



User Guide for Cisco Security MARS Local Controller

Release 4.2.x
June 2006

Corporate Headquarters

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
<http://www.cisco.com>
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

Customer Order Number:
Text Part Number: 78-17020-01



THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The following information is for FCC compliance of Class A devices: This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to part 15 of the FCC rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio-frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case users will be required to correct the interference at their own expense.

The following information is for FCC compliance of Class B devices: The equipment described in this manual generates and may radiate radio-frequency energy. If it is not installed in accordance with Cisco's installation instructions, it may cause interference with radio and television reception. This equipment has been tested and found to comply with the limits for a Class B digital device in accordance with the specifications in part 15 of the FCC rules. These specifications are designed to provide reasonable protection against such interference in a residential installation. However, there is no guarantee that interference will not occur in a particular installation.

Modifying the equipment without Cisco's written authorization may result in the equipment no longer complying with FCC requirements for Class A or Class B digital devices. In that event, your right to use the equipment may be limited by FCC regulations, and you may be required to correct any interference to radio or television communications at your own expense.

You can determine whether your equipment is causing interference by turning it off. If the interference stops, it was probably caused by the Cisco equipment or one of its peripheral devices. If the equipment causes interference to radio or television reception, try to correct the interference by using one or more of the following measures:

- Turn the television or radio antenna until the interference stops.
- Move the equipment to one side or the other of the television or radio.
- Move the equipment farther away from the television or radio.
- Plug the equipment into an outlet that is on a different circuit from the television or radio. (That is, make certain the equipment and the television or radio are on circuits controlled by different circuit breakers or fuses.)

Modifications to this product not authorized by Cisco Systems, Inc. could void the FCC approval and negate your authority to operate the product.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

CCSP, CCVP, the Cisco Square Bridge logo, Follow Me Browsing, and StackWise are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn, and iQuick Study are service marks of Cisco Systems, Inc.; and Access Registrar, Aironet, BPX, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Enterprise/Solver, EtherChannel, EtherFast, EtherSwitch, Fast Step, FormShare, GigaDrive, GigaStack, HomeLink, Internet Quotient, IOS, IP/TV, iQ Expertise, the iQ logo, iQ Net Readiness Scorecard, LightStream, Linksys, MeetingPlace, MGX, the Networkers logo, Networking Academy, Network Registrar, *Packet*, PIX, Post-Routing, Pre-Routing, ProConnect, RateMUX, ScriptShare, SlideCast, SMARTnet, The Fastest Way to Increase Your Internet Quotient, and TransPath are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or Website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0601R)

User Guide for Cisco Security MARS Local Controller
Copyright © 2006 Cisco Systems, Inc. All rights reserved.



Preface **xix**

Introduction	xix
The MARS Appliance	xix
The MARS Web Interface	xix
About This Manual	xx
Obtaining Documentation	xxi
Cisco.com	xxi
Documentation DVD	xxi
Ordering Documentation	xxii
Documentation Feedback	xxii
Cisco Product Security Overview	xxii
Reporting Security Problems in Cisco Products	xxiii
Obtaining Technical Assistance	xxiii
Cisco Technical Support Website	xxiii
Submitting a Service Request	xxiv
Definitions of Service Request Severity	xxiv
Obtaining Additional Publications and Information	xxv

CHAPTER 1

STM Task Flow Overview **1-1**

Checklist for Provisioning Phase	1-2
Checklist for Monitoring Phase	1-9
Strategies for Monitoring, Notification, Mitigation, Remediation, and Audit	1-16
Appliance-side Tuning Guidelines	1-17
Device Inventory Worksheet	1-18
User Role Worksheet	1-20

CHAPTER 2

Reporting and Mitigation Devices Overview **2-1**

Levels of Operation	2-1
Selecting the Devices to Monitor	2-2
Understanding Access IP, Reporting IP, and Interface Settings	2-8
Access IP	2-9
Reporting IP	2-9
Interface Settings	2-10

Selecting the Access Type	2-10
Configure SNMP Access for Devices in MARS	2-11
Configure Telnet Access for Devices in MARS	2-11
Configure SSH Access for Devices in MARS	2-12
Configure FTP Access for Devices in MARS	2-12
Bootstrap Summary Table	2-12
Adding Reporting and Mitigation Devices	2-16
Add Reporting and Mitigation Devices Individually	2-17
Edit a Device	2-18
Upgrade the Device Type to a Newer Version	2-18
Delete a Device	2-19
Delete All Displayed Reporting Devices	2-20
Add Multiple Reporting and Mitigation Devices Using a Seed File	2-20
Devices that Require Custom Seed Files	2-21
Devices that Require Updates After the Seed File Import	2-21
Seed File Header Columns	2-21
Load Devices From the Seed File	2-24
Adding Reporting and Mitigation Devices Using Automatic Topology Discovery	2-25
Verify Connectivity with the Reporting and Mitigation Devices	2-26
Discover and Testing Connectivity Options	2-26
Run a Reporting Device Query	2-27
Activate the Reporting and Mitigation Devices	2-27
Data Enabling Features	2-28
Layer 2 Discovery and Mitigation	2-29
Networks for Dynamic Vulnerability Scanning	2-29
Select a Network for Scanning	2-30
Create a Network IP Address for Scanning	2-30
Create a Network IP Range for Scanning	2-30
Understanding NetFlow Anomaly Detection	2-30
How MARS Uses NetFlow Data	2-31
Guidelines for Configuring NetFlow on Your Network	2-32
Enable Cisco IOS Routers and Switches to Send NetFlow to MARS	2-32
Configuring Cisco CatIOS Switch	2-34
Enable NetFlow Processing in MARS	2-34
Host and Device Identification and Detail Strategies	2-36
Configuring Layer 3 Topology Discovery	2-36
Add a Community String for a Network	2-37
Add a Community String for an IP Range	2-37
Add Valid Networks to Discovery List	2-38

Remove Networks from Discovery List	2-38
Discover Layer 3 Data On Demand	2-38
Scheduling Topology Updates	2-39
Schedule a Network Discovery	2-39
To edit a scheduled topology discovery	2-40
To delete a scheduled topology discovery	2-40
To run a topology discovery on demand	2-41
Configuring Resource Usage Data	2-41
Configuring Network Admission Control Features	2-42
Integrating MARS with 3 rd -Party Applications	2-43
Forwarding Alert Data to 3 rd -Party Syslog and SNMP Servers	2-43
MARS MIB Format	2-43
Relaying Syslog Messages from 3rd-Party Syslog Servers	2-44
Configure Syslog-ng Server to Forward Events to MARS	2-44
Configure Kiwi Syslog Server to Forward Events to MARS	2-45
Add Syslog Relay Server to MARS	2-45
Add Devices Monitored by Syslog Relay Server	2-46

CHAPTER 3

Configuring Router and Switch Devices 3-1

Cisco Router Devices	3-1
Enable Administrative Access to Devices Running Cisco IOS 12.2	3-1
Enable SNMP Administrative Access	3-2
Enable Telnet Administrative Access	3-2
Enable SSH Administrative Access	3-2
Enable FTP-based Administrative Access	3-2
Configure the Device Running Cisco IOS 12.2 to Generate Required Data	3-3
Enable Syslog Messages	3-3
Enable SNMP RO Strings	3-3
Enable NAC-specific Messages	3-4
Enable SDEE for IOS IPS Software	3-6
Add and Configure a Cisco Router in MARS	3-6
Cisco Switch Devices	3-9
Enable Communications Between Devices Running CatOS and MARS	3-9
Enable SNMP Administrative Access	3-10
Enable Telnet Administrative Access	3-10
Enable SSH Administrative Access	3-10
Enable FTP-based Administrative Access	3-10
Configure the Device Running CatOS to Generate Required Data	3-11
Enable SNMP RO Strings on CatOS	3-11

Enable Syslog Messages on CatOS	3-11
Enable L2 Discovery Messages	3-12
Add and Configure a Cisco Switch in MARS	3-13
Adding Modules to a Cisco Switch	3-14
Add Available Modules	3-14
Add Cisco IOS 12.2 Modules Manually	3-15
Extreme ExtremeWare 6.x	3-17
Configure ExtremeWare to Generate the Required Data	3-17
Add and Configure an ExtremeWare Switch in MARS	3-18
Generic Router Device	3-18
Add and Configure a Generic Router in MARS	3-19

CHAPTER 4

Configuring Firewall Devices 4-1

Cisco Firewall Devices (PIX, ASA, and FWSM)	4-1
Bootstrap the Cisco Firewall Device	4-2
Enable Telnet Access on a Cisco Firewall Device	4-4
Enable SSH Access on a Cisco Firewall Device	4-4
Send Syslog Files From Cisco Firewall Device to MARS	4-4
Add and Configure a Cisco Firewall Device in MARS	4-5
Add Security Contexts Manually	4-8
Add Discovered Contexts	4-10
Edit Discovered Security Contexts	4-11
NetScreen ScreenOS Devices	4-11
Bootstrap the NetScreen Device	4-12
Add the NetScreen Device to MARS	4-17
Check Point Devices	4-19
Determine Devices to Monitor and Restrictions	4-21
Bootstrap the Check Point Devices	4-22
Add the MARS Appliance as a Host in Check Point	4-23
Define an OPSEC Application that Represents MARS	4-24
Obtain the Server Entity SIC Name	4-27
Select the Access Type for LEA and CPMI Traffic	4-29
Create and Install Policies	4-31
Verify Communication Path Between MARS Appliance and Check Point Devices	4-32
Reset the OPSEC Application Certificate of the MARS Appliance	4-33
Add and Configure Check Point Devices in MARS	4-36
Add a Check Point Primary Management Station to MARS	4-37
Manually Add a Child Enforcement Module or Log Server to a Check Point Primary Management Station	4-41

Add a Check Point Certificate Server	4-44
Edit Discovered Log Servers on a Check Point Primary Management Station	4-45
Edit Discovered Firewall on a Check Point Primary Management Station	4-47
Define Route Information for Check Point Firewall Modules	4-47
Specify Log Info Settings for a Child Enforcement Module or Log Server	4-49
Verify Connectivity Between MARS and Check Point Devices	4-52
Remove a Firewall or Log Server from a Check Point Primary Management Station	4-52
Troubleshooting MARS and Check Point	4-53

CHAPTER 5**Configuring VPN Devices 5-1**

Cisco VPN 3000 Concentrator	5-1
Bootstrap the VPN 3000 Concentrator	5-1
Add the VPN 3000 Concentrator to MARS	5-2

CHAPTER 6**Configuring Network-based IDS and IPS Devices 6-1**

Cisco IDS 3.1 Sensors	6-1
Configure Sensors Running IDS 3.1	6-1
Add and Configure a Cisco IDS 3.1 Device in MARS	6-4
Cisco IDS 4.0 and IPS 5.x Sensors	6-5
Bootstrap the Sensor	6-5
Enable the Access Protocol on the Sensor	6-6
Enable the Correct Signatures and Actions	6-6
Add and Configure a Cisco IDS or IPS Device in MARS	6-6
Specify the Monitored Networks for Cisco IPS or IDS Device Imported from a Seed File	6-8
View Detailed Event Data for Cisco IPS Devices	6-9
Cisco IPS Modules	6-9
Enable DTM Support	6-10
Enable SDEE on the Cisco IOS Device with an IPS Module	6-10
Add an IPS Module to a Cisco Switch or Cisco ASA	6-11
ISS Site Protector	6-13
ISS RealSecure 6.5 and 7.0	6-17
Configure ISS RealSecure to Send SNMP Traps to MARS	6-18
Add an ISS RealSecure Device as a NIDS	6-19
Add an ISS RealSecure Device as a HIDS	6-20
IntruVert IntruShield	6-22
Extracting Intruvert Sensor Information from the IntruShield Manager	6-22
Configure IntruShield Version 1.5 to Send SNMP traps to MARS	6-23
Configure IntruShield Version 1.8 to Send SNMP Traps to MARS	6-23
Add and Configure an IntruShield Manager and its Sensors in MARS	6-25

Add the IntruShield Manager Host to MARS	6-26
Add IntruShield Sensors Manually	6-26
Add IntruShield Sensors Using a Seed File	6-27
Snort 2.0	6-28
Configure Snort to Send Syslogs to MARS	6-28
Add the Snort Device to MARS	6-28
Symantec ManHunt	6-29
Symantec ManHunt Side Configuration	6-29
MARS Side Configuration	6-30
Add Configuration Information for Symantec ManHunt 3.x	6-30
NetScreen IDP 2.1	6-31
IDP-side Configuration	6-31
MARS-side Configuration	6-31
Add Configuration Information for the IDP	6-31
Add NetScreen IDP 2.1 Sensors Manually	6-32
Enterasys Dragon 6.x	6-33
DPM/EFP Configuration	6-33
Configure the DPM or EFP	6-33
Host-side Configuration	6-34
Configure the syslog on the UNIX host	6-34
MARS-side Configuration	6-34
Add Configuration Information for the Enterasys Dragon	6-34
Add a Dragon NIDS Device	6-34

CHAPTER 7
Configuring Host-Based IDS and IPS Devices 7-1

Entercept Intercept 2.5 and 4.0	7-1
Extracting Entercept Agent Information into a CSV file (for Entercept Version 2.5)	7-1
Create a CSV file for Entercept Agents in Version 2.5	7-2
Define the MARS Appliance as an SNMP Trap Target	7-2
Specify the Events to Generate SNMP Traps for MARS	7-2
Add and Configure an Entercept Console and its Agents in MARS	7-3
Add the Entercept Console Host to MARS	7-3
Add Entercept Agents Manually	7-4
Add Entercept Agents Using a Seed File	7-4
Cisco Security Agent 4.x Device	7-5
Configure CSA Management Center to Generate Required Data	7-5
Configure CSA MC to Forward SNMP Notifications to MARS	7-6
Export CSA Agent Information to File	7-6
Add and Configure a CSA MC Device in MARS	7-7

Add a CSA Agent Manually	7-8
Add CSA Agents From File	7-9
Troubleshooting CSA Agent Installs	7-10

CHAPTER 8**Configuring Antivirus Devices 8-1**

Symantec AntiVirus Configuration	8-1
Configure the AV Server to Publish Events to MARS Appliance	8-1
Export the AntiVirus Agent List	8-7
Add the Device to MARS	8-7
Add Agent Manually	8-7
Add Agents from a CSV File	8-8
McAfee ePolicy Orchestrator Devices	8-8
Configure ePolicy Orchestrator to Generate Required Data	8-8
Add and Configure ePolicy Orchestrator Server in MARS	8-12
Cisco Incident Control Server	8-13
Configure Cisco ICS to Send Syslogs to MARS	8-14
Add the Cisco ICS Device to MARS	8-15
Define Rules and Reports for Cisco ICS Events	8-15

CHAPTER 9**Configuring Vulnerability Assessment Devices 9-1**

Foundstone FoundScan 3.0	9-1
Configure FoundScan to Generate Required Data	9-1
Add and Configure a FoundScan Device in MARS	9-2
eEye REM 1.0	9-3
Configure eEye REM to Generate Required Data	9-3
Add and Configure the eEye REM Device in MARS	9-4
Qualys QualysGuard Devices	9-5
Configure QualysGuard to Scan the Network	9-6
Add and Configure a QualysGuard Device in MARS	9-6
Schedule the Interval at Which Data is Pulled	9-8
Troubleshooting QualysGuard Integration	9-8

CHAPTER 10**Configuring Generic, Solaris, Linux, and Windows Application Hosts 10-1**

Adding Generic Devices	10-1
Sun Solaris and Linux Hosts	10-2
Configure the Solaris or Linux Host to Generate Events	10-2
Configure Syslogd to Publish to the MARS Appliance	10-2
Configure MARS to Receive the Solaris or Linux Host Logs	10-3

Microsoft Windows Hosts	10-4
Push Method: Configure Generic Microsoft Windows Hosts	10-5
Install the SNARE Agent on the Microsoft Windows Host	10-5
Enable SNARE on the Microsoft Windows Host	10-6
Pull Method: Configure the Microsoft Windows Host	10-6
Enable Windows Pulling Using a Domain User	10-7
Enable Windows Pulling from Windows NT	10-7
Enable Windows Pulling from a Windows 2000 Server	10-7
Enable Windows Pulling from a Windows Server 2003 or Windows XP Host	10-8
Configure the MARS to Pull or Receive Windows Host Logs	10-8
Windows Event Log Pulling Time Interval	10-10
Define Vulnerability Assessment Information	10-11
Identify Network Services Running on the Host	10-13

CHAPTER 11

Configuring Database Applications 11-1

Oracle Database Server Generic	11-1
Configure the Oracle Database Server to Generate Audit Logs	11-1
Add the Oracle Database Server to MARS	11-2
Configure Interval for Pulling Oracle Event Logs	11-3

CHAPTER 12

Configuring Web Server Devices 12-1

Microsoft Internet Information Sever	12-1
Install and Configure the Snare Agent for IIS	12-1
To configure IIS for web logging	12-2
MARS-side Configuration	12-5
To add configuration information for the host	12-5
Apache Web Server on Solaris or RedHat Linux	12-7
Sun Java System Web Server on Solaris	12-7
Generic Web Server Generic	12-7
Solaris or Linux-side Configuration	12-7
Install and Configure the Web Agent on UNIX or Linux	12-7
Web Server Configuration	12-8
To configure the Apache web server for the agent	12-8
To configure the iPlanet web server for the agent	12-8
MARS-side Configuration	12-9
To add configuration information for the host	12-9

CHAPTER 13

Configuring Web Proxy Devices 13-1

Network Appliance NetCache Generic	13-1

Configure NetCache to Send Syslog to MARS 13-1

Add and Configure NetCache in MARS 13-2

CHAPTER 14

Configuring AAA Devices 14-1

Supporting Cisco Secure ACS Server 14-2

Supporting Cisco Secure ACS Solution Engine 14-2

Bootstrap Cisco Secure ACS 14-2

Configure Cisco Secure ACS to Generate Logs 14-3

Define AAA Clients 14-5

Configure TACACS+ Command Authorization for Cisco Routers and Switches 14-6

Install and Configure the PN Log Agent 14-7

Upgrade PN Log Agent to a Newer Version 14-9

Application Log Messages for the PN Log Agent 14-10

Add and Configure the Cisco ACS Device in MARS 14-12

CHAPTER 15

Configuring Custom Devices 15-1

Adding User Defined Log Parser Templates 15-1

To add a custom Device/Application type: 15-1

To add Parser Templates for a Device/Application 15-3

CHAPTER 16

Policy Table Lookup on Cisco Security Manager 16-1

Overview of Cisco Security Manager Policy Table Lookup 16-1

More About Cisco Security Manager Device Lookup 16-3

More About Cisco Security Manager Policy Table Lookup 16-4

Prerequisites for Policy Table Lookup 16-4

Restrictions for Policy Table Lookup 16-5

Checklist for Security Manager-to-MARS Integration 16-6

Bootstrapping Cisco Security Manager Server to Communicate with MARS 16-12

Add a Cisco Security Manager Server to MARS 16-13

Procedure for Invoking Cisco Security Manager Policy Table Lookup from Cisco Security MARS 16-14

CHAPTER 17

Network Summary 17-1

Navigation within the MARS Appliance 17-1

Logging In 17-1

Basic Navigation 17-2

Help Page 17-4

Your Suggestions Welcomed 17-4

Summary Page 17-6

Dashboard	17-6
Recent Incidents	17-8
Sessions and Events	17-8
Data Reduction	17-9
Page Refresh	17-9
Diagrams	17-9
Manipulating the Diagrams	17-11
Display Devices in Topology	17-12
Network Status	17-12
Reading Charts	17-13
My Reports	17-15
To set up reports for viewing	17-15

CHAPTER 18

Case Management 18-1

Case Management Overview	18-1
Case Management Considerations for the Global Controller	18-3
Hide and Display the Case Bar	18-3
Create a New Case	18-4
Edit and Change the Current Case	18-5
Add Data to a Case	18-6
Generate and Email a Case Report	18-7

CHAPTER 19

Incident Investigation and Mitigation 19-1

Incidents Overview	19-1
The Incidents Page	19-2
Time ranges for Incidents	19-4
Incident Details Page	19-4
To Search for a Session ID or Incident ID	19-4
Incident Details Table	19-5
False Positive Confirmation	19-6
The False Positive Page	19-8
To Tune a False Positive	19-9
To Tune an Unconfirmed False Positive to False Positive	19-9
To Tune an Unconfirmed False Positive to True Positive	19-9
To Activate False Positive Drop Rules	19-10
Mitigation	19-10
802.1X Mitigation Example	19-11
Prerequisites for Mitigation with 802.1X Network Mapping	19-11

Procedure for Mitigation with 802.1X Network Mapping	19-11
Display Dynamic Device Information	19-15
Virtual Private Network Considerations	19-17
Layer 2 Path and Mitigation Configuration Example	19-17
Prerequisites for Layer 2 Path and Mitigation	19-17
Components Used	19-17
Network Diagram	19-18
Procedures for Layer 2 Path and Mitigation	19-19
Add the Cisco Catalyst 5000 with SNMP as the Access Type.	19-19
Add the Cisco Catalyst 6500 with SNMP as Access Type (Layer 2 only).	19-20
Add the Cisco 7500 Router with TELNET as the Access Type	19-21
Verify the Connectivity Paths for Layer 3 and Layer 2	19-22
Perform Mitigation	19-26

CHAPTER 20

Queries and Reports 20-1

Queries	20-1
To Run a Quick Query	20-2
To Run a Free-form Query	20-2
To Run a Batch Query	20-3
To Stop a Batch Query	20-4
To Resubmit a Batch Query	20-4
To Delete a Batch Query	20-5
Selecting the Query Type	20-5
Result Format	20-5
Order/Rank By	20-7
Filter By Time	20-8
Use Only Firing Events	20-8
Maximum Number of Rows Returned	20-8
Selecting Query Criteria	20-9
To Select a Criterion	20-9
Query Criteria	20-10
Source IP	20-10
Destination IP	20-11
Service	20-11
Event Types	20-11
Device	20-11
Severity/Zone	20-12
Operation	20-12
Rule	20-12

Action	20-12
Saving the Query	20-13
Viewing Events in Real-time	20-13
Restrictions for Real-time Event Viewer	20-13
Procedure for Invoking the Real-Time Event Viewer	20-13
Perform a Long-Duration Query Using a Report	20-17
View a Query Result in the Report Tab	20-19
Perform a Batch Query	20-19
Reports	20-22
Report Type Views: Total vs. Peak vs. Recent	20-23
Creating a Report	20-24
Working With Existing Reports	20-25

CHAPTER 21

Rules 21-1

Rules Overview	21-1
Prioritizing and Identifying	21-2
Think Like a Black Hat	21-2
Planning an Attack	21-2
Back to Being the Admin	21-3
Types of Rules	21-4
Inspection Rules	21-4
Global User Inspection Rules	21-4
Drop Rules	21-4
Constructing a Rule	21-5
Working Examples	21-16
Example A: Excessive Denies to a Particular Port on the Same Host	21-16
Example B: Same Source Causing Excessive Denies on a Particular Port	21-16
Example C: Same Host, Same Destination, Same Port Denied	21-16
Working with System and User Inspection Rules	21-17
Change Rule Status—Active and Inactive	21-17
Duplicate a Rule	21-17
Edit a Rule	21-18
Add an Inspection Rule	21-19
Working with Drop Rules	21-21
Change Drop Rule Status—Active and Inactive	21-21
Duplicate a Drop Rule	21-21
Edit a Drop Rule	21-22
Add a Drop Rule	21-22

Setting Alerts	21-23
Configure an Alert for an Existing Rule	21-24
Rule and Report Groups	21-24
Rule and Report Group Overview	21-25
Global Controller and Local Controller Restrictions for Rule and Report Groups	21-26
Add, Modify, and Delete a Rule Group	21-26
Add, Modify, and Delete a Report Group	21-29
Display Incidents Related to a Rule Group	21-31
Create Query Criteria with Report Groups	21-32
Using Rule Groups in Query Criteria	21-33

CHAPTER 22
Sending Alerts and Incident Notifications 22-1

Configure the E-mail Server Settings	22-4
Configure a Rule to Send an Alert Action	22-5
Create a New User—Role, Identity, Password, and Notification Information	22-10
Create a Custom User Group	22-12
Add a User to a Custom User Group	22-13

CHAPTER 23
Management Tab Overview 23-1

Activating	23-1
To activate a set of management additions or changes	23-1
Event Management	23-1
Search for an Event Description or CVE Names	23-1
To view a list of all currently supported CVEs	23-2
Event Groups	23-2
To filter by event groups or severity	23-2
Edit a Group of Events	23-2
Add a Group	23-2
IP Management	23-3
Search for an Address, Network, Variable, or Host	23-3
Filter by Groups	23-3
Edit a Group	23-3
Add a Group	23-4
Add a Network, IP Range, or Variable	23-4
Add a Host	23-4
Edit Host Information	23-6
Service Management	23-7
Search for a Service	23-7

Add a Group of Services	23-7
Edit a Group of Services	23-7
Add a Service	23-8
Edit a Service	23-8
Delete a Service	23-8
User Management	23-8
Add a New User	23-9
Add a Service Provider (Cell phone/Pager)	23-11
Search for a User	23-11
Edit or Remove a User	23-12
Create a User Group	23-12
Add or Remove a User from a User Group	23-12
Filter by Groups	23-13

CHAPTER 24

System Maintenance 24-1

Setting Runtime Logging Levels	24-1
Viewing the Appliance's Log Files	24-2
View the Back-end Log	24-2
Viewing the Audit Trail	24-3
View an Audit Trail	24-3
Retrieving Raw Messages	24-3
Retrieve Raw Messages From Archive Server	24-3
Retrieve Raw Messages From the Database of a Local Controller	24-5
Hard Drives	24-7
Status Lights	24-7
Partition Checking	24-7
Hotswapping Hard Drives	24-7
Remove a Hard Drive	24-7
Replace a Hard Drive	24-7
Replacing the Lithium Cell CMOS Battery	24-8
Replace the Lithium Cell CMOS Battery	24-8
Change the Default Password of the Administrator Account	24-8

APPENDIX A

Cisco Security MARS XML API Reference A-1

XML Overview	A-1
XML Incident Notification Data File and Schema	A-2
XML Incident Notification Data File Sample Output	A-2
XML Incident Notification Schema	A-4

Usage Guidelines and Conventions for XML Incident Notification A-4

APPENDIX B**Regular Expression Reference B-1**

PCRE Regular Expression Details B-1

Backslash B-2

Non-printing Characters B-3

Generic Character Types B-4

Unicode Character Properties B-5

Simple Assertions B-6

Circumflex and Dollar B-7

Full Stop (Period, Dot) B-8

Matching a Single Byte B-8

Square Brackets and Character Classes B-8

Posix Character Classes B-9

Vertical Bar B-10

Internal Option Setting B-10

Subpatterns B-11

Named Subpatterns B-12

Repetition B-12

Atomic Grouping and Possessive Quantifiers B-14

Back References B-15

Assertions B-16

Lookahead Assertions B-17

Lookbehind Assertions B-17

Using Multiple Assertions B-18

Conditional Subpatterns B-19

Comments B-20

Recursive Patterns B-20

Subpatterns as Subroutines B-21

Callouts B-22

APPENDIX C**Date/Time Format Specification C-1****GLOSSARY****INDEX**



Preface

Introduction

Thank you for purchasing the Cisco Security Monitoring, Analysis, and Response System (MARS) Local Controller. appliance. This guide will help you get the most value from your MARS Appliance.



