the Panorama support of firewalls and Log Collectors when FIPS-CC is enabled.

Panorama	Firewall		Log Collector	
FIPS-CC Enabled	FIPS-CC Enabled	FIPS-CC Disabled	FIPS-CC Enabled	FIPS-CC Disabled
	Supported	Supported	Supported	Supported
FIPS-CC Disabled	Not Supported	Supported	Not Supported	Supported

©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

Scrub the Swap Memory on Firewalls or Appliances Running in FIPS-CC Mode

You should ensure that sensitive information is removed from the swap memory before you decommission a firewall or appliance (in FIPS-CC mode) or before you send it in for repair. Use this procedure to remove all cryptographic security parameter (CSP) information from swap partitions.

If you send a firewall that is managed by Panorama in for repair, see Before Starting RMA Firewall Replacement.

STEP 1 | Open an SSH management session to the firewall or appliance.

STEP 2 | Run the following operational command:

```
request [restart | shutdown] system with-swap-scrub [dod | nnsa]
```

For example, to shut down the firewall or appliance and perform a Department of Defense (DoD) scrub, run the following command:

```
request shutdown system with-swap-scrub dod
```

STEP 3 | Press **Y** at the warning prompt to start the scrub.

STEP 4 | Verify that the scrub completed successfully. View the **System** log and filter on the word `swap`. The **System** log indicates the scrub status for each swap partition (either one or two partitions depending on the model) and also displays a log entry that indicates the overall status of the scrub. If the scrub completed successfully on all swap partitions, the **System** log shows `Swap space scrub was successful`.

If the scrub failed on one or more swap partitions, the **System** log shows `Swap space scrub was unsuccessful`. The following screen capture shows the log results for a firewall that has two partitions.

06/08 10:24:02	general	medium	general		Swap space scrub was successful
06/08 10:24:02	general	medium	general		Scrub performed on swap space /opt/panlogs /.secondary_swapfile
06/08 10:24:02	general	medium	general		Scrub performed on swap space /dev/sda7

To view the scrub logs using the CLI, run the ***show log system | match swap*** *command.*

If you initiate the scrub using the shutdown command, the firewall or appliance will power off after the scrub completes. Before you can power on the firewall or appliance, you must first disconnect and reconnect the power source.

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets

 ©2022 Palo Alto Networks, Inc.

https://t.me/learningnets