Note

The information in this document referring to a “MARS appliance” also applies to MARS use as Local Controller in a Global Controller architecture.

The MARS Appliance

The Cisco Security Monitoring, Analysis, and Response System Appliance (MARS Appliance)– the MARS 20, MARS 50, MARS 100, and MARS 200 – is a Security Threat Mitigation (STM) appliance. It delivers a range of information about your networks’ health as seen through the “eyes” and “ears” of the reporting devices in your networks. It takes in all of the raw events from your reporting devices, sessionizes them across different devices, fires default rules for incidents, determines false positives, and delivers consolidated information through diagrams, charts, queries, reports, and rules.

The MARS operates at distinct and separate levels based on how much information is provided about your networks’ devices. At its most basic level, MARS functions as a syslog server. As you add information about reporting devices, it starts sessionizing, and when fully enabled, it presents a bird’s-eye view of your networks with the ability to quickly drill-down to a specific MAC address.

The MARS Web Interface

The MARS user interface uses a tabbed, hyperlinked, browser-based interface. If you have used the Web, you have used similar Web pages.



Note

When using the MARS user interface, avoid using the Back and Forward arrows in the browser. Using these arrows can lead to unpredictable behavior.

About This Manual

This manual describes the features and functionality of the Local Controller. The layout of this manual is as follows:

- [Chapter 1, “STM Task Flow Overview,”](#) recommends a taskflow for planning and implementing your security threat mitigation system. It ties back to your corporate security policies and presents a structure deployment and configuration strategy based on two phases: provisioning and monitoring.

Part 1: Provisioning Phase. This part details provisioning your network devices to communicate with MARS. It involves performing device inventories, bootstrapping and configuring the reporting devices and mitigation devices to communicate with the MARS Appliance, and performing device-side tuning.

- [Chapter 2, “Reporting and Mitigation Devices Overview,”](#) discusses concepts important to a successful deployment of MARS. These concepts include selecting among the devices on your network, understanding the levels of operation, and performing those tasks that affect many devices, such as defining data pulling schedules.
- [Chapter 3, “Configuring Router and Switch Devices.”](#)
- [Chapter 4, “Configuring Firewall Devices.”](#)
- [Chapter 5, “Configuring VPN Devices.”](#)
- [Chapter 6, “Configuring Network-based IDS and IPS Devices.”](#)
- [Chapter 7, “Configuring Host-Based IDS and IPS Devices.”](#)
- [Chapter 8, “Configuring Antivirus Devices.”](#)
- [Chapter 9, “Configuring Vulnerability Assessment Devices.”](#)
- [Chapter 10, “Configuring Generic, Solaris, Linux, and Windows Application Hosts.”](#)
- [Chapter 11, “Configuring Database Applications.”](#)
- [Chapter 12, “Configuring Web Server Devices.”](#)
- [Chapter 13, “Configuring Web Proxy Devices.”](#)
- [Chapter 14, “Configuring AAA Devices.”](#)
- [Chapter 15, “Configuring Custom Devices.”](#)

Part II: Monitoring Phase. This part concepts important to successfully using MARS to monitor your network. These concepts include defining inspection rules and investigating incidents.

- [Chapter 16, “Policy Table Lookup on Cisco Security Manager”](#) explains how to integrate with Cisco Security Manager and use the policy lookup features in MARS.
- [Chapter 17, “Network Summary”](#) covers the Summary pages which includes the Dashboard, the Network Status, and the My Reports pages.
- [Chapter 18, “Case Management”](#) covers using cases to provide accountability and improve workflow.
- [Chapter 19, “Incident Investigation and Mitigation”](#) covers incidents and false positives and provides a starting point for configuring a Layer 2 path and mitigation to work with a MARS.
- [Chapter 20, “Queries and Reports”](#) covers working with scheduled and on-demand reports and queries. It also discussing using the real-time event viewer.
- [Chapter 21, “Rules”](#) covers defining and use inspection rules.

- [Chapter 22, “Sending Alerts and Incident Notifications”](#) explains how to configure the MARS to send an alert based on an inspection rule.
- [Chapter 23, “Management Tab Overview”](#) covers managing events, networks, variables, hosts, services, and MARS users.
- [Chapter 24, “System Maintenance”](#) covers some of the maintenance chores for the MARS.

Additionally, the following appendices are provided:

- [Appendix A, “Cisco Security MARS XML API Reference”](#) presents the XML schema used by MARS for XML-based notifications.
- [Appendix B, “Regular Expression Reference”](#) The syntax and semantics of the regular expressions supported by PCRE are described in this appendix.
- [Appendix C, “Date/Time Format Specification”](#) The date/time field parsing is supported using the Unix `strptime()` standard C library function.
- [Glossary](#) — A glossary of terms as they relate to MARS.

Obtaining Documentation

Cisco documentation and additional literature are available on Cisco.com. Cisco also provides several ways to obtain technical assistance and other technical resources. These sections explain how to obtain technical information from Cisco Systems.

Cisco.com

You can access the most current Cisco documentation at this URL:

<http://www.cisco.com/univercd/home/home.htm>

You can access the Cisco webbiest at this URL:

<http://www.cisco.com>

You can access international Cisco websites at this URL:

http://www.cisco.com/public/countries_languages.shtml

Documentation DVD

Cisco documentation and additional literature are available in a Documentation DVD package, which may have shipped with your product. The Documentation DVD is updated regularly and may be more current than printed documentation. The Documentation DVD package is available as a single unit.

Registered Cisco.com users (Cisco direct customers) can order a Cisco Documentation DVD (product number DOC-DOCDVD=) from the Ordering tool or Cisco Marketplace.

Cisco Ordering tool:

<http://www.cisco.com/en/US/partner/ordering/>

Cisco Marketplace:

<http://www.cisco.com/go/marketplace/>

Ordering Documentation

You can find instructions for ordering documentation at this URL:

http://www.cisco.com/univercd/cc/td/doc/es_inpk/pdi.htm

You can order Cisco documentation in these ways:

- Registered Cisco.com users (Cisco direct customers) can order Cisco product documentation from the Ordering tool:

<http://www.cisco.com/en/US/partner/ordering/>

- Nonregistered Cisco.com users can order documentation through a local account representative by calling Cisco Systems Corporate Headquarters (California, USA) at 408 526-7208 or, elsewhere in North America, by calling 1 800 553-NETS (6387).

Documentation Feedback

You can send comments about technical documentation to bug-doc@cisco.com.

You can submit comments by using the response card (if present) behind the front cover of your document or by writing to the following address:

Cisco Systems
Attn: Customer Document Ordering
170 West Tasman Drive
San Jose, CA 95134-9883

We appreciate your comments.

Cisco Product Security Overview

Cisco provides a free online Security Vulnerability Policy portal at this URL:

http://www.cisco.com/en/US/products/products_security_vulnerability_policy.html

From this site, you can perform these tasks:

- Report security vulnerabilities in Cisco products.
- Obtain assistance with security incidents that involve Cisco products.
- Register to receive security information from Cisco.

A current list of security advisories and notices for Cisco products is available at this URL:

<http://www.cisco.com/go/psirt>

If you prefer to see advisories and notices as they are updated in real time, you can access a Product Security Incident Response Team Really Simple Syndication (PSIRT RSS) feed from this URL:

http://www.cisco.com/en/US/products/products_psirt_rss_feed.html

Reporting Security Problems in Cisco Products

Cisco is committed to delivering secure products. We test our products internally before we release them, and we strive to correct all vulnerabilities quickly. If you think that you might have identified a vulnerability in a Cisco product, contact PSIRT:

- Emergencies—security-alert@cisco.com
- Nonemergencies—psirt@cisco.com

**Tip**

We encourage you to use Pretty Good Privacy (PGP) or a compatible product to encrypt any sensitive information that you send to Cisco. PSIRT can work from encrypted information that is compatible with PGP versions 2.x through 8.x.

Never use a revoked or an expired encryption key. The correct public key to use in your correspondence with PSIRT is the one that has the most recent creation date in this public key server list:

<http://pgp.mit.edu:11371/pks/lookup?search=psirt%40cisco.com&op=index&exact=on>

In an emergency, you can also reach PSIRT by telephone:

- 1 877 228-7302
- 1 408 525-6532

Obtaining Technical Assistance

For all customers, partners, resellers, and distributors who hold valid Cisco service contracts, Cisco Technical Support provides 24-hour-a-day, award-winning technical assistance. The Cisco Technical Support Website on Cisco.com features extensive online support resources. In addition, Cisco Technical Assistance Center (TAC) engineers provide telephone support. If you do not hold a valid Cisco service contract, contact your reseller.

Cisco Technical Support Website

The Cisco Technical Support Website provides online documents and tools for troubleshooting and resolving technical issues with Cisco products and technologies. The website is available 24 hours a day, 365 days a year, at this URL:

<http://www.cisco.com/techsupport>

Access to all tools on the Cisco Technical Support Website requires a Cisco.com user ID and password. If you have a valid service contract but do not have a user ID or password, you can register at this URL:

<http://tools.cisco.com/RPF/register/register.do>

**Note**

Use the Cisco Product Identification (CPI) tool to locate your product serial number before submitting a web or phone request for service. You can access the CPI tool from the Cisco Technical Support Website by clicking the **Tools & Resources** link under Documentation & Tools. Choose **Cisco Product Identification Tool** from the Alphabetical Index drop-down list, or click the **Cisco Product Identification Tool** link under Alerts & RMAs. The CPI tool offers three search options: by product ID

or model name; by tree view; or for certain products, by copying and pasting **show** command output. Search results show an illustration of your product with the serial number label location highlighted. Locate the serial number label on your product and record the information before placing a service call.

Submitting a Service Request

Using the online TAC Service Request Tool is the fastest way to open S3 and S4 service requests. (S3 and S4 service requests are those in which your network is minimally impaired or for which you require product information.) After you describe your situation, the TAC Service Request Tool provides recommended solutions. If your issue is not resolved using the recommended resources, your service request is assigned to a Cisco TAC engineer. The TAC Service Request Tool is located at this URL:

<http://www.cisco.com/techsupport/servicerequest>

For S1 or S2 service requests or if you do not have Internet access, contact the Cisco TAC by telephone. (S1 or S2 service requests are those in which your production network is down or severely degraded.) Cisco TAC engineers are assigned immediately to S1 and S2 service requests to help keep your business operations running smoothly.

To open a service request by telephone, use one of the following numbers:

Asia-Pacific: +61 2 8446 7411 (Australia: 1 800 805 227)

EMEA: +32 2 704 55 55

USA: 1 800 553-2447

For a complete list of Cisco TAC contacts, go to this URL:

<http://www.cisco.com/techsupport/contacts>

Definitions of Service Request Severity

To ensure that all service requests are reported in a standard format, Cisco has established severity definitions.

Severity 1 (S1)—Your network is “down,” or there is a critical impact to your business operations. You and Cisco will commit all necessary resources around the clock to resolve the situation.

Severity 2 (S2)—Operation of an existing network is severely degraded, or significant aspects of your business operation are negatively affected by inadequate performance of Cisco products. You and Cisco will commit full-time resources during normal business hours to resolve the situation.

Severity 3 (S3)—Operational performance of your network is impaired, but most business operations remain functional. You and Cisco will commit resources during normal business hours to restore service to satisfactory levels.

Severity 4 (S4)—You require information or assistance with Cisco product capabilities, installation, or configuration. There is little or no effect on your business operations.

Obtaining Additional Publications and Information

Information about Cisco products, technologies, and network solutions is available from various online and printed sources.

- Cisco Marketplace provides a variety of Cisco books, reference guides, and logo merchandise. Visit Cisco Marketplace, the company store, at this URL:

<http://www.cisco.com/go/marketplace/>

- *Cisco Press* publishes a wide range of general networking, training and certification titles. Both new and experienced users will benefit from these publications. For current Cisco Press titles and other information, go to Cisco Press at this URL:

<http://www.ciscopress.com>

- *Packet* magazine is the Cisco Systems technical user magazine for maximizing Internet and networking investments. Each quarter, Packet delivers coverage of the latest industry trends, technology breakthroughs, and Cisco products and solutions, as well as network deployment and troubleshooting tips, configuration examples, customer case studies, certification and training information, and links to scores of in-depth online resources. You can access Packet magazine at this URL:

<http://www.cisco.com/packet>

- *iQ Magazine* is the quarterly publication from Cisco Systems designed to help growing companies learn how they can use technology to increase revenue, streamline their business, and expand services. The publication identifies the challenges facing these companies and the technologies to help solve them, using real-world case studies and business strategies to help readers make sound technology investment decisions. You can access iQ Magazine at this URL:

<http://www.cisco.com/go/iqmagazine>

- *Internet Protocol Journal* is a quarterly journal published by Cisco Systems for engineering professionals involved in designing, developing, and operating public and private internets and intranets. You can access the Internet Protocol Journal at this URL:

<http://www.cisco.com/ipj>

- World-class networking training is available from Cisco. You can view current offerings at this URL:

<http://www.cisco.com/en/US/learning/index.html>



STM Task Flow Overview

This chapter describes the project phases and task flows that you should follow when you deploy MARS as a security threat mitigation (STM) system in your network. First, however, you must develop a set of policies that enables the application of security measures.

Your security policy should:

- Identify security objectives for your organization.
- Document the resources to protect.
- Identify the network infrastructure with current maps and inventories.
- Identify the critical resources (such as research and development, finance, and human resources) that require extra protection.

Your monitoring policy should:

- Identify the expected administrative traffic flows across your network, including user, source, destination, services, and hours of operation.
- Identify expected network traffic for security probing and vulnerability testing, including user, source, destination, services, and hours of operation.
- Identify the network infrastructure able to provide audit data in “network proximity” to the critical resources.
- Identify the various event logging levels available from the devices and hosts in the network infrastructure.
- Identify the devices and techniques used to investigate

Your mitigation policy should:


- Identify the choke points on your network relative to the critical resources.
- Define your process for documenting mitigated attacks on layer 2 and layer 3 devices.
- Define your process for documenting mitigated attacks at the host and application layer.
- Resolve corporate ownership issues among network operations, security operations, host owners and application owners on shared hosts.
- Identify your policy for notifying security response teams and remediation teams.
- Identify vendor detection tool prioritization process, such as IOS IPS Dynamic Attack Mitigation (DAM).
- Identify how you want to block detected attacks: block them temporarily or permanently, block them using MARS-generated rules, using custom rules defined by security operations team, etc.

Your remediation policy should:

- Identify the responses to detected but unmitigated attacks for each type of node in your network.
- Identify tool vendor update policies to ensure proper remediation of hosts and applications.
- Identify the policies and procedures for isolating infected legacy hosts where remediation options are unavailable. These procedures may include restoring from backups or network isolation.

After you develop your policies, they become the hub of the Cisco Security Wheel, (Figure 1-1).

Figure 1-1 Cisco Security Wheel



The spokes of the Cisco Security Wheel represent network security as a continual process consisting of four steps:

1. Secure your system.
2. Monitor the network for violations and attacks against your security policy and respond to them.
3. Test the effectiveness of the security safeguards in place.
4. Manage and improve corporate security.

You should perform all four steps continually, and you should consider each of them when you create and update your corporate security policy.

The remainder of this section details recommended task flows according to the following project phases:

- Provisioning (see [Checklist for Provisioning Phase, page 1-2](#)).
- Monitoring (see [Checklist for Monitoring Phase, page 1-9](#)).

Check out <http://www.cisco.com/web/about/security/intelligence/articles.html> for more planning ideas. Look closely at the SAFE information.

Checklist for Provisioning Phase

Provisioning deals with planning, setting up and configuring the hardware, software, and networks that actually provide access to the data and network resources for the MARS Appliance. This phase takes place after you successfully complete the installation, which was detailed in the *Install and Setup Guide for Cisco Security Monitoring, Analysis, and Response System*.

The following checklist describes the tasks required to understand the decision-making process and the basic flow required to provision MARS in the most productive manner. Each step might contain several substeps; the steps and substeps should be performed in order. The checklist contains references to the specific procedures used to perform each task.

✓	Task
☐	1. Inventory and review possible reporting devices, mitigation devices, and supporting devices. <i>Reporting devices</i> provide logs about user and network activities and device status and configuration. <i>Mitigation devices</i> can be used to respond to detected attacks. They also act as reporting devices. <i>Supporting devices</i> provide network services to reporting devices, mitigation devices, or a MARS Appliance. Identifying which devices on your network to monitor depends on multiple factors, including their placement, the reporting they can provide relative to other devices on the same network segment, and the level of operation that you want to achieve from your MARS Appliance. When considering which devices to declare as reporting devices and mitigation devices, be sure you know what data is provided to MARS by those devices. Simply adding all possible devices does not guarantee the best monitoring and mitigation strategy. Deliberate selection of the devices can reduce the MARS workload, resulting in improved detection and mitigation times, as well as improved false positive detection. Because MARS only considers monitored devices, you should take care in identifying which devices to monitor. The following are only a couple examples of considerations you should make when identifying devices. • Consider of the types of logs and data available from reporting devices on specific network segments, and select those logs that provide the most complete picture of the activity on your network. • Identify mitigation devices at natural chokepoints across each segment in your network. You are more likely to stop an attack if these mitigation devices are identified to MARS. When MARS identifies an attack, it studies the topology of your network to identify the best chokepoint; however, it only considers those devices that are monitored. Supporting devices can play an important role in the operation of your STM system. Therefore, you should inventory and review the supporting devices on your network, which include e-mail, AAA, DNS, and syslog servers, that will play a role in the envisioned STM system. <i>Result:</i> The list of devices that you want to monitor is complete. The details of each device include device name, reporting IP address, management IP address, management protocol, administrative account information, and the logging features, levels, and protocols to enable. For more information, see: • Selecting the Devices to Monitor, page 2-2 • Levels of Operation, page 2-1 • Deployment Planning Guidelines, page 2-1 in <i>Install and Setup Guide for Cisco Security Monitoring, Analysis, and Response System</i> • Device Inventory Worksheet, page 1-18

✓	Task
□	2. Identify and enable all required traffic flows. After you identify the devices, you must verify that the network services they use for management, reporting, and notification are permitted along the required traffic flows. Using the detailed Device Inventory Worksheet identified in Step 1., ensure that the management, logging, and notification traffic between the MARS Appliance and each supporting device, reporting device, and mitigation device is allowed by intermediate gateways. In addition, network services of supporting devices, such as DNS, e-mail, AAA, and NTP servers, must also be permitted to flow among the MARS Appliance, the supporting devices, and the reporting devices and mitigation devices on your network. MARS applies the device time to received events only. For all events pulled from devices such as IDS/IPS devices or Windows, MARS uses the reported time as long as that reported time falls within 3600 seconds of the time on the MARS Appliance. Tip It is a recommended security practice to have all devices, including MARS Appliances, synchronized to the same time. Also, since the MARS Appliance is an HTTPS server, it uses certificates which require the time, date, and time zone to be set properly. Otherwise, sessions and incidents are stamped incorrectly and you may experience “time out” errors when accessing the web interface. To limit troubleshooting, you should test each traffic flow from the source network segment to the destination segment. If possible, you should test all device-to- device flows for each protocol to ensure that best match versus first match semantics of various gateway ACLs do not hinder required traffic flows. As with any security devices on your network, enabled traffic flows should be restricted to the required protocols, ports, and source/destination pairs. <i>Result:</i> You have verified that all intermediate gateways permit the log, management, and notification traffic between the devices and the MARS Appliance. For more information, see: • Event Timestamps and Processing in <i>Top Issues for the Cisco Security Monitoring, Analysis, and Response System</i> • Deployment Planning Guidelines, page 2-1, in <i>Install and Setup Guide for Cisco Security Monitoring, Analysis, and Response System</i> • Supporting Devices, page 2-1, in <i>Install and Setup Guide for Cisco Security Monitoring, Analysis, and Response System</i> • Required Traffic Flows, page 2-2, in <i>Install and Setup Guide for Cisco Security Monitoring, Analysis, and Response System</i> • Specify the Time Settings, page 5-10, in <i>Install and Setup Guide for Cisco Security Monitoring, Analysis, and Response System</i> • Device Inventory Worksheet, page 1-18

✓	Task
☐	3. Bootstrap the reporting devices, mitigation devices, and supporting devices. For each device identified in the Device Inventory Worksheet, you must prepare, or bootstrap, that device to ensure that the desired communications with MARS occur. Bootstrapping a device involves configuring the settings for that device, as determined by its role within the STM system. Perform the following bootstrap tasks as applicable to a device type and its role: • Enable management of the device by the MARS Appliance for mitigation and access. • Install an agent that collects the correct logs for MARS Appliance. • Turn on the correct logging level and logging services. • Direct the logs to the MARS Appliance or identify the appliance to receive or pull those logs as needed. • Enable discovery of the device settings. • Enable the device to receive notifications from the MARS Appliance. Each device has a different required configuration to ensure that it assumes the role you have envisioned for it in the STM system. As you consider the devices, their expected role in your STM system will correlate directly with the configuration of the tasks listed above. In addition, you identify any restrictions imposed by MARS. For example, MARS may restrict the supported protocols for discovery of a specific device type. <i>Result:</i> The correct logging levels are enabled on the reporting devices and mitigation devices. The MARS Appliance can receive or pull any necessary logs from those devices, and it can retrieve configuration settings and push ACLS to the supported mitigation devices. Any devices that require notification of detected attacks are configured to receive such notifications from the MARS Appliance. While the MARS Appliance picks up and stores the events it receives, it does not inspect them until the reporting devices and mitigation devices are defined and activated in web interface. Tip Any events published by a device to MARS prior to adding and activating the device in the web interface can be queried using the reporting IP address of the device as a match criterion. This technique can be useful for verifying that the device is properly bootstrapped. For more information, see: • Device Inventory Worksheet, page 1-18 • Supported Reporting and Mitigation Devices, page 2 • Bootstrap Summary Table, page 2-12 • The log settings sections of the user guides for your reporting devices and mitigation devices

✓	Task
☐	4. Define the devices in MARS. After you identify and bootstrap the reporting devices and mitigation devices and enable the required traffic flows, you must represent those devices in MARS, which uses this information to communicate with the devices. You can do this by adding individual devices in the web interface or by importing a comma separated vector (CSV) file, which can define the required settings for basic device types and give you a headstart on defining the more complicated devices. In addition, you can use topology discovery to automatically discover reporting devices and mitigation devices and later go back to provide additional detail. For most device types, you must determine what access protocol to use for device discovery. The selection of this protocol determines what type of data you can discover and whether you can perform mitigation. Understanding the options helps you develop a consistent approach in compliance with your corporate policies. How you choose to add the devices depends on the number of devices on your network and whether there are CSV device keywords for the devices that you want to add. In addition, device types that use agents, modules, or sensors are defined in multiple steps, where you first define the base host or device, and then add the modules, sensors, and agents to the base device. For example, if you want to add an IPS module to a Cisco ASA device, you must first define the Cisco ASA device and then define the IPS module as a component of that device. In addition, many applications that are not dedicated appliances require that you first define the host (generic, Windows, Unix, or Linux) on which that application runs before you can associate the application with that host. After you add the devices, you must activate them by clicking Activate on any page in the web interface. To display all devices that are either added incorrectly or not activated in MARS, you can define one of two queries: • Select “Unknown Reporting Device” in the Devices field. This query returns the events only for those devices that are reporting events that do not matching the one of the reporting IPs defined in MARS. When MARS receives events, it first determines if the IP from which the events are received matches one of reporting IPs identified in the Reporting and Monitor Devices page. Only if MARS finds a match does it attempt to parse the events. Therefore, if the Reporting IP is defined incorrectly for a reporting device, the events from that device are not parsed. This query essentially identifies events that are not parsed. • Select the “Unknown Device Event Type” in the Events field. This query returns events from known devices that for some reason the event is not parsed by MARS (for example, if the MARS signature list is not current with the device event lists), and it returns events reported by unknown devices. These queries are a recommended good practice after adding the devices, especially when using a CSV seedfile or SNMP discovery. For both queries, if you are looking for a specific reporting IP address, enter that address in the Keyword field to filter the results down to those that include that IP address. <i>Result:</i> All reporting devices and mitigation devices are defined and activated in MARS. When the devices are bootstrapped and defined in MARS, MARS begins to inspect the logs received from the devices. Until the devices are added in MARS, MARS picks up and stores the events it receives without inspecting them. For more information, see: • Device Inventory Worksheet, page 1-18 • Selecting the Access Type, page 2-10 • Add Reporting and Mitigation Devices Individually, page 2-17 • Add Multiple Reporting and Mitigation Devices Using a Seed File, page 2-20 • Adding Reporting and Mitigation Devices Using Automatic Topology Discovery, page 2-25 • Supported Reporting and Mitigation Devices, page 2 (CSV Keyword column) • Verify Connectivity with the Reporting and Mitigation Devices, page 2-26 • Activate the Reporting and Mitigation Devices, page 2-27

✓	Task
	5. Configure global data collection settings and schedules in MARS. After you add the devices, you can enable the rich data collection features of MARS, which include: • Dynamic vulnerability scanning. When MARS detects an attack, it can probe the network to determine the likely success and severity of the attack. To allow this data collection in response to detected attacks, you must enable the feature and identify which networks to analyze. • NetFlow data collection. NetFlow data enables MARS to identify anomalies by profiling typical data flows across your network, allowing MARS to detect day-zero attacks, including worm outbreaks. Statistical profiling takes between four days and two weeks for a MARS Appliance to complete. When the profiles are developed, MARS begins detecting anomalous traffic flows and creates incidents in response to them. To configure NetFlow data collection, you must configure those devices that can generate NetFlow traffic, and you must configure MARS to listen on a shared community string. • Layer 3 topology discovery. A process-intensive operation that discovers the layer 3 network devices (that is, those devices operating at the IP layer). This layer 3 data is used to determine the attack path vector and to populate the Topology graphs. You can define the schedule for updating this information. • Layer 2 device discovery. This feature allows MARS to determine the attack path vector and to identify attacking hosts and targets by MAC address, which eliminates confusion caused by attacks that spoof IP addresses. This feature is typically configured when adding a switch and enabling mitigation. There are also several device types from which MARS periodically pulls data. For such devices, you can define the intervals at which the event logs are retrieved and processed. These update features are as follows: • Distributed Threat Mitigation (DTM) device updates. The DTM services poll Cisco IPS and Cisco IDS devices to determine the top firing signatures across the reporting devices. Based on this information, MARS generates the list of top signatures that are firing on the network so that Cisco IOS Routers running the DTM feature set can query MARS for the list of signatures they should be running. • Windows event logs. You can set the frequency by which MARS pulls audit trail records from Windows hosts and servers. This setting is global for all such hosts and has a default value of five minutes. • Oracle event logs. You can set the frequency by which MARS pulls audit trail records from Oracle database servers. This setting is global for all such servers and has a default value of five minutes. • Monitored device update scheduler. You can set the frequency by which MARS pulls data from specific reporting devices, such as Qualys QualysGuard, Foundstone Foundscan, and eEye REM. Schedules are set on a per IP address basis. After you define the settings, you must activate them by clicking Activate on any page in the web interface. <i>Result:</i> The schedules for updating cached data pulled from reporting, mitigation, and supporting devices are defined and activated in MARS. After these settings are defined, MARS can probe the network or pull updates from reporting, mitigation, and supporting devices. For more information, see: • Data Enabling Features, page 2-28 • Windows Event Log Pulling Time Interval, page 10-10 • Layer 2 Discovery and Mitigation, page 2-29 • Configure Interval for Pulling Oracle Event Logs, page 11-3 • Networks for Dynamic Vulnerability Scanning, page 2-29 • Understanding NetFlow Anomaly Detection, page 2-30 • Configuring Layer 3 Topology Discovery, page 2-36 • Technology Preview: Configuring Distributed Threat Mitigation with Intrusion Prevention System in Cisco Security MARS, page 1 • Scheduling Topology Updates, page 2-39

✓	Task
	6. Populate vulnerability assessment information for supporting devices and network assets. Vulnerability assessment information describes specific hosts on your network. You can detail this information for any host, whether it is a host representing a reporting device, a mitigation device, or an important asset on your network. This information identifies the operating system, patch levels, and the network services that run on the host. After you define the hosts, you must activate them by clicking Activate on any page in the web interface. <i>Result:</i> MARS understands more about the hosts on your network and the services that they run. For more information, see: • Host and Device Identification and Detail Strategies, page 2-36 • Device Inventory Worksheet, page 1-18 • IP Management, page 23-3 • Service Management, page 23-7

✓	Task
☐	7. Monitor and tune event generation and processing. As with all monitoring applications, tuning log generation and event processing is key to technical accuracy and performance. You can use two methods to tune which events are processed by MARS: • Device-side tuning. This method involves restricting event generation at the device level. MARS never receives events that are not relevant to security or device status. It also involves eliminating superfluous, duplicate data reported by multiple devices on the network, as well as eliminating those events that can be reproduced by reports or queries in MARS, such as traffic summary syslogs. • Appliance-side tuning. This method involves identifying events received by the MARS Appliance that represent normal or planned network activity. Drop rules are defined to prevent MARS from processing such events as part of potential security incidents. When defining such drop rules, you should be as precise in the definition as possible, for example, identify the source of expected ping sweeps by an IP address within an expected time period, which is much more difficult to spoof as it requires explicit knowledge of your network and administrative practices. You can further qualify the rules using a combination of seven conditions: source, destination, service type, event type, time range, reporting device, and event severity. You must choose whether to drop the event entirely or to drop it and log it to the database, where it can be used by queries and reports. Note Drop rules do not prevent MARS from storing the event data; they simply prevent the appliance from processing the events. Events affected by drop rules can still appear a query as they are being stored on the appliance. You are still storing them; just not processing them for inspection rules. Therefore, if appliance storage considerations are an issue, we recommend using device-side tuning. Tuning is an ongoing task to improve the identification of attacks, report quality surrounding truly suspicious activities, and the overall performance and accuracy of your STM solution. It involves a detailed study of traffic, which can be conducted and refined by evaluating the events that are coming into the appliance on a device-by-device basis. Tip In a lab network environment, use a MARS Appliance to study generated events and tuning options on an individual device type basis. By documenting your requirements in a controlled environment, you can eliminate much of the production network tuning by establishing valuable device-side tuning standards for each monitoring device type. <i>Result:</i> The events being processed by the MARS Appliance are restricted to those that provide value to the STM system. For more information, see: • Appliance-side Tuning Guidelines, page 1-17 • Configuring Logging Policies on Firewall Devices in <i>User Guide for Cisco Security Manager 3.0</i>

Checklist for Monitoring Phase

After you complete the provisioning phase, you must configure MARS to help you realize your broader security goals and requirements. During the monitoring phase, your primary goal is to effectively realize your monitoring, mitigation, and remediation policies. This phase involves defining the strategies, rules, reports, and other settings required to achieve this goal.



Note

You must prepare MARS to closely adhere to your corporate security policy before you begin monitoring traffic flows, as you must be prepared to react to detected attacks.

The following checklist describes the tasks required to understand the decision-making process and the basic flow required to operate MARS in the most productive manner. Each step might contain several substeps; the steps and substeps should be performed in order. The checklist contains references to the specific procedures used to perform each task.

✓	Task
☐	1. Develop monitoring, notification, mitigation, remediation, and audit strategies. These strategies are concerned less with desired traffic flows and generated events and focus more on what to do after MARS Appliance processes that data. These strategies are at the heart of how you will use MARS to protect your network, taking into account the short- and long-term requirements of monitoring and forensic analysis, as well as how to stop ongoing attacks and clean infected hosts. These strategies encompass not only your expected interaction with MARS, but the expectations of your reporting devices as well. Essentially, they identify the roles, tasks, and data requirements that you anticipate so that you can map events, rules, queries, and reports to those roles that provide the data required by the identified tasks. As with any security system, we recommend that users be assigned the lowest-level privilege required to perform their job. Admin-level privileges should be reserved for administrators of the MARS Appliance. <i>Result:</i> You have identified the users and roles required to effectively respond to detected attacks and device issues. You have defined clear guidance for responding to notifications and understand the information requirements of those such notifications and the expected format and delivery methods to be used. For more information, see: • Strategies for Monitoring, Notification, Mitigation, Remediation, and Audit, page 1-16 • Case Management, page 18-1s • User Management, page 23-8 • , page 23-13 • User Role Worksheet, page 1-20

✓	Task
☐	2. Define the notification services. This task prepares the notification services of MARS to notify your mitigation and remediation personnel and take other required actions. In MARS, notification services have three building blocks: • User accounts. Represent users who will receive reports or notifications or who will access the web interface for the purpose of monitoring or mitigation. Users can receive notifications in the form of e-mail, pager messages, or Short Message Service (SMS) messages. Users are assigned to one of four roles, admin, security analyst, operator, notification only, which determines their access privileges in the web interface. • Devices. Represent those devices that will receive notifications in the form of an SNMP message, a syslog message, or in the case of an IOS IPS device, a DAM message (equivalent to a shun). For more on defining devices, refer to Checklist for Provisioning Phase, page 1-2. • Actions. Actions are defined within inspection rules, and they represent the notifying action. Depending on the target of the notification, a user or a device, your action can provide guidance to your staff or instruct your devices to log or block an attack. Within MARS, any person or device that is expected to receive a notification must be identified in the system. Therefore, the first step is to define user accounts that map to the users or groups who must be notified based on specific event settings (see User Role Worksheet, page 1-20). You must also identify the devices that need to be notified or that need to take some action (see Device Inventory Worksheet, page 1-18). The next step is to define the notification service settings (actions), which can be one or more of e-mail, page, SMS, SNMP, Syslog, or Dynamic Attack Mitigation. Each of these settings includes the contact information and a message that you can define for each type of notification. There is not a separate interface for defining these settings. To define the notification service settings, you must edit an existing inspection rule and add new Action definitions. After you define these settings, they are available to all inspection rules. <i>Result:</i> All required personnel have been identified in MARS so that rules and reports can be customized to notify the correct personnel. For more information, see: • User Management, page 23-8 • Add or Remove a User from a User Group, page 23-12 • IP Management, page 23-3 • Adding Reporting and Mitigation Devices, page 2-16 • Forwarding Alert Data to 3rd-Party Syslog and SNMP Servers, page 2-43 • MARS MIB Format, page 2-43 • Inspection Rules, page 21-4 • Working with System and User Inspection Rules, page 21-17 • Setting Alerts, page 21-23 • Sending Alerts and Incident Notifications, page 22-1

✓	Task
□	3. Define custom inspection rules and refine system inspection rules. Inspection rules correlate events from disparate devices into meaningful sessions that reflect the end-to-end activities of an attack or other network session. By identifying the end-to-end view of attacks, MARS is better able to identify mitigation points in your network. However, you can define inspection rules to accomplish different goals: identification of an attack is just one possible goal. Other example goals include identifying use of priority assets, network health, and refining your network configuration based on usage analysis. MARS ships with over 100 system inspection rules; however, you may find that you cannot identify those sessions that are important to your corporate policies. For example, if you want to monitor the use of a custom or unsupported application, you can either define a new inspection rule that monitors traffic between a selected source and destination using a known protocol and port pair, or define a custom log parser that uniquely processes the events generated by that application to expose the data within the event that you want to track. Monitoring a known protocol port pair can provide summary data, such as number of sessions, where a custom log parser can enable detailed inspection of aspects of the traffic, such as resource utilization or failed logging attempts. To define a custom parser, you must know the message format used by that appliance and it must be published to MARS in clear text. Organizing the rules that you create into meaningful groups can help clarify your purpose and improves the learnability of the system. As you consider your specific goals, you should define a rule group (and a corresponding report group) to help you refine the strategies you identified in Step 1. Because rules can be members of multiple groups, you do not have to worry about creating multiple rules to address the same issue. The groups are merely available to help your organize your work and allow you to focus on one strategy at a time. <i>Result:</i> Any custom inspection rules are developed and existing inspection rules are configured to provide proper notification in compliance with your corporate policies. Any custom log parser and inspection rules are defined that enable the audit of the traffic flows of home-grown or unsupported applications or protocols. For more information, see: • Rule and Report Groups, page 21-24 • Event Management, page 23-1 • IP Management, page 23-3 • Service Management, page 23-7 • User Management, page 23-8 • Adding User Defined Log Parser Templates, page 15-1 • Inspection Rules, page 21-4 • Working with System and User Inspection Rules, page 21-17 • Setting Alerts, page 21-23 • Sending Alerts and Incident Notifications, page 22-1

✓	Task
☐	4. Define custom queries and reports. Queries and reports are forensic analysis tools. They help you analyze historical data and enable you to identify trends over longer periods of time than the real-time monitoring features of MARS. The relationship between queries and reports is essentially that queries are on-demand, refined inspections of data as defined by a report template. Reports are scheduled to run periodically, enabling you to define the periods and frequencies that you want to inspect on an ongoing basis. Queries allow you to narrow or broaden your search based on a report template by filtering the search criteria. While MARS provides many predefined report templates, you can define new report templates that focus on the incidents and events important to fulfilling your policies. This feature can be especially useful for adhering to compliance reporting requirements, as you can define a report, schedule it to be generated, and store the results as part of your audit records. As with overall access, you can restrict the ability to run or view reports and queries based on user role. Such safeguards can reduce accidental tampering with scheduled reports by other users of the system. In addition, you can configure your report templates so that users are notified of the report. Typically, e-mail is the primary method used for report notification, but all notification methods are supported. <i>Result:</i> The report templates required to realize your forensic analysis and audit goals are defined and assigned to user roles according to your least privilege policies. Any report groups that facilitate access or division of reports and queries among your staff are defined. For more information, see: • Queries and Reports, page 20-1 • Queries, page 20-1 • Perform a Batch Query, page 20-19 • Reports, page 20-22 • Creating a Report, page 20-24

✓	Task
	5. Monitor network and security activity. This task encompasses monitoring your network for attacks or issues and responding to them. How users interact with MARS depends on their role and your operational guidelines. For users who are expected to use the web interface to monitor traffic in near real-time, this task requires an in-depth understanding of the data that is correlated and displayed, as well as when and how to respond to suspicious or anomalous behavior. MARS provides two interfaces to network and security activity: the Summary tab and the Query/Reports tab. Each interface provides different views and tools to help you understand what is happening on your network. The Summary tab focuses on near real-time events, whereas the Query/Reports tab focuses on historical, forensic analysis as described in Step 4. The Summary tab organizes priority views of your network activity, displaying hot spot diagrams, recent events, charts of incidents, and a topology diagram, identifying recent activities. When you identify an incident that requires further investigation or mitigation, you can investigate the incident to determine whether it is a false positive or block attack using MARS. If you have choke points operating at layer 2, primarily switches, MARS will identify the appropriate device, provide recommended CLI changes, and allow you to push these changes to the device. If the choke point is a layer 3 device, MARS recommends CLI changes that you can copy and paste into an administrative session with the identified choke point. In this manner, you can monitor your network for suspicious behavior and respond to any detections. <i>Result:</i> Users understand the views and tools required to monitor, verify, and mitigate attacks on the network. For more information, see: • Network Summary, page 17-1 • Incident Investigation and Mitigation, page 19-1 • False Positive Confirmation, page 19-6 • Rule and Report Groups, page 21-24 • Event Groups, page 23-2 • Case Management, page 18-1 • The False Positive Page, page 19-8 • Retrieving Raw Messages, page 24-3

✓	Task
	6. Monitor system and network health. The STM system is more than your MARS Appliance; it includes all reporting devices and mitigation devices and any MARS Appliances. When assessing the health of the system, you should monitor the health of each of these devices. You can monitor your system health by using inspection rules that generate notifications for anomalous behavior, by generating system health queries and reports, and by manually reviewing the system logs of MARS. MARS provides reports about use of common resources, including CPU, bandwidth, and memory. To simplify the monitoring of system health, you can define a report group that organizes these reports into a meaningful collection. You can also restrict the presentation of those reports and queries to specific user roles. Because reports can be scheduled, you can notify the appropriate users each time the report is updated. Tip If you cannot view the resource usage of a reporting device, verify that you have enabled the Monitor Resource Usage option as part of that device definition in Admin > System Configuration > Security and Monitored Devices. For the list of devices that can be configured to provide this data, see Configuring Resource Usage Data, page 2-41. MARS also includes detailed logs about the status of the appliance itself, as well as several command-line utilities that present status on the health of the appliance. <i>Result:</i> The users responsible for monitoring the system and network health understand the tools and reports provided by MARS to perform these functions. For more information, see: • Rule and Report Groups, page 21-24 • Rule and Report Group Overview, page 21-25 • Configuring Resource Usage Data, page 2-41 • pnstatus, page A-22 • pnlog, page A-18 • Setting Runtime Logging Levels, page 24-1 • Viewing the Appliance's Log Files, page 24-2 • Viewing the Audit Trail, page 24-3 • Retrieving Raw Messages, page 24-3

✓	Task
□	7. Tune MARS processing. Tuning, which is an ongoing activity for any monitoring application, involves refining the sensitivity and accuracy of how events are processed. In MARS, you can do any of the following to effect such changes. • Use drop rules to enable or disable the processing of events by MARS. • Turn on or off event generation at the device. • Identify selected incidents as false positives. • Tune inspection rules to include or exclude specific networks, hosts, services, reporting devices, or traffic flows. • Tune the inspection of traffic by device type, such as IPS and IDS, refining the rule set they use to generate events. • Add or remove reporting devices to alter the reported event set or to provide supporting data that can be used to improve the self-tuning features of MARS, such as false positives, OS fingerprinting, and vulnerability assessment. • Describe the expected behavior on your network by describing the assets, services, and vulnerability assessment information. The more details MARS knows about your network, the better it can assess the incoming events. <i>Result:</i> The events being processed by the MARS Appliance are restricted or expanded to encompass those that provide the most value to the STM system. For more information, see: • Appliance-side Tuning Guidelines, page 1-17 • Working with Drop Rules, page 21-21 • False Positive Confirmation, page 19-6 • Selecting the Devices to Monitor, page 2-2

Strategies for Monitoring, Notification, Mitigation, Remediation, and Audit

STM requires the close coordination of multiple strategies in support of your corporate security policies:

- Monitoring involves the study of network activities and device status to identify anomalous activities or behavior.
- Notification involves alerting those parties responsible for responding to detected anomalies with the information necessary to respond.
- Mitigation involves responding to suspicious activity to prevent the spread of anomalies across your network.
- Remediation involves responding to successful exploits to clean infected hosts on your network.
- Audit involves logging and reporting activities that have taken place during other tasks. The goal of audit is to provide an account the activities and responses to support compliance audits and trend analysis.

The first decision you must make is who will be responsible for mitigation at the selected choke points. Often, organizations separate specialized security devices from the core network infrastructure devices along organizational divisions. As an example, two separate teams, security operations and network operations, may be responsible for different network components or different policies on shared devices. Before you roll MARS out on your network, ownership of a strategies for mitigation must be clearly defined in according with your corporate policies.

When it comes to a mitigation strategy, two options exist:

- You can rely on MARS to identify the choke point and accept the recommended CLI changes to block the detected attack.
- You can issue notifications and incident details to a responsible party who can evaluate the MARS recommendations, but ultimately that party will make the final decision about where and how to stop the detected attack.

Regardless of the option you choose, you should develop guidelines on how long an attack should be blocked, how to investigate an internal attack so that you can clean them, who is responsible for updating the policies after the required quarantine period, and how records of such events should be maintained for audit compliance (for example, is the case management feature of MARS tied to your ticket integration system?).

Next, you should make a distinction in the type of monitoring that you should perform: system monitoring versus security monitoring. *System monitoring* involves monitoring not only the status of the MARS Appliance but also the health and status of the reporting devices and mitigation devices that MARS manages. *Security monitoring* focuses on network and security activity.

For both types of monitoring, you must decide what predefined and custom queries and reports are required, the processes for evaluating and responding to the data they reveal, and guidelines on using the case management features of MARS to manage the responses and track changes.

The last phase involves determining who should be notified when specific incidents are detected. For example, who is notified of device status incidents versus security-related incidents. You must identify your mitigation and remediation personnel, identify those responsible for monitoring (across organizations if necessary), and determine how notifications are to be generated and what they should look like. This involves selecting among methods, including SMS, pager alert, and e-mail, as well as whether the notifications are based on incidents, queries, or reports.

Appliance-side Tuning Guidelines

Tuning on the MARS Appliance focuses on not inspecting traffic that is received from the reporting devices. Two primary techniques exist for appliance-side tuning:

- **Drop rules.** This technique involves dropping all events that match specific criteria received from a reporting device. This technique is the fastest and the least refined. As part of defining a drop rule, you can also specify whether to retain the event log in or simply discard it. The advantage of drop rules is that they events are not processed by any inspection rules, which speeds up the processing of the appliance by reducing the potential workload.
- **Removing devices from inspection.** This technique involves removing a device from inspection rules. This technique is specific to the events that trigger a specific type of alarm. The advantage of this technique is that it does not drop all events that match specific criteria received from a reporting device. In other words, your focus is on reducing a specific false positive rather than all incidents that are fired based on the events. In addition, the events are retained so that you can review them using queries and reports.

When using either of these techniques, remember that when you add or modify a rule, you must click **Activate** before the changes take effect.

Device Inventory Worksheet

The device inventory worksheet will help you collect the required information about the devices on your network. It includes the following information:

- **Device name.** Identifies the well-known name of the device. Typically, this name is the DNS name of the device. MARS uses this name in the topology graph, reports, and events.
- **Reporting IP address.** Identifies the IP address assigned to the network interface from which MARS will be receiving events. This address is used by MARS to map back to the device name and to uniquely identify messages and events originating from the device.
- **Management IP address.** Identifies the IP address assigned to the network interface to which MARS connects to discover the configuration settings for the device.
- **Username/password.** Identifies the account that has the correct authorization to connect to the management IP address and read or write information based on the role in the network. For reporting devices, this account must have privileges sufficient for MARS to read the existing configuration. For mitigation devices, specifically layer 2 switches, this account can enable MARS to publish actual CLI changes to the device to block detected attacks.
- **Role in system/segment.** Identifies whether this device is a reporting device or a mitigation device. It can also identify supporting devices, such as DNS and e-mail servers. In addition, the role should take into account this device's expected importance relative to the network segment, specifically relative to the other devices on the same segment. You can qualify this segment-level role using terms that fit your overall monitoring strategy, such as primary source, second opinion, attack identification, false positive assessment, session data, and endpoint/MAC address identification. Understanding the role that a device can or should play at a network segment level helps prioritize the required and tunable log settings.
- **Required protocols.** Identifies the protocols that this device uses to operate. The primary focus is on the management protocols, notification protocols, and protocols used to publish audit events.
- **Log settings/SNMP RO community string.** Identifies the specific settings with respect to event and log generation that are required for this device to satisfy the role that it will play in the MARS system. It also identifies the SNMP RO community string for this device.
- **Tunable.** Identifies whether you can perform device-side tuning of the log generation.
- **Notify.** Identifies whether this device can receive notifications from MARS.
- **Notification format.** Identifies the format for any notifications that are sent to this device.

[illegible]

User Role Worksheet

The user role worksheet will help you collect required information about administrators of your network. It includes the following information:

- **User Name.** Identifies the user by name.
- **User Role.** Identifies the role this user has with respect to your corporate security policies.
- **MARS Account.** Identifies the MARS user account and role—admin, security analyst, operator, or notification only—which determines access privileges in the web interface. For accountability and security, each user typically has a unique account. However, you can define group-based accounts.
- **Notification Settings.** Identifies the information required to contact this user when incident rules are fired. For users, notification settings include e-mail, pager messages, or SMS messages. For response teams, you may use group aliases. Users should be notified when inspection rules fire and scheduled reports are generated.
- **Device Ownership.** Identifies the reporting devices and mitigation devices on your network for which the user is responsible. This list is especially important when the user is a member of your mitigation or remediation team.
- **Inspection Rules.** Identifies any inspection rules required to meet the needs of this user role. These rules must to be defined or configured to notify the user when they fire.
- **Reports/Queries.** Identifies any reports and queries required to meet the needs of this user role. You must ensure that the user can access these reports and queries. Optionally, you may want to notify the user when scheduled reports are generated.

[illegible]



Reporting and Mitigation Devices Overview

After you complete the initial configuration of Local Controller as described in *Install and Setup Guide for Cisco Security Monitoring, Analysis, and Response System*, you must determine a monitoring strategy to use for your network. You must also determine a mitigation strategy, if you chose to take advantage of the MARS mitigation features. For guidance on how to determine the monitoring and mitigation strategies, see [STM Task Flow Overview, page 1-1](#).

This chapter assumes that you have made corporate-level policy decisions and that you are executing against the [Checklist for Provisioning Phase, page 1-2](#). This chapter provides the following:

- Guidance on selecting and configuring reporting devices and mitigation devices
- Discussion of the levels of operation that MARS supports
- Guidance on selecting a method for adding devices to Local Controller
- Discussion of those features that enable rich data collection

It contains the following sections:

- [Levels of Operation, page 2-1](#)
- [Selecting the Devices to Monitor, page 2-2](#)
- [Understanding Access IP, Reporting IP, and Interface Settings, page 2-8](#)
- [Selecting the Access Type, page 2-10](#)
- [Bootstrap Summary Table, page 2-12](#)
- [Adding Reporting and Mitigation Devices, page 2-16](#)
- [Data Enabling Features, page 2-28](#)
- [Integrating MARS with 3rd-Party Applications, page 2-43](#)

Before configuring the MARS Appliance to recognize reporting devices, you should understand the three levels of operation that MARS can achieve.

Levels of Operation

MARS operates at three discernible levels based on the type of data collected from reporting devices and the features such data enables for the system. These levels focus on the ability to identify attacks from end-to-end, and they are separate from the features enabled by specific types of reporting devices.

- **Basic.** At this level, MARS behaves like a smart syslog server. It collects reporting device logs and support basic queries and reports. To enable basic operation, you must complete the initial configuration of the MARS Appliance as described in *Install and Setup Guide for Cisco Security*

Monitoring, Analysis, and Response System. In addition, you must specify the device name and reporting IP addresses of the reporting devices as described in [Adding Reporting and Mitigation Devices, page 2-16](#).

- **Intermediate.** At this level, MARS processes events and performs session-based correlation, including resolving NAT and PAT translations at the IP address layer. To enable intermediate operation, you must provide more details about the devices you want to monitor, including access IP addresses, management access passwords, OS platforms and versions, and running services and applications, see [IP Management, page 23-3](#) for more information.
- **Advanced.** This level is a fully enabled MARS Appliance. When advanced operation is enabled, MARS Appliance discovers and displays the full topology, draws attack paths, and enables MAC address lookups of the hosts involved in an attack. To enable advanced operation, you must provide the SNMP community string information for your network. You must also enable topology discovery, as defined in [Scheduling Topology Updates, page 2-39](#).

Table 2-1 summarizes the levels, their configuration requirements, and the features enabled at that level.

Table 2-1 **Levels of Operation**

Level Of Operation	Configuration Requirements	Functionality Enabled
Level 1	MARS configured Reporting device names and reporting IP addresses added NetFlow enabled	Basic syslog functionality Event correlation Query, reports, and chart support NetFlow anomaly detection
Level 2	Access IP addresses and information added	Starts performing event and session-based correlation NAT and PAT resolution IP address lookup of attackers and targets
Level 3	Community strings and networks added	MAC address lookup of attackers and targets Topologies enabled

Selecting the Devices to Monitor

All monitoring strategies involve selecting the types of devices to monitor and how much data to provide the MARS Appliance. All devices on your network, be they hosts, gateways, security devices, or servers, provide some level of data that MARS can use to improve the accuracy of security incident identification. However, careful consideration of what data to provide can improve the attack identification response time by ensuring that MARS does not perform necessary or redundant event correlation and analysis. Unnecessary logging and reporting by reporting devices can also reduce the effectiveness of your network.

We recommend analyzing each network segment to identify the most data rich combination that you can achieve, while identifying and refining your configurations to reduce redundant data.

When determining a monitoring strategy, you must also determine the goals behind the monitoring. Is it just for attack detection? Attack detection and mitigation? Regulatory compliance? Your goals affect which devices you must monitor and what features you must configure on those devices.

Consider distinct goals:

- Attack detection

- Attack detection and mitigation
- Regulatory compliance
- Full NAC awareness
- Identify the devices/feature pairs that overlap on the same network segment, where a choice between device can reduce duplicity or prioritize device performance

Last, you must consider an event tuning method for your monitoring strategy. How you tune your MARS affects your overall operational costs proportionally to the number of device of a give type that are monitored. Essentially, if you have the bandwidth available, we recommend that you tune the events at the MARS Appliance, which reduces your operational costs by tuning at a single point in the network. However, if bandwidth is a precious commodity, you may chose to tune the event propagation at the reporting device level, preventing the events from going onto the network.

[Table 2-2](#) identifies the device types, describes what information they can provide, and recommends how to configure these devices within your network.

Table 2-2 **Device Types and Data Available**

Device Type	Data Available	Recommended Configurations
Router	The device discovery protocol is the one used for administrative access/mitigation. For example, if SSH is used to discover the device, then SSH is the protocol that used to pushed the mitigation command. The following data is pulled from routers: • hostname • static routes • ACL rules • static NAT rules • traffic flows • SNMP RO Community strings • NetFlow data • device status and resource utilization, such as memory, CPU, and interface/port statistics. • ARP cache table. Used to map IP address to MAC address.	Enable the following: • SNMP RO community strings • Syslog traffic • Device discovery via SSH or Telnet access
Switch	During investigation and mitigation, the ARP cache tables are reviewed to resolve the MAC addresses involved in the incident. This data is cached for 6 hours. SNMP RO Community strings Forwarding tables, used to map IP address to MAC address. Device status and resource utilization, such as memory, CPU, and interface/port statistics. NetFlow data 802.1x logs generated during NAC sessions	Enable the following: • SNMP RO community strings • Syslog traffic • Device discovery via SSH or Telnet access • Enable NetFlow data • Administrative access for mitigation push

Table 2-2 *Device Types and Data Available (continued)*

Device Type	Data Available	Recommended Configurations
Firewall	Interface configurations. Used to populate topology view and determine expected routes, which helps refine correlation of traffic traversing the firewall. NAT and PAT mappings. Used to identify the point of origin attackers and targets and trace attacks as they spread. Firewall policies. When discovering ASA, PIX, and FWSM, MARS parses ACLs and conduits (PIX only). For Check Point firewalls, it collects the firewall policy from policy table. MARS using this information only for path computation and mitigation recommendations. It is not used by any other components, such as rules, reports, and sessionization. Firewall logs. Accepted and denied sessions logs are used to identify false positives and determine if potential attacks were blocked before reaching their targets. Audit logs. Associates users with authentication sessions and assists in identifying exploited accounts and administrative sessions. ARP cache tables. Used to map IP address to MAC address. Device status and resource utilization information. Used to identify anomalous network activities based on memory, CPU, and interface and port statistics.	Enable the following: • SNMP RO community strings • Syslog messages • Device discovery
VPN	Remote user information. Provides username to IP address mapping. VPN client helps determine the person who logged in and performed specific actions. Clarifies the true point of origin by identifying the host, not the VPN concentrator. Login/logout records. Helps identifies worms by tracing outbreaks back to a specific user and provides network access periods. Device status information. Identifies whether the device is operational, which allows prediction of possible spread of potential attacks and worms. • SNMP RO Community strings	

Table 2-2 **Device Types and Data Available (continued)**

Device Type	Data Available	Recommended Configurations
Network IDS/IPS	Fired signature alerts. Identifies attacks and threats, which helps determine mitigation response, identify potential false positive information, and target vulnerability assessment probes conducted by MARS. Trigger packet information. Provides the payload of the packet that caused the signature to fire. Determine whether an attack was blocked at a specific device. Device status information	
Host IDSes	Provides host-level validation of exploits and blocked attacks, which improves the accuracy of false positive identification, which in turn improves the ability of the administrator to accurately prioritize the work required to contain attacks.	
Anti-Virus	Central anti-virus management servers provide information on which hosts are infected, which hosts have reported attempted infections, etc. The servers also provide the dat or signature file information for managed hosts, which improves the ability to determine whether an attack was likely to have succeeded.	
Vulnerability Assessment	Host OS and Patch Level. When a signature fires on an IDS and it is reported to MARS, MARS can either launch a targeted scan using Nessus, or query a vulnerability assessment system that helps determine whether the target was vulnerable.	Enable any vulnerability assessment solutions supported by MARS.

Table 2-2 *Device Types and Data Available (continued)*

Device Type	Data Available	Recommended Configurations
Host OSes	Microsoft Windows Hosts Events found in the security event log as well application event and system event log.	Install and configure SNARE, which pushes events to MARS in near real time, and scales more efficiently than pulling events from hosts.
	Solaris and Linux Hosts Incoming network session logs, via inted, and FTP transfer logs, via xferlog. In addition, any events that are written to the system log by applications and services running on host.	• Enable logging for the xferlog and inetd applications. • Enable syslog daemon. • Identify the MARS Appliance as syslog target.
	Generic Hosts (All OSes) Includes system-level information, such as privilege escalation and buffer overflow. Helps determine what attacks make it to the host layer. If MARS learns of activity at the host level, then it understands that the attack or exploit has successfully traversed the network. MARS correlates this data with the network level data to discover the whole incident and analyze the exploit method so the administrator can build a better defense. In some cases, MARS recommends actions for mitigating the attack. We recommend that you maintain these recommended blocks as long as similar attacks are expected. Typical blocking techniques, such as IPS shunning, often fail to identify the best chokepoint for containment. As part of the recommended action, MARS does identify the optimal chokepoint where the recommended action should be effected.	
Web Server	Same as hosts (SNARE and Perl script agents) need this when the hosts cannot send us the logs via syslog. agent is basically a transport.	
Web Proxy	Mapping from user to site, translations for the IP address mapping, tells us the real address of the host who is likely infected. URLs and also filtering...regulatory compliance.	
Database	Login/logout to determine the actual user (query report tab on the data). Privilege escalation, brute force crack type stuff, or maybe we want to do regulatory compliance.	

Table 2-2 **Device Types and Data Available (continued)**

Device Type	Data Available	Recommended Configurations
AAA Server	Login/logout and NAC functionality (deny a person due to privileges, it triggers NAC message) - passed authentication log - failed attempts log - RADIUS accounting log, including those events specific to NAC.	Supporting Cisco Secure ACS Server, page 14-2 Supporting Cisco Secure ACS Solution Engine, page 14-2
Generic Syslog	Same as host, provides support for additional customer devices.	
Generic SNMP	Same as host, provides support for additional customer devices.	
Cisco Security Manager	Mapping to any committed policy rules defined in Security Manager that match any ACL rules that could cause the generation of a specific syslog event by a reporting device. This policy lookup feature allows you to debug network issues and understand the cause/effect relationships between event messages and the device policies and traffic that resulted in the generation of the event.	Enable HTTPS on the Security Manager server. Define an administrative level account on the Security Manager server that CS-MARS can use for policy lookups.

Understanding Access IP, Reporting IP, and Interface Settings

When defining a reporting or mitigation device in the web interface, MARS allows (and at times, requires) you to specify several IP addresses. Understanding the purpose of the different addresses is important to effectively defining the devices that you want to monitor and manage. It is also important to understand their relationship to other settings that you can identify.

If a device has a single interface and a single IP address associated with that interface, the access and reporting IP addresses are the same as the address assigned to the interface. MARS collects this information separately to support those devices that have multiple interfaces, multiple IP addresses associated with a single interface, or both.



Note

Not all reporting devices support both an access and reporting IP address. Some devices use only access IP addresses to query the device for the required information (e.g., QualysGuard security service), while others have no settings that MARS can discover and only generate event messages for MARS to process (e.g., NetCache appliances). In addition, not all devices require the definition of interfaces.

This section discusses the following three addresses and their relationship to other settings:

- [Access IP, page 2-9](#)
- [Reporting IP, page 2-9](#)
- [Interface Settings, page 2-10](#)

Access IP

MARS uses the access IP address to either connect to the device for network-based administrative sessions or connect to a remote server on which a file containing the device's configuration is stored. The expected value is determined by the access type you select. Most devices also require that you explicitly identify the IP addresses of hosts allowed to administer them. The MARS Appliance must be listed among such hosts as part of the device preparation.

The protocol that MARS uses to connect to the device is defined by the access type value, which is a dependency for enabling administrative access. Once MARS has administrative access, it can perform device discovery, which includes settings such as ARP tables, NAT, routes, and active ACLs, all of which helps MARS understand the topology, perform attack path analysis, and identify false positive incidents. Discovery can be performed to varying degrees using any of the access types. For more information on access types, see [Selecting the Access Type, page 2-10](#).

MARS also uses SNMP RO and SNMPwalk to discover the device settings and topology information. However, the two methods of discovery are distinct and have distinct requirements. SNMPwalk requires the access IP address and the SNMP access type. SNMP RO discovery does not require the SNMP access type, but it does require the access IP address.

**Note**

MARS does not support the following characters in the SNMP RO community string: ' (single quote), " (double quote), < (less than symbol), and > (greater than symbol).

In addition, both SNMPwalk and SNMP RO are unrelated to SNMP notifications or SNMP traps. SNMPwalk and SNMP RO both require that MARS initiate the information request, whereas SNMP notifications are event notifications published by the reporting device, much the same as syslog messages are. As with syslog messages, SNMP notifications are published over the reporting IP address.

Reporting IP

The reporting IP is the source IP address of event messages, logs, notifications, or traps that originate from the device. MARS uses this address to associate received messages with the correct device. For single-homed devices, the reporting IP address is the same as the access IP; for dual- or multi-homed devices, this address must be explicitly associated with the syslog, NetFlow, and SNMP services running on the reporting device. Most devices also require, for each message type, that you explicitly identify the IP addresses of hosts to which messages should be published. These hosts are commonly referred to as target log servers. The MARS Appliance must be listed among such hosts as part of the device preparation.

The role in MARS of the reporting IP address differs from that of the access IP address in that the reporting IP address is treated passively from the MARS perspective. MARS does not query the device using this address. Such operations are performed using the access IP address and the access type.

MARS accepts only one reporting IP address per device. For devices supporting two message formats, such as NetFlow and syslog, you must ensure that both message formats are bound to the same source IP address (the reporting IP). In Cisco IOS devices, this common association is not the default so you must change either the syslog or the NetFlow reporting IP address to match the other. If the message types do not originate from a common IP address, one of them is seen as originating from an unreported device and MARS does not parse those events correctly.

The supported format of event data varies among reporting devices. Just because the device is able to generate syslog, NetFlow, and SNMP notifications does not mean that MARS processes all three formats. The document, [Supported Devices and Software Versions for Cisco Security MARS Local Controller 4.2.x](#), identifies the event retrieval protocol supported by each device type.

Interface Settings

Interface settings are exclusive to hosts and software applications running on hosts. While MARS can discover the settings of a reporting device that is a software application running on a host, it cannot discover settings about the host itself. The role of interface settings in MARS is different from that the access IP address and reporting IP address. Interface settings represent static information, not discovered or learned, about the host.

When correlating events specific to a host or reporting devices running on that host, MARS needs to understand the number of interfaces installed in the host, their names, and the IP addresses and networks associated with them. MARS uses the interface settings to guide discovery operations, to determine attack path vectors, and to perform Nessus vulnerability assessments.

Selecting the Access Type

The access type refers to the administrative protocol that MARS uses to access a reporting device or mitigation device. For most devices monitored by MARS, you can choose from among four administrative access protocols:

- **SNMP.** SNMP access provides administrative access to the device using a secured connection. It allows for the discovery of the settings using SNMPwalk, such as routes, connected networks, ARP tables, and address translations. If granted read-write access, SNMP also allows for mitigation on any L2 devices that support MIB2.
- **Telnet.** Telnet provides full administrative access to the device using an unsecured connection. It allows for the discovery of the settings, such as routes, connected networks, ARP tables, and address translations. It also allows for mitigation on L2 devices.
- **SSH.** SSH provides full administrative access to the device using a secured connection. It allows for the discovery of the settings, such as routes, connected networks, ARP tables, and address translations. It also allows for mitigation on L2 devices. This access method is recommended for DTM support; however, Telnet access can achieve the same results.
- **FTP.** FTP passive discovery of settings by providing MARS access to a file copy of the configuration running on the router. FTP does not support mitigation, DTM, or discovery of dynamic settings, such as NAT and ARP tables. In addition, if you select the FTP access type for device types, such as Cisco ASA and FWSM, you can only discover settings for the admin context. This access method is the least preferred and most limited access method. To enable configuration discovery using FTP access, you must place a copy the device's configuration file on an FTP server to which the MARS Appliance has access. This FTP server must have users authentication enabled.



Note TFTP is not supported. You must use an FTP server.

You can use any access scheme in conjunction with an SNMP RO community string. The division between Access IP and Reporting IP is clearly illustrated by an FTP access type example. Assume that you have SNMP RO access to a router, but your configuration discovery (access type) is restricted to a file stored on an FTP server.

When you define a device in MARS, the Access IP is the IP address of the FTP server (not the router), and the authentication information is used to access the FTP server. The Access Method is set to FTP. The Reporting IP is the IP address of the interface over which SNMP traps are published by the router.

The following topics describe how to configure each access type, identifying the fields that should be completed when a specific access type is selected. For efficiencies sake, these procedures are referenced throughout the specific device configuration topics, as they related to a specific device type.

- [Configure SNMP Access for Devices in MARS, page 2-11](#)
- [Configure Telnet Access for Devices in MARS, page 2-11](#)
- [Configure SSH Access for Devices in MARS, page 2-12](#)
- [Configure FTP Access for Devices in MARS, page 2-12](#)

Configure SNMP Access for Devices in MARS

This procedure assumes you are defining a reporting device or mitigation device and that you were referred to this procedure after selecting SNMP in the Access Type list. To select SNMP as the access type, you must provide MARS with SNMP read-write access.

**Note**

The SNMP access type is not required to enable the SMPO RO strings. In fact, no access type is required to support SNMP RO. SNMP RO uses a shared, read-only community string; it does not require a read-write community string as does the SNMP access type.

If you selected SNMP as the access type, follow these steps:

-
Step 1	In the Login field, enter the username of the administrative account to use when accessing the reporting device.
Step 2	In the Password field, enter the password associated with the username specified in the Login field.
Step 3	If this device supports an enable mode, enter that password in the Enable Password field.
-

Configure Telnet Access for Devices in MARS

This procedure assumes you are defining a reporting device or mitigation device and that you were referred to this procedure after selecting TELNET in the Access Type list.

If you selected TELNET as the access type, follow these steps:

-
Step 1	In the Login field, enter the username of the administrative account to use when accessing the reporting device.
Step 2	In the Password field, enter the password associated with the username specified in the Login field.
Step 3	If this device supports an enable mode, enter that password in the Enable Password field.
-

Configure SSH Access for Devices in MARS

This procedure assumes you are defining a reporting device or mitigation device and that you were referred to this procedure after selecting SSH in the Access Type list.

If you selected SSH as the access type, follow these steps:

-
- Step 1** From the list box to the right of the Access Type list, select **3DES**, **DES**, or **BlowFish** as the encryption cipher for SSH sessions between the MARS Appliance and the reporting device.
 - Step 2** In the Login field, enter the username of the administrative account to use when accessing the reporting device.
 - Step 3** In the Password field, enter the password associated with the username specified in the Login field.
 - Step 4** If this device supports an enable mode, enter that password in the Enable Password field.
-

Configure FTP Access for Devices in MARS

This procedure assumes you are defining a reporting device or mitigation device and that you were referred to this procedure after selecting FTP in the Access Type list.

If you selected **FTP** as the access type, follow these steps:

-
- Step 1** In the Login field, enter the username of the FTP server account to use when accessing the configuration file of the reporting device.
 - Step 2** In the Password field, enter the password associated with the username specified in the Login field.
 - Step 3** In the Config Path field, enter the path to the reporting device’s configuration file residing on the FTP server.

This path begins at the root of the FTP server’s published folder, not at the root directory of server.
 - Step 4** In the File Name field, enter the name of the reporting device’s configuration file residing on the FTP server.



Note

If you select **FTP**, you cannot enter an enable password.

Bootstrap Summary Table

[Table 2-3](#) summaries the settings that you must configure for reporting devices and mitigation devices. It also provides links to any required agent downloads and to detailed configuration information.

Table 2-3 Reporting and Mitigation Device Bootstrap Summary

Device Type/Name	Bootstrap Summary	Reference Information
Router/Switch		
Cisco Router	1. Access to IP address/interface by MARS.	Cisco Router Devices, page 3-1
Cisco Switch (IOS)	2. FTP, SNMP, Telnet or SSH access by MARS.	Cisco Switch Devices, page 3-9
Cisco Switch (CatOS)	3. Define SNMP RO community string. 4. Turn on syslog, define log level, and define MARS as target of syslog messages. 5. Enable NAC features.	
Extreme ExtremeWare	1. Access to IP address/interface by MARS.	Extreme ExtremeWare 6.x, page 3-17
Generic Router	2. (ExtremeWare only) Turn on syslog, define log level, and define MARS as target of syslog messages. 3. SNMP access by MARS. 4. Define SNMP RO community string.	Generic Router Device, page 3-18
Firewall Devices		
Cisco PIX	1. Access to access and reporting IP address/interface by MARS.	Bootstrap the Cisco Firewall Device, page 4-2
Cisco Adaptive Security Appliance (ASA)	2. FTP, Telnet, or SSH access by MARS.	
Cisco Firewall Services Module (FWSM)	3. Define SNMP RO community string. Note SNMP settings should be defined for the admin context on ASA and FWSM. You do not need to define these settings for each security context. 4. Turn on syslog, define log level, and define MARS as target of syslog messages.	
Cisco IOS Firewall Feature Set		
Juniper Netscreen		NetScreen ScreenOS Devices, page 4-11
Checkpoint Opsec NG and Firewall-1	1. Add the MARS Appliance as a host.	Bootstrap the Check Point Devices, page 4-22
Nokia Firewall (running Checkpoint)	2. Create and install an OPSEC Application object for the defined host.	
	3. Define policies to permit SIC traffic between the MARS Appliance, the Check Point management server, and any remote servers.	
	4. Define the log settings to push the correct events to the defined host.	
	5. Install the policies.	
VPN Devices		
Cisco VPN Concentrator		Cisco VPN 3000 Concentrator, page 5-1

Table 2-3 *Reporting and Mitigation Device Bootstrap Summary (continued)*

Device Type/Name	Bootstrap Summary	Reference Information
Network IDS		
Cisco Network IDS Cisco IDSM	1. Enable RDEP for IDS modules. 2. Configure the following signature actions: • Alert • (Optional) To view trigger packets, enable the “produce-verbose-alert”. • (Optional) To view IP logs, enable the alert or “produce-verbose-alert” and “log-pair-packets”.	Cisco IDS 3.1 Sensors, page 6-1 Cisco IDS 4.0 and IPS 5.x Sensors, page 6-5
Cisco Intrusion Prevention System (IPS), Network IPS	1. Enable SDEE for IPS modules. 2. Configure the following signature actions: • Alert • (Optional) To view trigger packets, enable the “produce-verbose-alert”. • (Optional) To view IP logs, enable the alert or “produce-verbose-alert” and “log-pair-packets”.	Cisco IDS 4.0 and IPS 5.x Sensors, page 6-5
Cisco IPS ASA module	1. Enable SDEE for IPS modules. 2. Configure the following signature actions: • Alert • (Optional) To view trigger packets, enable the “produce-verbose-alert”. • (Optional) To view IP logs, enable the alert or “produce-verbose-alert” and “log-pair-packets”.	Cisco IPS Modules, page 6-9
Cisco IOS IPS module	1. Enable SDEE for IPS modules. 2. Configure the following signature actions: • Alert • (Optional) To view trigger packets, enable the “produce-verbose-alert”. • (Optional) To view IP logs, enable the alert or “produce-verbose-alert” and “log-pair-packets”.	Cisco IPS Modules, page 6-9
McAfee Intrushield		IntruVert IntruShield, page 6-22
Juniper Netscreen IDP		NetScreen IDP 2.1, page 6-31
Symantec Manhunt		Symantec ManHunt, page 6-29
ISS RealSecure		ISS RealSecure 6.5 and 7.0, page 6-17
Snort		Snort 2.0, page 6-28
Enterasys Dragon		Enterasys Dragon 6.x, page 6-33

Table 2-3 Reporting and Mitigation Device Bootstrap Summary (continued)

Device Type/Name	Bootstrap Summary	Reference Information
Host IDS		
Cisco Security Agent		Cisco Security Agent 4.x Device, page 7-5
McAfee Enterecept		Entercept Entercept 2.5 and 4.0, page 7-1
ISS RealSecure Host Sensor		ISS RealSecure 6.5 and 7.0, page 6-17
Anti-virus		
Symantec AntiVirus		Symantec AntiVirus Configuration, page 8-1
Cisco Incident Control System (Cisco ICS), Trend Micro Outbreak Prevention Service (OPS)		Cisco Incident Control Server, page 8-13
McAfee ePolicy Orchestrator		McAfee ePolicy Orchestrator Devices, page 8-8
Network Associates VirusScan		McAfee ePolicy Orchestrator Devices, page 8-8
Vulnerability Assessment		
eEye REM		eEye REM 1.0, page 9-3
Qualys QualysGuard		Qualys QualysGuard Devices, page 9-5
Foundstone Foundscan		Foundstone FoundScan 3.0, page 9-1
Host Operating Systems		
Windows	Do one of the following: - Install and configure the SNARE agent - Create or edit an administrative account to ensure that it has permissions to pull the event data	Syslog (pushed by SNARE agent) or event data pull using MS-RPC Push Method: Configure Generic Microsoft Windows Hosts, page 10-5 Pull Method: Configure the Microsoft Windows Host, page 10-6
Solaris	—	Syslog (from Device) Sun Solaris and Linux Hosts, page 10-2
Redhat Linux	—	Syslog (from Device) Sun Solaris and Linux Hosts, page 10-2
Web Server		
Microsoft Internet Information Server	—	Syslog (from SNARE agent) Install and Configure the Snare Agent for IIS, page 12-1
Sun iPlanet	—	HTTP (from MARS Agent) Install and Configure the Web Agent on UNIX or Linux, page 12-7

Table 2-3 *Reporting and Mitigation Device Bootstrap Summary (continued)*

Device Type/Name	Bootstrap Summary	Reference Information
Apache	—	HTTP (from MARS Agent) Install and Configure the Web Agent on UNIX or Linux, page 12-7
Web Proxy		
NetApp NetCache	—	HTTP Network Appliance NetCache Generic, page 13-1
Database Server		
Oracle	TCP	SQLnet (from Host) Oracle Database Server Generic, page 11-1
AAA Server		
Cisco Secure Access Control Sever (ACS)	—	Syslog (from MARS Agent) Install and Configure the PN Log Agent, page 14-7 (Cisco Secure ACS)
Cisco Secure ACS Appliance	Install and configure remote log agent.	Syslog (from MARS Agent) on secondary host Supporting Cisco Secure ACS Solution Engine, page 14-2 Install and Configure the PN Log Agent, page 14-7 (Cisco Secure ACS)
SNMP and Syslog Servers		
Generic Syslog Server	Publish syslog messages to MARS Appliance. Enable SNMP access by MARS Appliance.	Adding Generic Devices, page 10-1
Generic SNMP Server	Enable SNMP access by MARS Appliance.	Adding Generic Devices, page 10-1
Other		
Cisco Security Manager	Enable HTTPS access by MARS Appliance	Checklist for Security Manager-to-MARS Integration, page 16-6 Bootstrapping Cisco Security Manager Server to Communicate with MARS, page 16-12 Add a Cisco Security Manager Server to MARS, page 16-13

Adding Reporting and Mitigation Devices

Three methods exist for adding reporting devices and mitigation devices to MARS:

- Manually add the devices one at a time.
- Add multiple devices using a seed file.
- Discover devices automatically using SNMP RO community strings.

From the Security and Monitor Devices page, you can add or edit the reporting devices and mitigation devices that MARS monitors. To access this page, click **Admin > System Setup > Security and Monitor Devices**. You can search for, add, edit, delete, change display status, and load devices from the seed file.

The device support is categorized into three categories:

- **HW-Based Security Devices.** Hardware-based devices represent routers, switches, and other dedicated security appliances. You can add such reporting devices by selecting the appropriate device.
- **SW-Based Security Devices.** Software-based devices represent applications that reside on a host, rather than a dedicated appliance. You can add reporting device on a new host by selecting **Add SW security apps on new host** or on an existing host by selecting **Add SW security apps on existing host**.
- **On-Demand Security Services.** Security services represent subscription-based services provided by vendors using a central security operations center (SOC) with remote monitoring nodes. These services, such as Qualys QualysGuard, represent systems from which MARS periodically pulls data. You can add such reporting devices by selecting the appropriate service. These devices also require you to define a schedule for pulling data (see [Scheduling Topology Updates, page 2-39](#)).

The complete list of supported devices is presented in the *Supported Devices and Software Versions for Cisco Security MARS Local Controller 4.2.x* document. Devices are added to this list on an ongoing basis via software upgrade packages. See *Install and Setup Guide for Cisco Security Monitoring, Analysis, and Response System* for details on how to upgrade your MARS Appliance.

MARS can also support any syslog or SNMP devices, even if they do not appear on the list of devices supported by the MARS. You can enter any syslog or SNMP device into the network topology, configure it to report data to the MARS, and query it using a free-form query. (See [To Run a Free-form Query, page 20-2](#).)

For more information on adding devices, see:

- [Add Reporting and Mitigation Devices Individually, page 2-17](#)
- [Add Multiple Reporting and Mitigation Devices Using a Seed File, page 2-20](#)
- [Adding Reporting and Mitigation Devices Using Automatic Topology Discovery, page 2-25](#)

Regardless of the method that you have used to add the devices, you should also perform the following tasks:

- [Verify Connectivity with the Reporting and Mitigation Devices, page 2-26](#)
- [Activate the Reporting and Mitigation Devices, page 2-27](#)

Add Reporting and Mitigation Devices Individually

In general, you have two choices for adding devices that you want to monitor into your MARS. You can create a seed file or you can add each device manually. Seed file support is limited to a few device types, see [Column E, page 2-23](#) for the devices supported.

When manually configuring devices, select the devices that are most interesting to you. Once added, you can come back and edit them as necessary. Manual configuration is also useful when you add or change a single security device on your network.

**Note**

Remember that you do not have to add all of the devices configuration information at once. You can start by adding the device's name and its access IP address. You can always return later, when the MARS starts to report to you, and provide more details.

To add a device manually, follow these steps

-
- Step 1** Click **Admin > System Setup > Security and Monitor Devices > Add**.
 - Step 2** Select the device from the list.
 - Step 3** Enter the information needed to communicate with the device.
 - Step 4** Click **Submit**.
 - Step 5** Once add a device, you must click **Activate** for MARS to correctly process events received from that device. For more information, see [Activate the Reporting and Mitigation Devices, page 2-27](#).
-

Edit a Device

-
- Step 1** Check the box next to the device.
 - Step 2** Edit the device's settings.
 - Step 3** Click **Submit**.
-

Upgrade the Device Type to a Newer Version

You can change the Device Type version setting of a hardware-based security device. You cannot upgrade the version for software applications running on a host. To upgrade the software appliance version, you must remove the application from the host and add the newer one.

This version change feature applies only to device types with the same vendor and model but different versions. Specifically, you can change the version for the following device types:

- Cisco IDS
- Cisco PIX
- Cisco VPN Concentrator
- NetScreen ScreenOS

For example, you could change the settings for the device type Cisco PIX 6.1 to Cisco PIX 7.0 without having to delete the device and add it again. The benefit of matching the version setting to the deployed device is that it allows MARS to correlate any event types introduced in the more recent version. It also allows you to incrementally upgrade your reporting devices without having to worry about when to add that reporting device to MARS. The events that are correlated under one device type will be associated with the newer device type version when you make the change in MARS.

To change the version of a device, follow these steps:

-
- Step 1** Click **Admin > System Setup > Security and Monitor Devices**.
- Step 2** Select the checkbox to the left of the device for which you want to change the version, and click **Change Version**.
- The Change the Device Type Version page appears, displaying the device name, vendor, model, and old version type information.
- Step 3** Select the new version in the New Device Type Version list.
- Step 4** To change the version of the device to the new version, click **Submit**.
- If any additional changes are available due to the version change, the Edit page appears.
- Step 5** If the Edit page appears, make any desired changes and click **Submit**.
- Step 6** Once you change the version setting for a device, you must click **Activate** for MARS to correctly process events received from that device. For more information, see [Activate the Reporting and Mitigation Devices, page 2-27](#).
-

Delete a Device

When you define a reporting device in MARS, this device is added in two separate pages of the web interface. It appears where you have defined it, on the Admin > Security and Monitoring Devices page, as well as under the general device identification page under Management > IP Management. This duplication of content is based on the different functions that each of these pages serves.

The Security and Monitoring Devices page configures the contact and device type information, whereas the IP Management page is used by the parser module to correlate known devices versus unknown devices. Typically when you delete a device from the Security and Monitoring Device page, you still want to retain the knowledge of that device in MARS so that historical incidents and events and cases can resolve to a known device; therefore, the device is not deleted from the IP Management page.



Note

Deleting a device does disassociate any historical incidents and events from the IP address. In other words, once you delete the device, you will not be able to find historical events for that device even if you re-add the device at a later date.

However, if you need to delete and re-add a device to MARS, you must delete the device from both pages before you attempt to re-add the device.

In addition, as devices are discovered on your network, they are added to the list of devices in the IP Management page. If you want to add a reporting device and find that you cannot, review the list of devices in the IP Management page to ensure that the device has not been auto-populated. If it has, you must first delete that device, then you can add it as a reporting device on the Security and Monitoring Devices page.

To delete a device, follow these steps:

-
- Step 1** Select one of the following pages:
- Admin > Security and Monitoring Devices
 - Management > IP Management
- Step 2** Check the box next to each device you want to delete.

- Step 3** Click **Delete**.
- Step 4** On the confirmation page, click **Submit**.
- The device is deleted from the selected page.

Delete All Displayed Reporting Devices

You can perform this procedure from the Admin > Security and Monitoring Devices page.

To delete all devices displayed on a page, follow these steps:

- Step 1** On the Admin > Security and Monitoring Devices, select the checkbox to the left of the Device Name column heading at the top left of the table.
- All displayed devices are selected.
- Step 2** Click **Delete**.
- A page appears prompting you to confirm that you want to delete the list of devices.
- Step 3** Click **Submit** to delete all the selected devices.

Add Multiple Reporting and Mitigation Devices Using a Seed File

The seed file is a comma-delimited file with the file extension .csv (comma-separated value). Most spreadsheet programs let you import and export files as CSV files.

The following is a sample seed file as exported from a popular spreadsheet program:

```
10.1.1.1,,,PIX,TELNET,,,cisco,,,,,,,,,
192.168.229.241,,,IOS,TELNET,,,csRv$12*,EcsRv$12$,,,,,,,,,,
10.1.1.83,,,PIX,SSH,pix,Vpnspn12,,vPfwlne,,,,,,,,,
192.168.151.169,,,PIX,SSH,pix,lpt$12,,pot$1*d1,,,,,,,,,
10.4.2.4,,,NETSCREEN,SSH,netscreen,nt*$scn25,,,,,,,,,
10.4.2.3,,,NETSCREEN,SSH,netscreen,nt*$scn10,,,,,,,,,
10.1.1.241,,,IOS,TELNET,,,cisco,cisco,,,,,,,,,
10.4.2.1,,,IOS,TELNET,,,Qa$1*5ft,gt*$j15,,,,,,,,,
10.4.2.2,,,IOS,TELNET,,,Qa$1*5ft,gt*$j15,,,,,,,,,
wanRouter,public,,,IOS,SNMP,,,,,,,,,
myPix63,,,PIX,SSH,pix,test1,,test1234,,,,,,,,,10.2.3.1
MyPc,,,WINDOWS,RPC,myname,mypass,,,,,,,,,
myPix70,,,PIX7X,SSH,,,,,,,,,
myids40,,,CiscoIDS4x,SSL,,,,,,,,,
myids50,,,CiscoIPS5x,SSL,,,,,,,,,
myASA70,,,ASA,SSH,,,,,,,,,
myWindowsNT,,,WindowsNT,RPC,,,,,,,,,
myFWSM23,,,FWSM,SSH,,,,,,,,,
```

With the CSV file, you can enter the values, passwords, and information for each device that you want the MARS Appliance to monitor in its appropriate row and column. While the seed file is useful for getting the MARS Appliance started processing event data for most devices, you may need to use the Admin > System Setup > Security and Monitoring Devices page to fine-tune the device manually. In addition, you must Activate the devices that you add using a seed file (see [Activate the Reporting and Mitigation Devices, page 2-27](#)).

Devices that Require Custom Seed Files

Some reporting devices represent the management consoles for the actual host- or node-based reporting devices. These consoles often represent centralized log servers for the devices they manage. However, for MARS to correctly correlate the logs for these centralized log servers, you must identify those host- or node-based reporting device. In some cases, MARS is able to dynamically learn of the hosts or nodes by parsing the logs. In other cases, you must use a seed file generated by management console to identify each of the managed reporting devices.

Once you generate the seed file, you must import that seed file under the host that represents the management console in the MARS web interface to load the sensor module information from the CSV or seed file. The device types that use a custom seed file are as follows:

- **Entercept.** For more information, see [Extracting Entercept Agent Information into a CSV file \(for Entercept Version 2.5\)](#), page 7-1.
- **IntruVert IntruShield.** For more information, see [Extracting Intruvert Sensor Information from the IntruShield Manager](#), page 6-22.
- **Cisco Security Agent.** While MARS can learn of the CSA agents dynamically, you can also import the initial list of agents using a custom seed file. For more information, see [Export CSA Agent Information to File](#), page 7-6.
- **Symantec AntiVirus.** While MARS can learn of the Symantec AntiVirus agents dynamically, you can also import the initial list of agents using a custom seed file. For more information, see [Export the AntiVirus Agent List](#), page 8-7.

Devices that Require Updates After the Seed File Import

When you add specific reporting devices using a seed file, you must edit them to complete the definition of the device before you can monitor them. Typically, these devices are IDS/IPS devices that monitor specific networks. The device types that you must update are as follows:

- **Cisco IDS 4.x Devices.** These sensors are defined by importing a MARS-specific seed file as defined in [Load Devices From the Seed File](#), page 2-24. However, once you import a sensor, you must identify the monitored networks that it monitors. For more information, see [Specify the Monitored Networks for Cisco IPS or IDS Device Imported from a Seed File](#), page 6-8.
- **Cisco IPS 5.x Devices.** These sensors are defined by importing a MARS-specific seed file as defined in [Load Devices From the Seed File](#), page 2-24. However, once you import a sensor, you must identify the monitored networks that it monitors. For more information, see [Specify the Monitored Networks for Cisco IPS or IDS Device Imported from a Seed File](#), page 6-8.
- **IntruShield Sensors.** These sensors are defined by importing a custom seedfile; however, once you import the sensors, which appear as children of the IntruShield Manager host, you must identify the monitored networks for each sensor. For more information, see [Add IntruShield Sensors Using a Seed File](#), page 6-27.

Seed File Header Columns

[Table 2-4](#) describes the columns in the seed files and identifies valid values. If you do not enter a value for a given column, you must enter a comma to delineate that column.

**Note**

Remember that you do not have to add all of the devices' configuration information at once. You can start by adding the device's name and its access IP address. You can always return later, when the MARS starts to report to you, and provide more details.

Table 2-4 **Seed File Column Description**

Column	Type	Entry
Column A	NAME OR IP	The device's name or IP address. (Mandatory) If the device name is provided and Column U is empty, MARS performs a DNS lookup to identify the address which will be used to populate the Access and Reporting IP fields Note If an IP address appears in Column U, that address overrides any address or derived address specified in Column A. However, the name value specified in Column A is used.
Column B	SNMP RO/RW Community	The device's SNMP RO community name here. Note MARS does not support the following characters in the SNMP RO community string: ' (single quote), " (double quote), < (less than symbol), and > (greater than symbol).
Column C	EMPTY	Empty placeholder column.
Column D	EMPTY	Empty placeholder column.

Table 2-4 Seed File Column Description (continued)

Column	Type	Entry
Column E	DEVICE TYPE	The device type designator. Note Some of the devices supported in the GUI cannot be entered via a CSV file. Use the following strings represent the desired device type: • ASA: for Cisco ASA devices • CiscoIDS4x: for appliance running Cisco IPS 4.x (not modules) • CiscoIPS5x: for appliance running Cisco IPS 5.x (not modules) • FWSM: for Cisco FWSM 1.1 • PIX: for Cisco PIX 6.0, 6.1, 6.2, and 6.3 devices • PIX7x: for Cisco PIX 7.0 devices • IOS: for Cisco IOS 12.2 (default) • SWITCH-CATOS: for Cisco Switch in Hybrid Mode • SWITCH-IOS: for Cisco Switch in Native Mode • EXTREME: for Extreme ExtremeWare 6.x • NETSCREEN: for ScreenOS 4.0 and 5.0 • WINDOWS: for Window host • Windows2000: for Windows 2000 host • Windows2003: for Windows 2003 host • WindowsNT: for Windows NT 4.x host. • SOLARIS: for Solaris host • LINUX: for Linux host Note In the case of host files, Linux, Solaris, and Windows, MARS is configured by default to receive events from the hosts specified in a seed file. However, for a Windows host where the RPC settings are also specified in the seed file, MARS will both pull and receive logs from the host by default.
Column F	ACCESS TYPE	The Access Type for this device. Your choices are: • TELNET • FTP • SSH • SNMP (default) • RPC (Windows only) In the RPC case, the username field (Column G) should be non-empty. The password can be provided in Column H. If RPC access type and username are given, the PULL flag is set by the backend in addition to the default RECEIVE flag.

Table 2-4 **Seed File Column Description (continued)**

Column	Type	Entry
Column G	USER NAME	The TELNET, SSH, FTP, or RPC user name. This column is only valid if you have used TELNET, SSH, or FTP in Column F .
Column H	SSH/FTP/RPC PASSWORD	The SSH or FTP Password for the device. This column is only valid if you have used SSH or FTP in Column F .
Column I	TELNET PASSWORD	The Telnet password for the device.
Column J	ENABLE PASSWORD	The enable password (applicable only with FWSM, PIX, or IOS devices).
Columns K	EMPTY	Empty placeholder column.
Column L	EMPTY	Empty placeholder column.
Column M	EMPTY	Empty placeholder column.
Column N	EMPTY	Empty placeholder column.
Column O	EMPTY	Empty placeholder column.
Column P	EMPTY	Empty placeholder column.
Column Q	EMPTY	Empty placeholder column.
Column R	EMPTY	Empty placeholder column.
Column S	EMPTY	Empty placeholder column.
Column T	FTP LOCATION [if Access Type =FTP]	The location for the FTP file. This location starts from the FTP root, not the sysroot. If, for example, the file is at <ftproot>/configdata/router1.txt, using ./configdata/router1.txt is correct.
Column U	Access/Reporting IP [optional]	The Access IP and Reporting IP address to use when populating this device. The MARS Appliance uses this address to communicate with the device. See Understanding Access IP, Reporting IP, and Interface Settings, page 2-8

Load Devices From the Seed File

Once you have completed the seed file, you must the CSV file on to the FTP server where you want to upload it.

To load the file into the MARS, follow these steps:

-
- Step 1** Click **Admin > System Setup > Security and Monitor Devices > Load From Seed File**.
- Step 2** Enter the FTP Server's IP address, the user name and password for the FTP server, the path, and the file name for the seed file.
- The FTP path starts from the FTP root, not from the sysroot for the configuration path.
- Step 3** Click **Submit**.
- Once you have loaded devices from the seed file, return to each device. Continue to configure the devices and to add information such as reporting IP addresses, and SNMP information.
- Step 4** Once add a device, you must click **Activate** for MARS to correctly process events received from that device. For more information, see [Activate the Reporting and Mitigation Devices, page 2-27](#).

**Note**

Using a seed file to define the reporting devices replaces the manual definition of the device; however, the topology information will not be available. After adding the reporting devices via a seed file, you must either manually discover each device by selecting the device, clicking Edit, and then clicking the Discover button or by scheduling a topology discovery. In addition, some device types required that you define additional settings (see [Devices that Require Updates After the Seed File Import, page 2-21](#)).

Adding Reporting and Mitigation Devices Using Automatic Topology Discovery

On the Admin page, under the Topology Discovery Information section, three links exist, allowing you to define the settings required to discover reporting and mitigation devices automatically. These links are:

- **Community String and Networks.** Allows you to define SNMP RO community strings on a per network or IP range basis. Networks and SNMP RO strings can overlap. At least one SNMP string must be defined before discovery is attempted.
- **Valid Networks.** Identifies the set of networks and IP ranges that you want to discover. You should also define one or more SNMP targets. If no SNMP targets are defined, MARS uses its own gateway as the SNMP target. SNMP targets should be layer 3 gateway devices, such as a router or firewall with SNMP RO community strings defined and discovery permitted; they should also be defined on a per network or per range basis if you wish to separate the discovery using schedule rules. At least one valid network must be defined before discovery is attempted.
- **Topology/Monitored Device Update Scheduler.** While not required for discovery, it does allow you to increase the frequency of topology discovery and further refine the potential depth of a discovery based on a particular schedule rule. The default schedule rule is once a month for all valid networks. However, if no valid networks are defined, the process wakes up, sees no valid networks are defined, and quits. Each schedule rule allows you to select which networks, as defined within the list of valid networks and ranges, that should be discovered according to frequency also specified in the rule. As connected networks often exist, you can refine which networks are discovered by ensuring that separate schedule rule exists for each network that you do not want to be automatically discovered as part of a connected network.

Based on the networks defined within the schedule rules, MARS starts with the first SNMP target associated with those networks or ranges as defined under Valid Networks and attempts to discover that device using SNMP discovery. The discovery process continues as long as the target device provides additional routes and the addresses of such routes are part of the networks in another schedule rule. The process also iterates through each SNMP target that is defined. The entire discovery process is limited based on the schedule rule's bounding networks, the SNMP targets, the valid network and IP ranges, and the SNMP RO community strings, which are defined on a per network basis. Networks and SNMP RO community strings can overlap, in which case MARS tries each string against the gateway addresses discovered within that network. The discovery process only discovers Layer 3 gateway devices, such as routers and firewalls. It does not discover hosts, unless those hosts are defined as the explicit target within a schedule rule (see [Scheduling Topology Updates, page 2-39](#)).

As the discovery process identifies supported reporting and mitigation devices, it adds those devices to the Monitoring and Security Devices list (Admin > Monitoring and Reporting Devices), identifying them by the Reporting IP. You can later edit these discovered devices to provide Access IP information and perform more thorough device-level discovery. Once a device is listed under Monitoring and Reporting Devices, it may be rediscovered, but it will not be added again unless it has been properly deleted (see [Delete a Device, page 2-19](#)).

For more information on these settings, see:

- [Configuring Layer 3 Topology Discovery, page 2-36](#)
- [Scheduling Topology Updates, page 2-39](#)


Note

Once the discovery process is complete, you must click **Activate** for MARS to correctly process events received from that device. For more information, see [Activate the Reporting and Mitigation Devices, page 2-27](#).

Verify Connectivity with the Reporting and Mitigation Devices

After loading the seed file or manually adding devices, you can verify that the devices were loaded by clicking **Admin > System Setup > Security and Monitor Devices**. You should see the devices that you have added populating this page.

You can test the devices by checking the box next to the name of the device and clicking **Edit**. On the device's page, click **Discover** or **Test Connectivity**. The UI displays a “holding pattern” screen while it connects to the device. When complete, it shows you the device's discovery screen.


Note

Some devices cannot be checked for connectivity nor can be discovered. The next section, [Discover and Testing Connectivity Options, page 2-26](#), contains a list of devices that can be checked or discovered.

Discover and Testing Connectivity Options

When you add a device, you should check its connectivity or perform the discovery. Checking a device's connectivity or discovery analyzes the device's configuration, checks that MARS can process its events, and that MARS can understand its NAT information.

You can test these devices for connectivity or perform discovery:

- Cisco IOS
- Cisco PIX
- Cisco ASA
- Cisco Switch CatOS
- Cisco Switch IOS
- Cisco IDS
- Cisco IDSM
- Cisco FWSM
- Cisco Security Manager server
- Cisco VPN Concentrator 4.x
- Check Point
- Extreme ExtremeWare 6.x
- NetScreen

Run a Reporting Device Query

Another method to see the added devices, is to run a query with the display format: **Reporting Device Ranking**.



Note

You might not see all of the devices that you loaded using the seed file right away because of lag, network size and traffic. If you do not see a device after waiting, it could be due to input error.

To run a reporting device ranking query, follow these steps:

-
- Step 1** Click the **Queries / Reports** tab.
 - Step 2** On the Queries page, in the Query Event Data table, click **Event Type** in the Display Format column.
 - Step 3** Select **Reporting Device Ranking**.
 - Step 4** Click **Apply**.
 - Step 5** Click **Submit** to run the query.
-

Activate the Reporting and Mitigation Devices

After you have added reporting devices and mitigation devices to MARS, you must activate those devices before MARS begins to fully process the data provided by those devices. This processing is different from those devices discovered on the network, where the logs sent to the appliance are stored, but your ability to interact with that data is limited to queries and reports. Typically, MARS runs inspection rules and generates notifications only against the data retrieved from activated devices.

Once a device is known to the MARS Appliance, all data provided by that particular device can be normalized and sessionized, which enables that device's data to be used to fire an incident



Note

Default installations of MARS do not fire incidents based on data received from unknown devices. However, you can still enable this by creating one or more rules that use keyword search. A device must be defined for the MARS to be able to parse and sessionize the event data. The act of parsing the event data correctly is what ensures rules fire more accurately.



Tip

You must click **Activate** whenever you add or modify rules, drop rules, reports, or add or modify any options or settings under in the Admin tab other than those on the User Management subtab. Otherwise, the changes that you make will not take effect.

To activate added devices, follow these steps:

-
- Step 1** For each device that you want to add, provide the device details and click **Submit** to add the device.
The Submit action stores the device details in the database. Once you click Submit, your work is saved, even if you drop the administrative connection before clicking **Activate**.
 - Step 2** Once you have all of the devices desired for this administrative session, click **Activate**.

The Activate action differs from Submit in that MARS begins to inspect and generate notifications about the data provided by the devices.

**Tip**

If you are adding or editing several devices, it is better for the system to click **Activate** for several changes rather than for each individual change.

Data Enabling Features

Adding a the reporting devices and mitigation devices is the primary method of providing MARS with the data required to study the activities on your network. However, other features, both within the web interface and as part of configuring the devices, can provide MARS with additional data, which is used to refine the views it provides and to assist in the improving the overall effectiveness of the system. We think of these features as data enabling features.

This section contains the following topics:

- [Layer 2 Discovery and Mitigation, page 2-29](#)

Enable SNMP community strings to support the discovery the network topology. Allows for mapping to the port level for switches. Combined with 802.1x support required by NAC, this setting can resolve MAC address level settings for attached and wireless nodes on the network.

- [Networks for Dynamic Vulnerability Scanning, page 2-29](#)

Enables a Nessus-based scan of the targeted hosts. Nessus also uses nmap for OS fingerprinting and port scanning during a vulnerability assessment scan. These scans are conducted in response to suspicious activity to determine whether the attempted attack is successful or likely to succeed based on information such as target operating system type, patch level, and open ports on the host.

- [Understanding NetFlow Anomaly Detection, page 2-30](#)

By enabling NetFlow, MARS can detect anomalies in traffic and network usage by comparing new events with summary data. When anomalies are detected, MARS begins to store full NetFlow data. By default, full NetFlow data is not stored by MARS unless an incident is identified.

- [Host and Device Identification and Detail Strategies, page 2-36](#)

Details about reporting devices and the hosts that are on your network aids in the elimination of false positives, as well as improves the performance of MARS in assessing events.

- [Configuring Layer 3 Topology Discovery, page 2-36](#)

Layer 3 topology discovery aids in attack path analysis, as well as the population of the topology graph in the web interface.

- [Scheduling Topology Updates, page 2-39](#)

Topology update schedules are a critical part of many of the data enabling features, including discovery of Layer 2 and Layer 3 devices, as well as pulling information from specific reporting devices.

- [Configuring Resource Usage Data, page 2-41](#)

MARS can collect additional data from a select set of reporting devices, which is used to provide reports about CPU utilization, memory utilization, and device saturation. This data can be helpful in detecting anomalies as well in network capacity planning.

- [Configuring Network Admission Control Features, page 2-42](#)
Describes how to accomplish full NAC awareness, what it provides, and what products are required.
- [Configuring Distributed Threat Mitigation](http://www.cisco.com/en/US/products/ps6241/products_configuration_example09186a008067a2b0.shtml)
http://www.cisco.com/en/US/products/ps6241/products_configuration_example09186a008067a2b0.shtml
Describes how to accomplish full DTM awareness, the features it provides, and what products are required.
- [MARS MIB Format, page 2-43](#)
Describes the format of the MARS MIB, which helps integrate with other SNMP-based management applications on your network.

Layer 2 Discovery and Mitigation

Make sure that all the L2 devices have the SNMP RO community strings specified in the web interface for L2 mitigation, even if the access type is not SNMP. (See [False Positive Confirmation, page 19-6](#) for more information on mitigating an attack.)

The SNMP RO community string is *always* required on Layer 2 devices for L2 mitigation. L2 devices must be added manually—there is no automatic discovery for these device.



Note

MARS does not support the following characters in the SNMP RO community string: ' (single quote), " (double quote), < (less than symbol), and > (greater than symbol).

MARS does not discover L2 devices automatically as it does with L3 devices.



Note

L2 devices must be added manually; there is no automatic discovery for these devices. Make sure all the L2 devices (switches) have the SNMP RO community strings specified in the web interface, even if the access type is not SNMP. The SNMP RO community string is always required on L2 devices for L2 mitigation.

You can specify which L3 devices to discover by specifying networks and SNMP RO community values, as defined in [Configuring Layer 3 Topology Discovery, page 2-36](#).

The reason is MARS does not scan the network for devices. Therefore, you must manually add L2 devices using the web interface or a CSV file. Assuming that device discovery permission has been provided, L3 devices are discovered automatically using the route information provided by monitored gateways. Once devices are loaded/added in the web interface, user can use the topology scheduler feature to update the configuration of both L2 and L3.

For L2 devices SNMP access type is sufficient with RO community. But for mitigation, MARS requires SNMP RW community access. If SNMP RW community is not possible, select TELNET/SSH access type with SNMP RO Community.

Networks for Dynamic Vulnerability Scanning

With dynamic vulnerability scanning, the MARS probes the networks that you have specified for weaknesses. These automatic scans commence after a rule has fired that indicates an attack is in progress. Once an attack is underway, these scans accomplish the following:

- return information that determines if the attack failed
- return information that determines if the attack likely succeeded
- return false positive information
- assign severity to firing events and incidents

Select a Network for Scanning

To select a network for scanning, follow these steps:

-
- Step 1** Click the **Select** radio button.
 - Step 2** Click a network to scan.
 - Step 3** Click **Add**.
 - Step 4** Repeat [Step 1](#) through [Step 3](#).
 - Step 5** Click **Submit** when ready.
-

Create a Network IP Address for Scanning

To create a network address that you can use to define the scan settings, follow these steps:

-
- Step 1** Click the **Network IP** radio button.
 - Step 2** Enter the Network IP address and Mask.
 - Step 3** Click **Add**.
-

Create a Network IP Range for Scanning

To create a range of network addresses that you can use to define the scan settings, follow these steps:

-
- Step 1** Click the **IP Range** radio button.
 - Step 2** Enter the range of IP addresses.
 - Step 3** Click **Add**.
-

Understanding NetFlow Anomaly Detection

NetFlow is a Cisco technology that supports monitoring network traffic and is supported on all basic IOS images. NetFlow uses an UDP-based protocol to periodically report on flows seen by the Cisco IOS device. A *flow* is a Layer 7 concept that consists of a session set up, data transfer, and session teardown. For every flow, a NetFlow-enabled device record several flow parameters including

- Flow identifiers, specifically source and destination addresses, ports, and protocol

- Ingress and egress interfaces
- Packets exchanged
- Number of bytes transferred

Periodically, a collection of flows and its associated parameters are packaged in an UDP packet according to the NetFlow protocol and sent to any identified collection points. Because data about multiple flows is recorded in a single UDP packet, NetFlow is an efficient method of monitoring high volumes of traffic compared to traditional methods, including SYSLOG and SNMP.

The data provided by NetFlow packets is similar to that provided by SYSLOG, SNMP, or Checkpoint LEA as reported by enterprise-level firewalls, such as Cisco PIX, NetScreen ScreenOS, and Checkpoint Firewall-1. The difference being that NetFlow much more efficient. To receive comparable syslog data from a firewall device, the syslog logging level on the firewall must be set to DEBUG, which degrades firewall throughput at moderate to high traffic loads.

If NetFlow-enabled reporting devices are positioned correctly within your network, you can use NetFlow to improve the performance of the MARS Appliance and your network devices, without sacrificing MARS's ability to detect attacks and anomalies. In fact, NetFlow data and firewall traffic logs are treated uniformly as they both represent traffic in the network.

This section contains the following topics:

- [How MARS Uses NetFlow Data, page 2-31](#)
- [Guidelines for Configuring NetFlow on Your Network, page 2-32](#)
- [Enable Cisco IOS Routers and Switches to Send NetFlow to MARS, page 2-32](#)
- [Configuring Cisco CatIOS Switch, page 2-34](#)
- [Enable NetFlow Processing in MARS, page 2-34](#)

How MARS Uses NetFlow Data

When MARS is configured to work with NetFlow, you can take advantage of NetFlow's anomaly detection using statistical profiling, which can pinpoint day zero attacks like worm outbreaks. MARS uses NetFlow data to accomplish the following:

- Profile the network usage to determine a usage baseline
- Detect statistically significant anomalous behavior in comparison to the baseline
- Correlate anomalous behavior to attacks and other events reported by network IDS/IPS systems

After being inserted into a network, MARS studies the network usage for a full week, including the weekend, to determine the usage baseline. Once the baseline is determined, MARS switches to detection mode where it looks for statistically significant behavior, such as the current value exceeds the mean by 2 to 3 times the standard deviation.

By default, MARS does not store the NetFlow records in its database because of the high data volume. However, when anomalous behavior is detected, MARS does store the full NetFlow records for the anomalous entity (host or port). These records ensure that the full context of the security incident, such as the infected source and destination port, is available to the administrator. This approach to data collection provides the intelligence required by an administrator without affecting the performance of the MARS Appliance. Storing all NetFlow records consumes unnecessary CPU and disk resources.



Note

MARS only supports NetFlow version 5 and version 7.

Guidelines for Configuring NetFlow on Your Network

Ideally NetFlow should be collected from the core and distribution switches in your network. These switches, together with the NetFlow from Internet-facing routers or SYSLOG from firewalls, typically represent the entire network. With this in mind, review the following guidelines before deploying NetFlow in your network:

- MARS normalizes NetFlow and SYSLOG events to prevent duplicate event reporting from the same reporting device.
- Review VLANs in switches and pick several VLANs for which the traffic volume is low. This approach allows you to slowly integrate NetFlow and become comfortable with using it in your environment.
- Be aware of existing CPU utilization on NetFlow capable devices. For more information on understanding how NetFlow affects the performance of routers and network throughput, see the following link:

http://www.cisco.com/en/US/tech/tk812/technologies_white_paper0900aecd802a0eb9.shtml

- Consider using a sampling of NetFlow data 10:1 100:1 ratio's in highly utilized VLANs.
- Be selective in using NetFlow, you do not need to enable it on all NetFlow-capable devices. In fact, such usage can create duplicate reporting of events, further burdening the MARS Appliance.
- MARS uses NetFlow versions 5 and 7. Ensure that the version of Cisco IOS software or Cisco CatOS running on your reporting devices supports at least one of these NetFlow versions.

The taskflow for configuring NetFlow to work with MARS is as follows:

1. Identify the reporting devices on which to enable NetFlow.
2. Enable NetFlow on each identified reporting device and direct the NetFlow data to the MARS Appliance responsible for that network segment.
3. Verify that all reporting devices are defined in the MARS web interface.
4. Enable NetFlow processing in the MARS web interface.
5. Allow MARS to study traffic for a week to develop a usage baseline before it begins to generate incidents based on detected anomalies.

The following tasks provide guidance on the required device configuration:

- [Enable Cisco IOS Routers and Switches to Send NetFlow to MARS, page 2-32](#)
- [Enable NetFlow Processing in MARS, page 2-34](#)

Enable Cisco IOS Routers and Switches to Send NetFlow to MARS

For more information on NetFlow and configuring the settings in Cisco IOS, refer to:

http://www.cisco.com/univercd/cc/td/doc/product/software/ios124/124cg/hnf_c/nfb_ov.htm

Before you configure NetFlow from MARS, you must first configure it on the router or switch.

To enable NetFlow on a Cisco IOS router or switch and to push those events to the MARS Appliance, follow these steps:

-
Step 1	Log in to the Cisco IOS router or switch with administrator's privileges.
Step 2	Enter the following commands:

Command	Purpose
enable	Turn on enable mode.
configure terminal	Enter global configuration mode.
	Note Commands in this mode are written to the running configuration file as soon as you enter them (using the Enter key/Carriage Return).
ip flow-export destination <i><MARS_IP_address></i> <i><UDP_port></i>	Enables the data export to the MARS Appliance on UDP port 2055 (assuming the default port is used). <i>MARS_IP_address</i> is the IP address of the MARS Appliance that is responsible for processing the NetFlow events for this reporting device. <i>UDP_port</i> is the default UDP port to send NetFlow (the default port is 2055).
ip flow-export source <i><syslog_interface_name></i>	Set the source IP for the interface to send the NetFlow. The <i>syslog_interface_name</i> value should be the interface attached to the network through which the MARS Appliance is reachable, and it must equal the syslog source interface name.
ip flow-export version <i><version_number></i>	Identifies which version of NetFlow, 5 or 7, to use when generating events. Cisco recommends using version 5 if supported. <i>version_number</i> is either 5 or 7. MARS only supports NetFlow versions 5 and 7.
ip flow-cache timeout active 5	Configures the flow timeout. This timeout value breaks up long-lived flows into 5-minute segments. You can choose any number of minutes between 1 and 60; however, selecting the default of 30 minutes will result in spikes appearing in utilization reports.
ip flow-cache timeout inactive 15	Ensures that those flows that have finished are exported in a timely manner.

Step 3 For each interface in the device, enter the following commands:

Command	Purpose
interface <i><interface_name></i>	Specifies the interface for which you want to enable NetFlow and it enters the interface configuration mode. <i>interface_name</i> is the name of the interface to which the MARS is connected. This command varies based on the device type. For example, interface <i>type slot/port-adapter/port</i> (Cisco 7500 series routers) interface <i>type slot/port</i> (Cisco 7200 series routers)
ip route-cache flow	Enables NetFlow for the selected interface.

Step 4 To verify that NetFlow is enabled correctly, enter the following commands:

show ip flow export

show ip cache flow

Step 5 To exit enable mode, enter the following command:

exit

Configuring Cisco CatIOS Switch

Some Cisco Catalyst switches support a different implementation of NetFlow that is performed on the supervisor. With the cache-based forwarding model, which is implemented in the Catalyst 55xx running the Route Switch Module (RSM) and NetFlow Feature Card (NFFC), the RSM processes the first flow and the remaining packets in the flow are forwarded by the Supervisor. This support is also implemented in the early versions of the 65xx with MSFC. The deterministic forwarding model used in the 65xx with MSFC2 do not use NetFlow to determine the forwarding path, the flow cache is only used for statistics as in the current IOS implementations. In all of these configurations, flow exports arrive from both the RSM/MSFC and the Supervisor engines as distinct streams.

The router-side running IOS is configured as specified in [Enable Cisco IOS Routers and Switches to Send NetFlow to MARS](#), page 2-32. However, to configure the he CatIOS NetFlow Data Export, use the following commands:

```
set mls flow full
```

```
set mls nde version 5
```

```
set mls nde <MARS_IP_address> 2055
```

```
set mls nde enable
```

From a user's perspective, the switch is only running IOS when the 65xx is running in Native mode.

Enable NetFlow Processing in MARS

Once you have enabled NetFlow on your routers or switches and you have directed those devices to publish NetFlow data to the MARS Appliance, you must configure the appliance to process that data. This configuration involves determining how to store data, as well as identifying which networks you want to process for anomalous behavior. Both of these options can affect the rate at which MARS can process events: storing the full event data rather than summary data burdens the system with writing large volumes of data rather than processing new incoming events. Also, by not specifying a select set of networks, MARS studies all networks.

Step 1 Click **Admin > System Setup > NetFlow Config Info**.

NetFlow Configuration

Global NetFlow UDP Port:	2055
Enable NetFlow Processing:	Yes ☒ No ☐
Always Store NetFlow Records:	Yes ☐ No ☒

NetFlow Valid Network Addresses

	Add	☐ Network IP:	
	Remove	Mask:	
	☐ IP Range:		

143229

- Step 2** Under **NetFlow Configuration**, enter the NetFlow **Global NetFlow UDP Port**. This is the default port for MARS to listen for NetFlow; the default value is 2055.

**Note**

This value must match the value you entered in the “ip flow-export destination” command when configuring the router (see [Enable Cisco IOS Routers and Switches to Send NetFlow to MARS, page 2-32](#)). Also, verify you have enabled this traffic to flow between the router or switch and the MARS Appliance on any intermediate gateways, such as routers and firewalls.

- Step 3** Choose whether to **Enable NetFlow Processing**.

- **Yes** tells MARS to process the NetFlow logs.
- **No** disables the processing of NetFlow data into the MARS.

- Step 4** Choose whether to **Always Store NetFlow Records**.

- **Yes** tells MARS to store all of the NetFlow events in the database. Selecting this option can slow down the system by greatly decreasing the number of events per second that MARS is able to process.
- **No** tells MARS to store only anomalies. The MARS detects anomalies by using two dynamically generated watermarks comparing the previous data against current data. When the data breaches the first watermark, MARS starts to save that data. When the data rises above the second watermark, MARS creates an incident.

- Step 5** Under **NetFlow Valid Network Addresses**, you can enter one or more for networks you want to monitor and use the << **Add** button to add them.

- Specifying one or more networks causes MARS to generate NetFlow-based anomalies that occur only on the specified networks.

**Note**

To reduce the memory usage and increase performance of the appliance, you can configure MARS to profile hosts belonging to a set of valid networks.

- Leaving this value blank (not specifying any networks) causes MARS to examine all networks for anomalous behavior based on the NetFlow events.

Step 6 Click **Submit** to save your changes.

Step 7 To enable NetFlow processing by the MARS Appliance, click **Activate**.

Before MARS can start detecting anomalies based on NetFlow data, it must first develop a baseline for network behavior. It takes a full week, including the weekend, for MARS to develop such a baseline. After this period has elapsed, MARS can start generating incidents based on NetFlow's anomaly detection.

Host and Device Identification and Detail Strategies

MARS studies many events at the network layer, relying on firewalls, routers, and IPS devices to identify anomalies and suspected incidents at a layer above the endpoint hosts that are the source or destination of network sessions. If operating exclusively at this network layer, MARS can generate a number of false positive incidents that must be manually investigated. However, several features exist that allow you to provide host-level details to MARS:

- Enable event reporting from the hosts on your network. MARS can receive, and in some cases, pull event data directly from the hosts on your network. This additional data allows MARS to verify the success of some attacks, as well as to report issues with the operation of the host, such as including them in “device down” reports if they are inaccessible. For more information on configuring the hosts and MARS to pull or receive data from those hosts, see the following topics:
 - [Adding Generic Devices, page 10-1](#)
 - [Sun Solaris and Linux Hosts, page 10-2](#)
 - [Microsoft Windows Hosts, page 10-4](#)
- Manually identify the operating system type and network services running on discovered hosts. For more information, see [Define Vulnerability Assessment Information, page 10-11](#) and [Identify Network Services Running on the Host, page 10-13](#)
- Manually identify common hosts and nodes in your network by adding other devices via Management > IP Management. This additional data allows you to identify those hosts that are likely to be involved in network sessions without having to configure the hosts to provide event data directly to MARS. This open allows you to provide vulnerability assessment information to assist in the reduction of false positives. For more information on adding hosts manually, see [Add a Host, page 23-4](#).

Configuring Layer 3 Topology Discovery

For the MARS to reach full operability, you must specify its community strings and select the networks that you want to discover. Once the appliance discovers these networks, you get a more accurate view of MAC addresses, end-point lookup (attack paths), and network topology. Topology discovery enables operation level three, see [Levels of Operation, page 2-1](#) for more information.

See [Figure 17-13 on page 17-7](#) for a view of the topologies.

**Note**

Remember to activate additions and changes to your community strings and valid networks by clicking **Activate**.

Add a Community String for a Network

To add a community string for a network IP, follow these steps:

- Step 1** To open the Community Strings and Networks page, click **Admin > Community Strings and Networks**.

Community Strings and Networks

- Step 2** Click the **Network IP** radio button.
- Step 3** Enter the Community String, Network IP address, and Mask.
- Step 4** Click **Add**.
- Step 5** Repeat [Step 2](#) through [Step 4](#) for all the community strings that you want to add.
- Step 6** Click **Submit** to commit these additions.

Add a Community String for an IP Range

To add a community string for an IP range, follow these steps:

- Step 1** To open the Community Strings and Networks page, click **Admin > Community Strings and Networks**.
- Step 2** Click the **IP Range** radio button.
- Step 3** Enter the Community String and its IP Range.
- Step 4** Click **Add**.
- Step 5** Repeat [Step 2](#) through [Step 4](#) for all the community strings that you want to add.
- Step 6** Click **Submit** to commit these additions.

You can add multiple community strings for the same network by adding similar entries.

Add Valid Networks to Discovery List

Adding valid networks confines the MARS to discover the networks that you want. MARS uses this information to create topologies, find MAC addresses, and for end-point lookup (attack paths).



Note

You can only specify networks for the zone where the MARS Appliance operates.

To add valid networks, follow these steps:

-
- Step 1** Click **Admin > Valid Networks** to open the Valid Networks page.
 - Step 2** Enter the **SNMP Target's** IP address.

The SNMP target is the entry-point where the MARS starts discovering devices on a network. It typically identifies an address on a default gateway of the network.
 - Step 3** Click either **Network IP** or **Network Range** to define the scope of the scan.
 - Step 4** Enter the appropriate information.
 - Step 5** Click **Submit**.
-

Remove Networks from Discovery List

To remove a network, follow these steps:

-
- Step 1** Click **Admin > Valid Networks** to open the Valid Networks page.
 - Step 2** Click the network that you want to remove.
 - Step 3** Click **Remove**.
-

Discover Layer 3 Data On Demand

You can schedule topology discovery, as defined in [Scheduling Topology Updates, page 2-39](#). However, you can also initiate an on-demand discovery.

To perform an on-demand discovery, follow these steps:

-
- Step 1** Click **Admin > Valid Networks** to open the Valid Networks page.
 - Step 2** Verify that the list of Valid Network Addresses contains the networks that you want to discover.
 - Step 3** Click **Discover Now**.
-

Scheduling Topology Updates

You can configure MARS to run automatic topology updates on devices, networks, and groups of networks. Scheduling topology updates is a critical part of keeping the MARS Appliance abreast of changes in the network and of changes to the configuration settings of the reporting devices and mitigation devices. This operation is similar to clicking Discover when defining a reporting device.

Configuration discovery depends on the device type, proper authorization, an access type, such as Telnet or SSH, and an access IP address. When device discovery is performed, MARS contacts the device and conducts a topology and configuration discovery. This discovery collects all of the route, NAT, and ACL-related information for the device or admin context. In addition, the name of the device may change to `hostname.domain` format if it was not already entered as such. If discovering a device that supports them, MARS also discovers information about modules, admin contexts, and security contexts. Another effect of scheduled updates is that MARS keeps the network diagram and attack paths current in the Dashboard.

This feature also allows you to pull data from those devices that require interval-based polling. The list to devices that require such polling are:

- Qualys QualysGuard
- eEye REM
- FoundStone FoundScan
- Check Point log servers

Figure 2-1 **Example Scheduled Update for eEye REM**

Name:

10.1.1.86/255.255.255.255

Add

Select:

Network IP:

Mask:

IP Range:

Remove

Schedule	Time of Day	Days
☒ Run On Demand Only		
☐ Daily	12:00 Midnight	
☐ Weekly	12:00 Midnight	☐ Sun ☐ Mon ☐ Tue ☐ Wed ☐ Thu ☐ Fri ☐ Sat
☐ Monthly	12:00 Midnight	☐ 1st ☐ 2nd ☐ 3rd ☐ 4th ☐ 5th ☐ 6th ☐ 7th ☐ 8th ☐ 9th ☐ 10th ☐ 11th ☐ 12th ☐ 13th ☐ 14th ☐ 15th ☐ 16th ☐ 17th ☐ 18th ☐ 19th ☐ 20th ☐ 21st

Schedule a Network Discovery

To add a network for scheduled discovery, follow these steps:

- Step 1** Click **Admin > Topology/Monitored Device Update Scheduler**.

The Topology/Monitored Device Update Scheduler page displays.

- Step 2** Click **Add**.
- Step 3** Enter a name for the network (or group of networks).
- Step 4** Select or enter your networks:
- Click the **Select** radio button, and select a network from the list.
 - Click the **Network IP** radio button, and enter the IP address and Mask.
 - Click the **IP Range** radio button, and enter the IP ranges.
- Step 5** Click **Add** to move the network into the selected field.
- To remove an item in the selected field, click it to highlight it, and click **Remove**.
- Step 6** In the schedule table, select the appropriate radio button and its time criteria:
- **Run On Demand Only**
 - **Daily** and the Time of Day
 - **Weekly**, the Time of Day, and the Days
 - **Monthly**, the Time of the Day, and the Dates
- Step 7** Click **Submit**.
-

To edit a scheduled topology discovery

- Step 1** Check the box next to the Topology Group.
- Step 2** Click **Edit**.
- Step 3** Click **Add** to move the network into the selected field.
- To remove an item in the selected field, click it to highlight it, and click **Remove**.
- Step 4** In the schedule table, select the appropriate radio button and its time criteria:
- **Run On Demand Only**
 - **Daily** and the Time of Day
 - **Weekly**, the Time of Day, and the Days
 - **Monthly**, the Time of the Day, and the Dates
- Step 5** Click **Submit**.
-

To delete a scheduled topology discovery

- Step 1** Check the box next to the Topology Group.
- Step 2** Click **Delete**.
-

To run a topology discovery on demand


Note

You can run any scheduled or on-demand topology discoveries at any time.

Step 1 Check the box next to the Topology Group.

Step 2 Click **Run Now**.

Configuring Resource Usage Data

While the Monitor Resource Usage box appears on every host and reporting device, only three device types actually provide resource usage data to MARS:

- Cisco IOS routers and switches
- Cisco PIX
- Check Point devices

For these three devices, MARS can provide data about CPU utilization, memory utilization, and device saturation.

To enable the collection of resource usage data, you must ensure that the resource usage-specific events are logged by the reporting devices, that the SNMP RO community string is set, that those devices forward the events to MARS, and that the device is defined in the web interface as a reporting device or mitigation device. In addition, you must select **Yes** in the Monitor Resource Usage box of the General tab for each supported reporting device.

Once configured, MARS uses SNMP to poll the device every 5 minutes for the following SNMP OIDs:

- Bytes in/out of every interface on the device (Cisco IOS, Cisco PIX)
- Number of current connections (Cisco PIX, Check Point)
- CPU of last second and last 60 seconds (Cisco IOS, Cisco PIX)
- Memory free/used (Cisco IOS, Cisco PIX)

It also detects anomalous resource utilization if the consumption is significantly higher than the hourly average.

The following resource usage data reports are available:

- Resource Utilization: Bandwidth: Inbound - Top Interfaces
- Resource Utilization: Bandwidth: Outbound - Top Interfaces
- Resource Utilization: CPU - Top Devices
- Resource Utilization: Concurrent Connections - Top Devices
- Resource Utilization: Errors: Inbound - Top Interfaces
- Resource Utilization: Errors: Outbound - Top Interfaces
- Resource Utilization: Memory - Top Devices

You can define custom rules, reports, and queries about resource usage based on the following events:

- CPU Utilization Higher Than 50%

- CPU Utilization Higher Than 75%
- CPU Utilization Higher Than 90%
- CPU Utilization Abnormally High
- Memory Utilization Higher Than 50%
- Memory Utilization Higher Than 75%
- Memory Utilization Higher Than 90%
- Memory Utilization Abnormally High

There is also a pre-defined resource utilization inspection rule:

- System Rule: DoS: Network Device - Success Likely
- System Rule: DoS: Network - Success Likely
- System Rule: Resource Issue: Network Device

Configuring Network Admission Control Features

Network Admission Control (NAC) is a Cisco Systems sponsored industry initiative that uses the network infrastructure to enforce security policy compliance on all devices seeking to access network computing resources, thereby limiting damage from viruses and worms.

Using NAC, organizations can provide network access to endpoint devices such as PCs, PDAs, and servers that are verified to be fully compliant with established security policy. NAC can also identify noncompliant devices and deny them access, place them in a quarantined area, or give them restricted access to computing resources.

MARS supports the NAC initiative by storing and reporting about the NAC-based events generated by the various reporting devices on your network. The devices include:

- Cisco Trust Agent. While CTA does not report to MARS, it does report discovered settings to the Cisco network devices, from which MARS collects events.
- 3rd-party 802.1x Supplicants.
- Cisco IOS routers running Cisco IOS Software, Release 12.3(8)T with security.
- Cisco VPN 3000 Concentrators
- Cisco Secure ACS
- Cisco Security Agent

To enable NAC reporting on your network, you must ensure that the NAC-specific events are logged by the reporting devices, that those devices forward the events to MARS, and that the device is defined in the web interface as a reporting device or mitigation device.

The following reports are available to support NAC:

- Activity: Security Posture Not Up To Date - All Events
- Activity: Security Posture Not Up To Date - Top Users
- Activity: Security Posture Up To Date - Top Users
- Activity: Security Posture Validation Failure - Top Users
- Activity: Security Posture w/o Credentials - Top Hosts

For information on configuring reporting devices and mitigation devices with NAC support, see [Enable NAC-specific Messages, page 3-4](#).

Integrating MARS with 3rd-Party Applications

MARS provides multiple integration methods with 3rd-party applications. The following topics describe how to integrate using these methods:

- [Forwarding Alert Data to 3rd-Party Syslog and SNMP Servers, page 2-43](#)
- [MARS MIB Format, page 2-43](#)
- [Relaying Syslog Messages from 3rd-Party Syslog Servers, page 2-44](#)

Forwarding Alert Data to 3rd-Party Syslog and SNMP Servers

You can forward alert data from MARS to third-party syslog and SNMP servers. The data is forwarded on a per rule basis. In other words, you must configure those rules for which you want to forward alert data to include either SNMP, syslog, or both as a notification methods. When a rule fires, the notifications will be sent in the selected formats to the specified recipients, which should be the desired servers in the case of SNMP and syslog.

For more information on configuring notification methods for a rule, see [Setting Alerts, page 21-23](#). To learn more about the SNMP MIB format sent by MARS, see [MARS MIB Format, page 2-43](#).

MARS MIB Format

The MARS management information base (MIB) is defined for all MARS releases. The SNMP notification contains the same content as the syslog generated by MARS.

The MARS MIB definition is as follows:

```
enterprises.16686.1.0 string "MARS-1-101"
enterprises.16686.2.0 string "<alert_content>"
```

The MARS private enterprise number is 16686 and <alert_content> is defined as follows:

```
<<priorityInfo>> <current_time> %MARS-1-101: Rule <ruleid> (<rulename>) fired and caused
<color> Incident <incidentId>, starting from <starttime> to <endtime>.
```

In the following example of the SNMP notification output, 10.1.1.1 is the IP address of the MARS Appliance:

```
SNMPv2-SMI::enterprises.16686 10.1.1.1 SNMPv2-SMI::enterprises.16686.1.0 "MARS-1-101"
SNMPv2-SMI::enterprises.16686.2.0 "<34>Mon Apr 28 20:11:43 2003 %MARS-1-101: Rule 45513
(Nimda Attack) fired and caused red Incident 12265001, starting from Mon Apr 28 19:58:47
2003 to Mon Apr 28 20:11:21 2003."
```



Note

Notifications are sent only from the Local Controller.

Relaying Syslog Messages from 3rd-Party Syslog Servers

You can rapidly deploy MARS by forwarding messages from existing syslog-ng or Kiwi syslog servers. This feature eliminates the network and device changes required to insert MARS into an operational network. You are no longer required to configure each network device to publish its syslog messages directly to MARS, which saves time, avoids device change approval processes, preserves packet processing performance of the network devices, and ensures daily network operations proceed without interruption. This relay feature also allows the correlation and inspection of syslog messages from reporting devices, such as those on the DMZ, for which corporate policies might prohibit the existence of or connection to configuration information.

If your network devices already publish syslog messages to syslog-ng or Kiwi syslog servers, you can configure those servers to forward messages to the MARS Appliance and identify the syslog servers in MARS. Currently, MARS parses the syslog messages generated by the following devices: Cisco PIX, Cisco IOS, Cisco CatOS, Cisco ICS, Cisco ASA, Cisco FWSM, Cisco VPN 3000, Cisco Secure ACS, Snort IDS, Juniper/Netscreen firewalls, Solaris, Linux, and Microsoft Internet Information Server (ISS), Microsoft Windows running the SNARE agent. For other devices, you can define custom log parsers.

The MARS Appliance can begin processing and storing the events while you define the reporting devices using the MARS user interface. You are still required to define the reporting device by IP address and device type in MARS to ensure proper event correlation; however, you are not required to configure device to publish syslog messages directly to MARS.

To configure MARS to work with a syslog relay server, perform the following tasks:

1. Configure the syslog relay server to forward correctly formatted messages to MARS. See [Configure Syslog-ng Server to Forward Events to MARS, page 2-44](#) or [Configure Kiwi Syslog Server to Forward Events to MARS, page 2-45](#).
2. Identify the MARS Appliance as a forward target.
3. Add the syslog relay server to the MARS user interface. See [Add Syslog Relay Server to MARS, page 2-45](#).
4. Add the reporting devices monitored by the syslog relay server to the MARS user interface. See [Add Devices Monitored by Syslog Relay Server, page 2-46](#).

Configure Syslog-ng Server to Forward Events to MARS

We recommend the following settings in the configuration options of the syslog-ng.conf file to ensure good integration of syslog-ng with MARS:

```
options { long_hostnames(off); use_dns(0); keep_hostname(yes); };
```

where

- The long_hostnames(off) setting conforms to RFC 3164, which recommends that the HOSTNAME does not contain domain name.
- The use_dns(0) setting ensures that the IP address is used in HOSTNAME rather than the hostnames.
- The keep_hostname(yes) setting preserves the original sending device's HOSTNAME even when it is relayed more than once.

In addition to configuring the message format, you must specify that the MARS Appliance is a destination loghost on UDP port 514. The following lines must appear in the syslog-ng.conf file:

```
destination loghost { udp("IP address of MARS Appliance" port(514)); };  
log { source(src); destination(loghost); };
```

```
log { source(net); destination(loghost); };
```

Configure Kiwi Syslog Server to Forward Events to MARS

We recommend the following settings in the configuration options of the Kiwi Syslog Daemon to ensure good integration of Kiwi with MARS:

-
- Step 1** Expand the **File > Setup > Rules > Actions** tree.
- Step 2** Right on **Actions** and click **Add an Action**.
- Step 3** Enter a name for the action, such as “Forward to pncop”.
- Step 4** For the following fields, enter the following values:
- **Destination IP address or hostname** — Enter the IP address of the MARS Appliance.
 - **Protocol** — UDP
 - **New Facility** — No Change
 - **New Level** — No Change
 - **Port** — 514
 - **Send with RFC 3164 header information** — Selected if the syslog server receives syslog messages directly from the source devices only. Clear if the syslog server also receives syslog messages from relays. Do not configure mixed relays.
- This additional header is necessary for the supported device types that do *not* have HOSTNAME in the syslog messages; thereby allowing MARS to correctly identify the original sending device. However, this option cannot be used on a Kiwi relay of relay. To support a Kiwi relay of relay in MARS, the first relay must have this option selected and must receive syslog messages only from the source devices, and all other relays must have this option cleared and must only receive syslog messages from other Kiwi relays, not directly from devices.
- **Retain the original source address of the message** — Cleared.
- Step 5** If you are using SNARE agents, click **Setup > Modifiers** and clear “Replace non printable characters with <ASCII value>”
- If this value is selected, tabs appear as <009> in the Windows event logs, which prevents MARS from parsing the events correctly.
- Step 6** Save your changes.
-

Add Syslog Relay Server to MARS

In addition to representing each of the potential reporting devices, you must define the syslog relay server so that MARS knows for which messages it should attempt to discover the true reporting device. To add a syslog relay server, you must add it as a security software application running on a host.

To add a syslog relay server, follow these steps:

-
- Step 1** Select **Admin > System Setup > Security and Monitor Devices > Add**.
- Step 2** Do one of the following:
- Select **Add SW Security apps on a new host** from the Device Type list, and continue with [Step 3](#)

- Select **Add SW security apps on existing host** from the Device Type list. Select the device to which you want to add the software application and click Add. Continue with [Step 6](#).
- Step 3** Specify values for the following fields:
- **Device Name** — Enter the hostname of the syslog relay server. MARS maps this name to the reporting IP address. This name is used in topology maps, queries, and as the primary management station in the Security and Monitoring Device list.
 - **Reporting IP** — Enter the IP address of the interface in the syslog relay server from which MARS will receive syslog messages.

This address represents the physical IP address of the syslog relay server. To learn more about the reporting IP address, its role, and dependencies, see [Understanding Access IP, Reporting IP, and Interface Settings, page 2-8](#).
- Step 4** Under Enter interface information, enter the interface name, IP address, and netmask value of the interface in the syslog relay server from which syslog messages will be received.

This address represents the physical IP address of the syslog relay server. To learn more about the interface settings, its role, and dependencies, see [Understanding Access IP, Reporting IP, and Interface Settings, page 2-8](#).
- Step 5** Click **Apply** to save these settings.
- Step 6** Click **Next** to access the Reporting Applications tab.
- Step 7** Select **Generic Syslog Relay ANY** from the Select Application list, and click **Add**.
- Step 8** Click **Submit** to add this application to the host.

Result: Generic Syslog Relay ANY appears in the Device Type list.
- Step 9** Click the **Vulnerability Assessment Info** link to define the host information that MARS uses to determine false positive attacks against this host. Continue with [Define Vulnerability Assessment Information, page 10-11](#).
- Step 10** Click **Done** to save the changes.

Result: The host appears in the Security and Monitoring Information list.
- Step 11** To activate the device, click **Activate**.
-

Add Devices Monitored by Syslog Relay Server

While you do not have to configure each reporting device to forward syslog messages to the MARS Appliance, you must define the device to MARS so that when it parses the syslog messages forwarded by the relay server, then it is able to match the true reporting IP address to that of a known reporting device type. By knowing the reporting device type, MARS can correctly parse the events.

The process for adding these reporting devices is the same as if there were no syslog relay server except that you do not configure the reporting device to forward events to the MARS Appliance. In the MARS web interface, you should still configure the reporting devices so that MARS can discover their settings and to perform any mitigation operations.



Configuring Router and Switch Devices

This chapter describes how to bootstrap routers and switches and add those reporting devices and mitigation devices to MARS. It also describes how to configure NetFlow, NAC's EAP over UDP and 802.1x logging, and the Layer 2 (L2) mitigation features of switches.

Routers and switches provide MARS with data about traffic flows and the network topology, including address translations, endpoint devices, connected networks, and accepted and rejected sessions. Routers and switches also support modules that enable features common to specialty security appliances, such as firewalls and intrusion detection or prevention systems (IDS/IPS). This chapter does not describe how to enable the features on routers and switches that enable the modules or how to configure these modules for use by MARS. Such discussions are provided in [Configuring Firewall Devices, page 4-1](#), and [Configuring Network-based IDS and IPS Devices, page 6-1](#).

This chapter explains how to bootstrap and add the following router and switch devices to MARS:

- [Cisco Router Devices, page 3-1](#)
- [Cisco Switch Devices, page 3-9](#)
- [Extreme ExtremeWare 6.x, page 3-17](#)
- [Generic Router Device, page 3-18](#)

Cisco Router Devices

To configure Cisco routers running Cisco IOS Software Release 12.2 to communicate with a MARS Appliance, you must perform three tasks:

- [Enable Administrative Access to Devices Running Cisco IOS 12.2, page 3-1](#)
- [Configure the Device Running Cisco IOS 12.2 to Generate Required Data, page 3-3](#)
- [Add and Configure a Cisco Router in MARS, page 3-6](#)

Enable Administrative Access to Devices Running Cisco IOS 12.2

You must enable administrative access by the MARS Appliance to any Cisco routers or switches running Cisco IOS Software release 12.2 or later. The type of access that you must enable depends on whether modules are installed in your Cisco router or switch and the role of the device in your network. MARS uses this administrative access to discover the device's configuration and, at times, to make changes to the device's running configuration. For information on selecting an administrative access method, see [Selecting the Access Type, page 2-10](#).

Before you add a Cisco router to MARS, make sure that you have enabled SNMP, Telnet, SSH, or FTP access to the router. The following sections provide guidance on configuring each supported access method:

- [Enable SNMP Administrative Access, page 3-2](#)
- [Enable Telnet Administrative Access, page 3-2](#)
- [Enable SSH Administrative Access, page 3-2](#)
- [Enable FTP-based Administrative Access, page 3-2](#)

Enable SNMP Administrative Access

To enable configuration discovery using SNMP access to the Cisco router or switch, refer to your device documentation or the following URL:

http://cisco.com/en/US/products/sw/iosswrel/ps1835/products_configuration_guide_chapter09186a008030c762.html#wp1001217

Enable Telnet Administrative Access

To enable configuration discovery using Telnet access to the Cisco router or switch, refer to your device documentation or the following URL:

http://cisco.com/en/US/products/sw/iosswrel/ps1818/products_configuration_example09186a0080204528.shtml

Enable SSH Administrative Access

To enable configuration discovery using SSH access to the Cisco router or switch, refer to your device documentation or the following URL:

http://cisco.com/en/US/products/sw/iosswrel/ps1834/products_feature_guide09186a008007fed9.html

Enable FTP-based Administrative Access

To enable configuration discovery using FTP access, you must place a copy the Cisco router's or switch's configuration file on an FTP server to which the MARS Appliance has access. This FTP server must have user authentication enabled.



Note

TFTP is not supported. You must use an FTP server.

You must copy the running configuration from the Cisco router or switch. For information on copying the running configuration, refer to your device documentation or the following URL:

http://www.cisco.com/en/US/products/sw/iosswrel/ps1835/products_tech_note09186a008020260d.shtml

Configure the Device Running Cisco IOS 12.2 to Generate Required Data

Cisco routers and switches that are running Cisco IOS Software release 12.2 can be configured to provide different types of data to MARS:

- **Syslog messages.** The syslog messages provide information about activities on the network, including accepted and rejected sessions.
- **SNMP traffic.** SNMP RO community strings support the discovery of your network's topology.
- **NAC-specific data.** NAC logs events that are specific to its configuration, including Extensible Authentication Protocol (EAP) over UDP messages and 802.1x accounting messages.
- Access lists or NAT statements. You must enable SSH or Telnet access if the configuration on the Cisco router or switch includes access lists or NAT statements.
- **Spanning tree messages** (Switch only). You must have STP (spanning tree protocol) configured correctly on the switches to enable L2 discovery and mitigation. STP provides MARS with access to the L2 MIB, which is required to identify L2 re-routes of traffic and to perform L2 mitigation. MARS also uses the MIB to identify trunks to other switches, which are used to populate VLAN information used in L2 path calculations. STP, which is enabled by default on Cisco Switches, should remain enabled, as it is required for L2 mitigation.

The following topics describe how to configure these settings:

- [Enable Syslog Messages, page 3-3](#)
- [Enable SNMP RO Strings, page 3-3](#)
- [Enable NAC-specific Messages, page 3-4](#)
- [Enable L2 Discovery Messages, page 3-12](#)
- [Enable SDEE for IOS IPS Software, page 3-6](#)

Enable Syslog Messages

To send syslog messages to the MARS Appliance from a device running Cisco IOS Software Release 12.2, follow these steps:

Step 1 Log in to the Cisco IOS device with enabled password.

Step 2 Enter the commands:

```
Router(config)#logging source-interface <interface name>
Router(config)#logging trap <logging level desired>
Router(config)#logging <IP address of MARS Appliance>
```

Enable SNMP RO Strings

To enable SNMP RO strings for topology discovery on the Cisco IOS device, you must enable the SNMP server, define the RO community, and then direct the SNMP server to forward SNMP traps to the MARS Appliance.

To configure the SNMP RO string settings, follow these steps:

Step 1 Enter configuration mode:

```
Router> enable
Password: <password>
Router#
```

Step 2 Enter the **configure terminal** command to enter configuration mode:

```
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)#
```

Step 3 Set the SNMP read community string as follows:

```
Router(config)# snmp-server community <read community> RO <ACL name if required>
```

**Note**

This information is required to retrieve the MAC addresses and associated L2 information.

Step 4 Set the SNMP write community string as follows:

```
Router(config)# snmp-server community <write community> RW
```

Enable NAC-specific Messages

Cisco routers and switches that are running Cisco IOS Software release 12.2 or CatOS can enable network Admission Control (NAC) specific data. This data includes:

- **Client logs.** These logs relate the activities of the client software.
- **RADIUS server logs.** These logs relate the authorization communications between clients and the posture validation servers.
- **Network access device logs.** These logs relate connection attempts by clients and final authorizations provided by the AAA server enforcing the NAC policies.

For more information on the events that are logged as part of NAC, see the *Monitoring and Reporting Tool Integration into Network Admission Control* white paper at the following URL:

http://www.cisco.com/en/US/netsol/ns617/networking_solutions_white_paper0900aecd801dee49.shtml

This section contains the following two topics, which address the NAC configuration settings specific to each device type:

- [Cisco Routers, page 3-4](#)
- [Cisco Switches, page 3-5](#)

Cisco Routers

To configure the NAC Phase I data on a Cisco router to work with MARS, you must allow EAP over UDP and allow an IP address in the AAA station-id field of the packets. In addition, you must enable logging of these events, which are published as syslog messages.

To enable the NAC-specific data on a Cisco router, enter the following commands:

```
Router(config)#eou allow ip-station-id
Router(config)#eou logging
```

For more information on these commands and related commands, see the Network Admission Control feature document at the following URL:

http://www.cisco.com/univercd/cc/td/doc/product/software/ios123/123newft/123t/123t_8/gt_nac.htm

Cisco Switches

NAC Phase II enables Cisco switches to act as network access devices. To support this new feature, you must configure the Cisco switch to initiate 802.1x authentication when the link state changes from down to up and periodically if the port remains up but unauthenticated. NAC requires that hosts use 802.1x supplicants, or clients, to authenticate to the Cisco Secure ACS server before gaining access to network services. Enabling the 802.1x messages on your network helps you troubleshoot supplicant failures because connection attempts are logged, which you can analyze.

Configuring the Cisco switch to act as proxy between the Cisco Secure ACS server and the 802.1x supplicants is a multi-step process. First, the switch must be defined as a AAA client (RADIUS) in the Cisco Secure ACS server. For information on defining a AAA client, see [Define AAA Clients, page 14-5](#). Second, the switch must be configured to use a RADIUS server. Then, you must enable the following features on each interface installed in the switch:

- **802.1X port-based authentication.** The device requests the identity of the client and begins relaying authentication messages between the client and the authentication server. Each client attempting to access the network is uniquely identified by the system by using the client's MAC address.
- **802.1x reauthentication.** The device re-authenticates the supplicants after the reauthentication timeout value is reached, which is 3600 seconds by default.
- **802.1x accounting.** The device logs authentication successes and failures, as well as link down events and users logging off. The switch publishes these audit records to the Cisco Secure ACS server for logging.
- **DHCP snooping.** The device filters DHCP requests, safeguarding against spoof attacks. This feature ensures that MARS receives reliable data and identifies the port number of the 802.1x supplicant.

The following URLs detail how to configure these features:

Dot1x and Radius Sever

IOS Software:

<http://www.cisco.com/univercd/cc/td/doc/product/lan/cat3750/12225sec/3750scg/sw8021x.htm>

CatOS Software:

http://www.cisco.com/univercd/cc/td/doc/product/lan/cat6000/sw_8_4/config_gd/8021x.htm

DHCP Snooping

IOS Software:

<http://www.cisco.com/univercd/cc/td/doc/product/lan/cat3750/12225sec/3750scg/swdhcp82.htm>

CatOS Software:

http://www.cisco.com/univercd/cc/td/doc/product/lan/cat6000/sw_8_4/config_gd/dhcp.htm

After you configure the switch to act as proxy and it is defined as a AAA client in Cisco Secure ACS, you must ensure that the authentication messages are sent to the MARS Appliance. For 802.1x accounting records, you must ensure that the audit records are written to the RADIUS log on the Cisco Secure ACS server. To configure these settings, refer to [Configure Cisco Secure ACS to Generate Logs, page 14-3](#).

Enable SDEE for IOS IPS Software

Before you enable SDEE, you must enable either Telnet or SSH as the access type for configuration discovery on a Cisco IOS device. You must also enable SDEE on the device that supports the IOS IPS software feature. SDEE is used to publish events to MARS about signatures that have fired.

To enable SDEE protocol on the Cisco IOS device that supports IOS IPS, follow these steps:

-
- Step 1** Log in to the Cisco IOS device using the enable password.
- Step 2** Enter the following commands to enable MARS to retrieve events from the IOS IPS software:

```
Router(config)#ip http secure-server
Router(config)#ip ips notify sdee
Router(config)#ip sdee subscriptions 3
Router(config)#ip sdee events 1000
Router(config)#no ip ips notify log
```

**Note**

The “no ips notify log” causes the IOS IPS software to stop sending IPS events over syslog.

Add and Configure a Cisco Router in MARS

Cisco routers provide data about the network and its activities in the form of syslog messages and SNMP RO MIBs. In addition, MARS can discover settings, such as network address translations, attached networks, and active access rules, that improve the accuracy of false positive identification, attack path analysis, and L3 network discovery.

To add a Cisco router running Cisco IOS 12.2 or later, follow these steps:

-
- Step 1** Select **Admin > System Setup > Security and Monitor Devices > Add**.
- Step 2** Select **Cisco IOS 12.2** from the Device Type list.

Device Type:

→ *Device Name:

→ Access IP:

→ Reporting IP:

→ *Access Type:

Login:

Password:

Enable Password:

Config Path:

File Name:

SNMP RO Community:

→ Monitor Resource Usage:

143635

- Step 3** Enter the name of the device in the Device Name field.
- MARS maps this name to the reporting IP address. This name is used in topology maps, queries, and in the Security and Monitoring Device list. For devices that support the discovery operation, such as routers and firewalls, MARS renames this field's value to match the name discovered in the device configuration, which typically uses the *hostname.domain* format. For devices that cannot be discovered, such as Windows and Linux hosts and host applications, MARS uses the provided value.
- Step 4** (Optional) To enable MARS to discover settings from this device, enter the administrative IP address in the Access IP field.
- To learn more about the access IP address, its role, and dependencies, see [Understanding Access IP, Reporting IP, and Interface Settings, page 2-8](#).
- Step 5** Enter the IP address of the interface that publishes syslog messages, SNMP notifications, NetFlow MIBs, or any combination of the three, in the Reporting IP field.
- To learn more about the reporting IP address, its role, and dependencies, see [Understanding Access IP, Reporting IP, and Interface Settings, page 2-8](#).
- Step 6** If you entered an address in the Access IP field, select **SNMP**, **TELNET**, **SSH**, or **FTP** from the Access Type list, and continue with the procedure that matches your selection:
- [Configure SNMP Access for Devices in MARS, page 2-11](#)
 - [Configure Telnet Access for Devices in MARS, page 2-11](#)
 - [Configure SSH Access for Devices in MARS, page 2-12](#)
 - [Configure FTP Access for Devices in MARS, page 2-12](#)
- For more information on determining the access type, see [Selecting the Access Type, page 2-10](#).
- Step 7** (Optional) To enable MARS to retrieve MIB objects for this reporting device, enter the device's read-only community string in the SNMP RO Community field.

Before you can specify the SNMP RO string, you must define an access IP address. MARS uses the SNMP RO string to read MIBs related to a reporting device's CPU usage, network usage, and device anomaly data and to discover device and network settings.

- Step 8** (Optional) To enable MARS to monitor this device for anomalous resource usage, select **Yes** from the Monitor Resource Usage list.

Result: MARS monitors the device for anomalous consumption of resources, such as memory and CPU. If anomalies are detected, MARS generates an incident. Resource utilization statistics are also used to generate reports. For more information, see [Configuring Resource Usage Data, page 2-41](#).

- Step 9** (Optional) If this router has the IOS IPS feature and SDEE access enabled and you have configured the router to accept HTTPS connections from the MARS Appliance, click **Add IPS** to provide the username and password required to pull SDEE events.



Note

IOS IPS does *not* refer to an IPS module. It refers to a software feature in the IOS software. The IOS IPS feature is required to enable the DTM functionality in MARS. See [Technology Preview: Configuring Distributed Threat Mitigation with Intrusion Prevention System in Cisco Security MARS, page 1](#) for more information.

Result: The IOS IPS Information page appears.

IOS IPS Information

Reporting IP: 192.168.20.1

User Name:

password:

Port:

Test Connectivity Cancel Submit

- Enter the username that has HTTPS access to this device in the User Name field.
- Enter the corresponding password in the Password field.
- In the Port field, verify the port used for SDEE communications with this device.

MARS pulls data using SDEE over HTTPS. The default port number for HTTPS/SDEE is 443. This access allows MARS to retrieve XML files that contain the events generated by the IOS IPS feature.

Result: MARS can query the router for SDEE events.

- Step 10** (Optional) If you defined an access IP and selected and configured an access type, click **Discover** to determine the device settings, including the IOS IPS settings.

Result: If the username and password are correct and the MARS Appliance is configured as an administrative host for the device, the "Discovery is done." dialog box appears when the discovery operation completes. Otherwise, an error message appears. After the initial pull, the MARS Appliance pulls based on the schedule that you define. For more information, see [Scheduling Topology Updates, page 2-39](#).

- Step 11** To add this device to the MARS database, click **Submit**.

Result: The submit operation records the changes in the database tables. However, it does not load the changes into working memory of the MARS Appliance. The activate operation loads submitted changes into working memory.

Step 12 Click **Activate**.

Result: MARS begins to sessionize events generated by this device and evaluate those events using the defined inspection and drop rules. Any events published by the device to MARS before activation can be queried using the reporting IP address of the device as a match criterion. For more information on the activate action, see [Activate the Reporting and Mitigation Devices, page 2-27](#).

Cisco Switch Devices

You can manage Cisco switches that run either CatOS or Cisco IOS Software Release 12.2 or later. The configuration of the switch varies between these two operating system, as does the addition of the device in MARS. Adding a Cisco switch involves three steps:

1. Configure the switch to enable MARS to discover the its settings.
2. Configure the switch to generate the data required by MARS.
3. Add and configure the switch in MARS.
4. Add modules to the switch.

To prepare a Cisco switch running Cisco IOS Software Release 12.2 or later, refer to the following procedures:

- [Enable Administrative Access to Devices Running Cisco IOS 12.2, page 3-1](#)
- [Configure the Device Running Cisco IOS 12.2 to Generate Required Data, page 3-3](#)

To prepare a Cisco switch running CatOS, refer to the following procedures:

- [Enable Communications Between Devices Running CatOS and MARS, page 3-9](#)
- [Configure the Device Running CatOS to Generate Required Data, page 3-11](#)

Adding a Cisco switch running to MARS has two distinct steps. First, you add the base module of the switch, providing administrative access to that device. Second, you add any modules that are running in the switch. For instructions on performing these two steps, refer to the following topics:

- [Add and Configure a Cisco Switch in MARS, page 3-13](#)
- [Adding Modules to a Cisco Switch, page 3-14](#)

Enable Communications Between Devices Running CatOS and MARS

Before you add a Cisco switch running CatOS to MARS, make sure that you have enabled SNMP, Telnet, SSH, or FTP access to the switch. First, you must configure the MARS Appliance as an IP address that is permitted to access the switch.

For information on permitting IP addresses and specifying the access type, see the following URL:

http://www.cisco.com/univercd/cc/td/doc/product/lan/cat6000/sw_8_4/config_gd/ip_perm.htm#wp1019819

Next, you must ensure that your switch is configured to enable the correct access method. The following sections provide guidance on configuring each supported access method:

- [Enable SNMP Administrative Access, page 3-10](#)
- [Enable Telnet Administrative Access, page 3-10](#)
- [Enable SSH Administrative Access, page 3-10](#)
- [Enable FTP-based Administrative Access, page 3-10](#)

Enable SNMP Administrative Access

To enable configuration discovery using SNMP access to the Cisco switch, refer to your device documentation or the following URL:

IP Access

http://www.cisco.com/univercd/cc/td/doc/product/lan/cat6000/sw_8_4/config_gd/ip_perm.htm#wp1019819

Configure SNMP

http://www.cisco.com/univercd/cc/td/doc/product/lan/cat6000/sw_8_4/config_gd/snmp.htm

Enable Telnet Administrative Access

To enable configuration discovery using Telnet access to the Cisco switch, refer to your device documentation or the following URL:

IP Access

http://www.cisco.com/univercd/cc/td/doc/product/lan/cat6000/sw_8_4/config_gd/ip_perm.htm#wp1019819

Enable SSH Administrative Access

To enable configuration discovery using SSH access to the Cisco router or switch, refer to your device documentation or the following URL:

IP Access

http://www.cisco.com/univercd/cc/td/doc/product/lan/cat6000/sw_8_4/config_gd/ip_perm.htm#wp1019819

Enable FTP-based Administrative Access

To enable configuration discovery using FTP access, you must place a copy the Cisco router's or switch's configuration file on an FTP server to which the MARS Appliance has access. This FTP server must have user authentication enabled.



Note

TFTP is not supported. You must use an FTP server.

You must copy the running configuration from the Cisco switch. For information on copying the running configuration, refer to your device documentation or the following URL:

http://www.cisco.com/univercd/cc/td/doc/product/lan/cat6000/sw_8_4/config_gd/admin.htm#wp1040556

Configure the Device Running CatOS to Generate Required Data

You can configure the following message types:

- SNMP RO strings
- NAC messages (802.1x)
- L2 discover settings
- Syslog message

For information on configuring these settings, refer to the following topics:

- [Enable SNMP RO Strings on CatOS, page 3-11](#)
- [Enable NAC-specific Messages, page 3-4](#)
- [Enable L2 Discovery Messages, page 3-12](#)
- [Enable Syslog Messages on CatOS, page 3-11](#)

Enable SNMP RO Strings on CatOS

If the supervisor SNMP server is not configured, you must perform this procedure.

To configure the supervisor SNMP server and enabled SNMP traps on the Catalyst switch, follow these steps:

Step 1 Enter configuration mode:

```
switch> enable
Enter password: <password>
switch> (enable)
```

Step 2 Set the SNMP read community string as follows:

```
switch> (enable) set snmp community read-only <read community>
```

Step 3 Set the SNMP write community string as follows:

```
switch> (enable) set snmp community read-write <write community>
switch> (enable) set snmp community read-write-all <write community>
```

Step 4 To collect RMON Ethernet statistics, RMON data collection must be enabled in the CatOS agent (this is not required in Native IOS). To enable RMON collection, enter the following:

```
switch> (enable) set snmp rmon enable
```

Step 5 Exit configuration mode as follows:

```
switch> (enable) exit
```

Enable Syslog Messages on CatOS

To configure a Cisco switch running CatOS to send syslog information to MARS, follow these steps:

Step 1 To enable the syslog server on the switch, enter:

```
set logging server enable
```

Step 2 To identify the MARS Appliance as a destination for syslog messages, enter the following command:

```
set logging server <IP address of MARS Appliance>
```

Step 3 The remaining commands tell the switch what kinds of logging information to provide and at what level. The commands in the following example can be changed to suit your requirements.

```
set logging level cdp 7 default
set logging level mcast 7 default
set logging level dtp 7 default
set logging level dvlan 7 default
set logging level earl 7 default
set logging level fddi 7 default
set logging level ip 7 default
set logging level pruning 7 default
set logging level snmp 7 default
set logging level spantree 7 default
set logging level sys 7 default
set logging level tac 7 default
set logging level tcp 7 default
set logging level telnet 7 default
set logging level tftp 7 default
set logging level vtp 7 default
set logging level vmpls 7 default
set logging level kernel 7 default
set logging level filesys 7 default
set logging level drip 7 default
set logging level pagp 7 default
set logging level mgmt 7 default
set logging level mls 7 default
set logging level protfilt 7 default
set logging level security 7 default
set logging server facility SYSLOG
set logging server severity 7
set logging buffer 250
set logging timestamp enable
```

Enable L2 Discovery Messages

To enable L2 discovery on your Cisco switches, you must enable the spanning tree protocol (STP) and provide the SNMP RO community string. All L 2 devices must support SNMP STP MIB (IETF RFC 1493). The discovered information includes interfaces, Layer 3 (L3) routes, L2 spanning trees, L2 forwarding tables, MAC addresses, and so on.



Note

STP is enabled by default on all Cisco switches. Therefore, unless you have altered this setting, no changes are necessary.

For more information on configuring STP, select **Spanning Tree Protocol** in the View Documents by topics list at the following URLs:

http://www.cisco.com/en/US/partner/products/hw/switches/ps708/prod_configuration_examples_list.html

http://www.cisco.com/univercd/cc/td/doc/product/lan/cat6000/sw_8_4/config_gd/spantree.htm

Add and Configure a Cisco Switch in MARS

MARS monitors Cisco switches running either CatOS or Cisco IOS 12.2.

To add the configuration information that MARS uses to monitor a Cisco switch running Cisco IOS 12.2 or later, follow these steps:

-
- Step 1** Select **Admin > System Setup > Security and Monitor Devices > Add**.
- Step 2** Do one of the following:
- If the switch is running any version of CatOS, select **Cisco Switch-CatOS ANY** from the Device Type list.
 - If the switch is running Cisco IOS 12.2 or later, select **Cisco Switch-IOS 12.2** from the Device Type list.
- Step 3** Enter the name of the device in the Device Name field.
- MARS maps this name to the reporting IP address. This name is used in topology maps, queries, and in the Security and Monitoring Device list. For devices that support the discovery operation, such as routers and firewalls, MARS renames this field's value to match the name discovered in the device configuration, which typically uses the *hostname.domain* format. For devices that cannot be discovered, such as Windows and Linux hosts and host applications, MARS uses the provided value.
- Step 4** (Optional) To enable MARS to discover settings from this device, enter the administrative IP address in the Access IP field.
- To learn more about the access IP address, its role, and dependencies, see [Understanding Access IP, Reporting IP, and Interface Settings, page 2-8](#).
- Step 5** Enter the IP address of the interface that publishes syslog messages, SNMP notifications, NetFlow MIBs, or any combination of the three, in the Reporting IP field.
- To learn more about the reporting IP address, its role, and dependencies, see [Understanding Access IP, Reporting IP, and Interface Settings, page 2-8](#).
- Step 6** If you entered an address in the Access IP field, select **SNMP**, **TELNET**, **SSH**, or **FTP** from the Access Type list, and continue with the procedure that matches your selection:
- [Configure SNMP Access for Devices in MARS, page 2-11](#)
 - [Configure Telnet Access for Devices in MARS, page 2-11](#)
 - [Configure SSH Access for Devices in MARS, page 2-12](#)
 - [Configure FTP Access for Devices in MARS, page 2-12](#)
- For more information on determining the access type, see [Selecting the Access Type, page 2-10](#).
- Step 7** (Optional) To enable MARS to retrieve MIB objects for this reporting device, enter the device's read-only community string in the SNMP RO Community field.
- Before you can specify the SNMP RO string, you must define an access IP address. MARS uses the SNMP RO string to read MIBs related to a reporting device's CPU usage, network usage, and device anomaly data and to discover device and network settings.
- Step 8** (Optional) To enable MARS to monitor this device for anomalous resource usage, select **Yes** from the Monitor Resource Usage list.
- Result:* MARS monitors the device for anomalous consumption of resources, such as memory and CPU. If anomalies are detected, MARS generates an incident. Resource utilization statistics are also used to generate reports. For more information, see [Configuring Resource Usage Data, page 2-41](#).

- Step 9** (Optional) If you defined an access IP and selected and configured an access type, click **Discover** to determine the device settings
- Result:* If the username and password are correct and the MARS Appliance is configured as an administrative host for the device, the “Discovery is done.” dialog box appears when the discovery operation completes. Otherwise, an error message appears. After the initial pull, the MARS Appliance pulls based on the schedule that you define. For more information, see [Scheduling Topology Updates, page 2-39](#).
- Step 10** To add this device to the MARS database, click **Submit**.
- Result:* The submit operation records the changes in the database tables. However, it does not load the changes into working memory of the MARS Appliance. The activate operation loads submitted changes into working memory.
- Step 11** Click **Activate**.
- Result:* MARS begins to sessionize events generated by this device and evaluate those events using the defined inspection and drop rules. Any events published by the device to MARS before activation can be queried using the reporting IP address of the device as a match criterion. For more information on the activate action, see [Activate the Reporting and Mitigation Devices, page 2-27](#).
- After submitting, you can add modules. See [Adding Modules to a Cisco Switch, page 3-14](#).
-

Adding Modules to a Cisco Switch

In MARS, you can represent, discover, and monitor modules that are installed in Cisco switches. These modules perform special purpose security functions for the switch, such as firewall or intrusion detection and prevention. MARS recognizes the following switch modules and versions:

- Cisco FWSM 1.1, 2.2, and 2.3
- Cisco IDS 3.1 and 4.0
- Cisco IPS 5.x
- Cisco IOS 12.2

To add a module, you must first add the base module, which is the Cisco switch. After the base module is defined in the web interface, you can discover the modules that are installed in the switch (click **Add Available Module**) or add them manually (click **Add Module**).

For instructions on adding and configuring a firewall services module (FWSM), see [Cisco Firewall Devices \(PIX, ASA, and FWSM\), page 4-1](#).

For instructions on adding and configuring an intrusion detection or prevention services module (IDSM or IPSM), see [Cisco IPS Modules, page 6-9](#).

This section contains the following topics:


- [Add Available Modules, page 3-14](#)
- [Add Cisco IOS 12.2 Modules Manually, page 3-15](#)

Add Available Modules

When you perform a discovery operation on a base module, MARS lists the discovered modules. From this list, you can select the modules to monitor using MARS.

To add available modules, follow these steps:

Step 1 Click **Add Available Module**.



The screenshot shows a web interface with four buttons at the top: 'Add Module', 'Edit Module', 'Remove Module', and 'Add Available Module'. Below the buttons is a table with two columns: 'Module Name' and 'Module Type'. The table contains two rows of data, each with a checkbox in the first column.

Module Name	Module Type
☐ HQ-SW-1-msfc	Cisco IOS 12.2
☐ HQ-SW-1-idsm	Cisco IDS 3.1

143216

If modules are installed in the switch, a list of the modules appears.

Step 2 Select a module from the Select list.

Step 3 Click **Add**.

Step 4 Repeat for other modules.

Step 5 After you add the desired modules, verify the configuration information of each. For example, verify that the SNMP RO community string matches that defined for use by MARS. To verify these settings, select a module and click **Edit Module**.

Basic guidance for editing these settings can be found in the topics that discuss manually adding these modules. See the following topics for more information:

- [Add Cisco IOS 12.2 Modules Manually, page 3-15](#)
- [Cisco Firewall Devices \(PIX, ASA, and FWSM\), page 4-1](#)
- [Cisco IPS Modules, page 6-9.](#)

Step 6 To add these modules to the base module defined in the MARS database, click **Submit**.

Result: The submit operation records the changes in the database tables. However, it does not load the changes into working memory of the MARS Appliance. The activate operation loads submitted changes into working memory.

Step 7 Click **Activate**.

Result: MARS begins to sessionize events generated by this device and the selected modules and evaluate those events using the defined inspection and drop rules. Any events published by the device or its modules to MARS before activation can be queried using the reporting IP address of the device or module as a match criterion. For more information on the activate action, see [Activate the Reporting and Mitigation Devices, page 2-27](#).

Add Cisco IOS 12.2 Modules Manually

To add a module manually, follow these steps:

Step 1 Click **Add Module**.

Step 2 Select **Cisco IOS 12.2** from the Device Type list.

Device Type: Cisco FWSM 1.1

→ *Device Name: [Text Field]

→ Access IP: [Text Field]

→ Reporting IP: [Text Field]

→ *Access Type: Select [3DES]

Login: [Text Field]

Password: [Text Field]

Enable Password: [Text Field]

Config Path: [Text Field]

File Name: [Text Field]

SNMP RO Community: [Text Field]

→ Monitor Resource Usage: NO

- Step 3** Enter the name of the module in the Device Name field.
- MARS maps this name to the reporting IP address. This name is used in topology maps, queries, and in the Security and Monitoring Device list. For modules that support the discovery operation, such as router and firewall modules, MARS renames this field's value to match the name discovered in the device configuration, which typically uses the *hostname.domain* format.
- Step 4** (Optional) To enable MARS to discover settings from this device, enter the administrative IP address in the Access IP field.
- To learn more about the access IP address, its role, and dependencies, see [Understanding Access IP, Reporting IP, and Interface Settings, page 2-8](#).
- Step 5** Enter the IP address of the interface that publishes syslog messages, SNMP notifications, NetFlow MIBs, or any combination of the three, in the Reporting IP field.
- To learn more about the reporting IP address, its role, and dependencies, see [Understanding Access IP, Reporting IP, and Interface Settings, page 2-8](#).
- Step 6** If you entered an address in the Access IP field, select **TELNET**, **SSH**, or **FTP** from the Access Type list, and continue with the procedure that matches your selection:
- [Configure Telnet Access for Devices in MARS, page 2-11](#)
 - [Configure SSH Access for Devices in MARS, page 2-12](#)
 - [Configure FTP Access for Devices in MARS, page 2-12](#)
- For more information on determining the access type, see [Selecting the Access Type, page 2-10](#).
- Step 7** (Optional) To enable MARS to retrieve MIB objects for this reporting device, enter the device's read-only community string in the SNMP RO Community field.
- Before you can specify the SNMP RO string, you must define an access IP address. MARS uses the SNMP RO string to read MIBs related to a reporting device's CPU usage, network usage, and device anomaly data and to discover device and network settings.
- Step 8** (Optional) To enable MARS to monitor this device for anomalous resource usage, select **Yes** from the Monitor Resource Usage list.

Result: MARS monitors the module for anomalous consumption of resources, such as memory and CPU. If anomalies are detected, MARS generates an incident. Resource utilization statistics are also used to generate reports. For more information, see [Configuring Resource Usage Data, page 2-41](#).

- Step 9** (Optional) If you defined an access IP and selected and configured an access type, click **Discover** to determine the module settings.

Result: If the username and password are correct and the MARS Appliance is configured as an administrative host for the module, the “Discovery is done.” dialog box appears when the discovery operation completes. Otherwise, an error message appears. After the initial pull, the MARS Appliance pulls based on the schedule that you define. For more information, see [Scheduling Topology Updates, page 2-39](#).

- Step 10** To add this module to the device in the MARS database, click **Submit**.

Result: The submit operation records the changes in the database tables. However, it does not load the changes into working memory of the MARS Appliance. The activate operation loads submitted changes into working memory.

Extreme ExtremeWare 6.x

MARS can use Extreme ExtremeWare switches to enforce L2 mitigation. To configure MARS to communicate with an ExtremeWare switch, you must configure the switch to publish SNMP notifications to the MARS Appliance. In addition, you must add and configure the switch in the web interface.

This section contains the following topics:

- [Configure ExtremeWare to Generate the Required Data, page 3-17](#)
- [Add and Configure an ExtremeWare Switch in MARS, page 3-18](#)

Configure ExtremeWare to Generate the Required Data

To bootstrap an ExtremeWare switch, you must configure two features. First, you must configure the switch to send syslog messages to the MARS Appliance. Next, you must configure the SNMP RO community for MARS to access available L2 information.

To prepare the ExtremeWare device to generate the data required by MARS, follow these steps:

- Step 1** For syslog configuration, add this command:

```
configure syslog add <MARS's IP address> local7 debug
enable syslog
```

- Step 2** For SNMP configuration add these commands:

```
enable snmp dot1dTpFdbTable
configure snmp delete community readonly all
configure snmp delete community readwrite all
configure snmp add community readonly encrypted <encrypted community string>
configure snmp add community readwrite encrypted <encrypted community string>
```

Add and Configure an ExtremeWare Switch in MARS

To add and configure an ExtremeWare switch in MARS, follow these steps:

-
- Step 1** Select **Admin > System Setup > Security and Monitor Devices > Add**.
- Step 2** Select **Extreme ExtremeWare 6.x** from the Device Type list.
- Step 3** Enter the name of the device in the Device Name field.
- MARS maps this name to the reporting IP address. This name is used in topology maps, queries, and in the Security and Monitoring Device list. For devices that support the discovery operation, such as routers and firewalls, MARS renames this field's value to match the name discovered in the device configuration, which typically uses the *hostname.domain* format. For devices that cannot be discovered, such as Windows and Linux hosts and host applications, MARS uses the provided value.
- Step 4** (Optional) To enable MARS to discover settings from this device, enter the administrative IP address in the Access IP field.
- To learn more about the access IP address, its role, and dependencies, see [Understanding Access IP, Reporting IP, and Interface Settings, page 2-8](#).
- Step 5** Enter the IP address of the interface that publishes syslog messages, SNMP notifications, or both in the Reporting IP field.
- To learn more about the reporting IP address, its role, and dependencies, see [Understanding Access IP, Reporting IP, and Interface Settings, page 2-8](#).
- Step 6** If you entered an address in the Access IP field, select **SNMP** from the Access Type list.
- For more information on understanding the access type, see [Selecting the Access Type, page 2-10](#).
- Step 7** (Optional) To enable MARS to retrieve MIB objects for this reporting device, enter the device's read-only community string in the SNMP RO Community field.
- Before you can specify the SNMP RO string, you must define an access IP address. MARS uses the SNMP RO string to read MIBs related to a reporting device's CPU usage, network usage, and device anomaly data and to discover device and network settings.
- Step 8** To add this device to the MARS database, click **Submit**.
- Result:* The submit operation records the changes in the database tables. However, it does not load the changes into working memory of the MARS Appliance. The activate operation loads submitted changes into working memory.
- Step 9** Click **Activate**.
- Result:* MARS begins to sessionize events generated by this device and evaluate those events using the defined inspection and drop rules. Any events published by the device to MARS before activation can be queried using the reporting IP address of the device as a match criterion.
-

Generic Router Device

You can add any L2 or L3 device to the MARS as long as SNMP is enabled on the device. A generic router refers to any L2 or L3 device that is not listed in the *Supported Devices and Software Versions for CS-MARS Local Controller 4.1*.

Add and Configure a Generic Router in MARS

To add and configure a generic router device in MARS, follow these steps:

-
- Step 1** Select **Admin > System Setup > Security and Monitor Devices > Add**.
- Step 2** Select **Generic Router version unknown** from the Device Type list.
- Step 3** Enter the name of the device in the Device Name field.
- MARS maps this name to the reporting IP address. This name is used in topology maps, queries, and in the Security and Monitoring Device list. For devices that support the discovery operation, such as routers and firewalls, MARS renames this field's value to match the name discovered in the device configuration, which typically uses the *hostname.domain* format. For devices that cannot be discovered, such as Windows and Linux hosts and host applications, MARS uses the provided value.
- Step 4** (Optional) To enable MARS to discover settings from this device, enter the administrative IP address in the Access IP field.
- To learn more about the access IP address, its role, and dependencies, see [Understanding Access IP, Reporting IP, and Interface Settings, page 2-8](#).
- Step 5** Enter the IP address of the interface that publishes syslog messages, SNMP notifications, or both in the Reporting IP field.
- To learn more about the reporting IP address, its role, and dependencies, see [Understanding Access IP, Reporting IP, and Interface Settings, page 2-8](#).
- Step 6** If you entered an address in the Access IP field, select **SNMP** from the Access Type list.
- For more information on understanding the access type, see [Selecting the Access Type, page 2-10](#).
- Step 7** (Optional) To enable MARS to retrieve MIB objects for this reporting device, enter the device's read-only community string in the SNMP RO Community field.
- Before you can specify the SNMP RO string, you must define an access IP address. MARS uses the SNMP RO string to read MIBs related to a reporting device's CPU usage, network usage, and device anomaly data and to discover device and network settings.
- Step 8** To add this device to the MARS database, click **Submit**.
- Result:* The submit operation records the changes in the database tables. However, it does not load the changes into working memory of the MARS Appliance. The activate operation loads submitted changes into working memory.
- Step 9** Click **Activate**.
- Result:* MARS begins to sessionize events generated by this device and evaluate those events using the defined inspection and drop rules. Any events published by the device to MARS before activation can be queried using the reporting IP address of the device as a match criterion.
-



Configuring Firewall Devices

This chapter describes how to bootstrap firewall devices and add them to MARS as reporting devices. Firewall devices come in several form factors: hardware appliances, software applications running on a host, modules that are installed in switches and routers, and modules that install in multifunction security devices.

Multifunction security devices, such as the Cisco Adaptive Security Appliance (ASA), also support non-firewall modules, such as intrusion detection or prevention systems (IDS/IPS). This chapter does not focus on configuring non-firewalling modules. Such discussions are provided in [Configuring Network-based IDS and IPS Devices](#), page 6-1.

This chapter explains how to bootstrap and add the following firewall devices to MARS:

- [Cisco Firewall Devices \(PIX, ASA, and FWSM\)](#), page 4-1
- [NetScreen ScreenOS Devices](#), page 4-11
- [Check Point Devices](#), page 4-19

Cisco Firewall Devices (PIX, ASA, and FWSM)

MARS support for Cisco firewall devices includes the following:

- PIX Security Appliance, including PIX software releases 6.0, 6.1, 6.2, 6.3 and 7.0.
- Cisco Adaptive Security Appliance (ASA) 7.0
- Cisco Firewall Services Modules (FWSM), versions 1.1, 2.2, and 2.3.

Because these PIX software is mostly backward compatible, the commands required to bootstrap PIX security appliance remain consistent across the releases. In addition, Cisco ASA and FWSM have much in common with PIX command set.

The taskflow required to configure MARS to monitor a Cisco firewall device is as follows:

1. Configure the Cisco firewall device to accept administrative sessions from MARS (to discover settings).

For Cisco ASA, PIX 7.0, and FWSM device types, you configure the admin context to accept these sessions.



Note

To be monitored by MARS, the Cisco ASA, PIX 7.0, and FWSM device types have the following two requirements: each context requires a unique routable IP address for sending syslog messages to MARS, and each context must have a unique name (hostname+ domain name).

2. Configure the Cisco firewall device to publish its syslog events to MARS.

For Cisco ASA, PIX 7.0, and FWSM device types, you must configure the admin context and each security context.

**Note**

MARS uses syslog events to discover information about the network topology. It uses SNMP to discover CPU utilization and related information.

3. Within MARS, define the Cisco firewall device by providing the administrative connection information.

**Note**

Before you can add an FWSM module in a switch, you must add and configure the base module (the Cisco switch) in MARS. For more information, [Cisco Switch Devices, page 3-9](#).

For Cisco ASA, PIX 7.0, and FWSM, the basic device type represents the admin context. However you must also define or discover each security context and any installed Advanced Inspection and Prevention (AIP) modules running IPS 5.0.

To configure MARS to accept syslog event data and to pull device configurations settings from a Cisco firewall device, you must perform the following tasks:

- [Bootstrap the Cisco Firewall Device, page 4-2](#)
- [Add and Configure a Cisco Firewall Device in MARS, page 4-5](#)

Bootstrap the Cisco Firewall Device

You should configure your Cisco firewall devices to act as reporting devices and manual mitigation devices because they perform multiple roles on your network. MARS can benefit from the proper configuration of specific features:

- **IDS/IPS signature detection.** While it does not boast the most efficient or comprehensive set of signatures, the built-in IDS and IPS signature matching features of the Cisco firewall device can be critical in detecting an attempted attack.
- **Accept/Deny Logs.** The logging of accepted as well as denied sessions aids in false positive analysis.
- **Administrative Access.** Administrative access ensure MARS access to several key pieces of data:
 - *Route and ARP tables*, which aid in network discovery and MAC address mapping.
 - *NAT and PAT translation tables*, which aid in address resolution and attack path analysis, exposing the real instigator of attacks.
 - *OS Settings*, from which MARS determines the correct ACLs to block detected attacks, which paste into a management session with the Cisco firewall device.

To bootstrap the Cisco firewall device, you must identify the MARS Appliance as an administrative host. Enabling administrative access allows MARS to discover the Cisco firewall device configuration settings. To enable administrative access, you must make sure that the MARS Appliance is granted Telnet or SSH administrative access to the firewall device. If you use FTP access type, make sure that you have added its configuration file to an FTP server to allow MARS access to the FTP server.

In addition to configuring specific event types and administrative access, syslog messages should be sent to the MARS Appliance. To prepare the Cisco firewall device to send these messages to the MARS Appliance, you must configure the logging settings associated with each firewall device on your network. To prepare a firewall device to generate the syslog messages and direct them to a specific MARS Appliance, you must:

1. Enable logging on the firewall device.

Before a firewall device can generate syslog messages, you must enable logging for one or more interfaces. In addition, if you configured your firewall device in a failover pair, you can specify the standby firewall device to generate syslog messages as well. You can enable the device to ensure that the standby unit's syslog messages stay synchronized if failover occurs. However, this option results in twice as much traffic on the MARS Appliance.

2. Select the log facility and queue size.

To generate meaningful reports about the network activity of a firewall device and to monitor the security events associated with that device, you must select the appropriate logging level. The logging level generates the syslog details required to track session-specific data. After you select a logging level, you can define a syslog rule that directs traffic to the MARS Appliance.

3. Select the log level to debug.

This setting generates syslog messages that assist you in debugging. It also generates logs that identify the commands issued during FTP sessions and the URLs requested during HTTP sessions. It includes all emergency, alert, critical, error, warning, notification, and information messages.

**Note**

Full URLs, such as `www.cisco.com/foo.html`, are included in HTTP session logs and FTP command data is logged only if web filtering (N2H2\SecureComputing or WebSense) is enabled on the reporting device. If web filtering is not enabled, then the HTTP session log does not include the hostname (although the destination host's IP and the Request-URI are included, such as `192.168.1.1:/foo.htm`) and FTP command data is not logged at all. Caveats exist with HTTP session logging, such as if the HTTP session request is broken across packets, then the hostname data might not be included in the log data.

4. Identify the target MARS Appliance and the protocol and port pair that it listens on.

By directing syslog messages generated by a firewall device to MARS, you can process and study the messages.

**Tip**

When monitoring a failover pair of Cisco firewall devices, you should designate the primary Cisco firewall device as the device to be monitored. If failover occurs, the secondary device assumes the IP address of the primary, which ensures that session correlation is maintained after the failover. The same focus on the primary is true for performing any bootstrap operations. The secondary device will synchronize with the configuration settings of the primary.

To enable administrative connections to the firewall device, select from the following options:

- [Enable Telnet Access on a Cisco Firewall Device, page 4-4](#)
- [Enable SSH Access on a Cisco Firewall Device, page 4-4](#)
- [Send Syslog Files From Cisco Firewall Device to MARS, page 4-4](#)

To configure log settings, see [Send Syslog Files From Cisco Firewall Device to MARS, page 4-4](#).

Enable Telnet Access on a Cisco Firewall Device

-
- Step 1** Log in to the Cisco firewall device with administrator's privileges.
- Step 2** Enter the command:
- ```
telnet <MARS IP address> <netmask of MARS IP address> <interface name>
```
- where *interface* name can be inside, outside, DMZ.
- 

## Enable SSH Access on a Cisco Firewall Device

- 
- Step 1** Log in to the Cisco firewall device with administrator's privileges.
- Step 2** Enter the command:
- ```
ssh <MARS IP address> <netmask of the MARS IP address> <interface name>
```
- where *interface* name can be inside, outside, DMZ.
-

Send Syslog Files From Cisco Firewall Device to MARS

When preparing a Cisco firewall device to publish syslog messages, consider the following restrictions:

- **Do not** customize the priority of any syslog messages. If you do, MARS fails to parse those messages.
- **Do not** configure EMBLEM format for syslog messages. Make sure neither of the following commands are used in the configuration:

```
logging host <interface name> <PN-MARS's IP address> format EMBLEM
logging emblem
```

To send syslog messages to the MARS Appliance, you must enable logging, select the log facility and queue size, and specify the log level to debug.

-
- Step 1** Log in to the Cisco firewall device with administrator's privileges.
- Step 2** To enable logging, enter one of the following commands:
- (PIX and Cisco ASA) **logging enable**
 - (FWSM) **logging on**
- Step 3** To specify the MARS Appliance as a target logging host, enter the following command:
- ```
logging host <interface name> <MARS IP address>
```
- Step 4** To set the log level to debug, which ensures that HTTP and FTP session logs are generated, enter the following command:
- ```
logging trap debugging
```

The debug messages contain the HTTP URL address information. Therefore, you can create keyword-based rules matching against the firewall message itself. For example, if the debug messages are enabled and users were logging to “http://mail.cisco.com”, you could create keyword-based rules that matched against “mail.yahoo.com.”

**Note**

Full URLs, such as `www.cisco.com/foo.html`, are included in HTTP session logs and FTP command data is logged only if web filtering (N2H2\SecureComputing or WebSense) is enabled on the reporting device. If web filtering is not enabled, then the HTTP session log does not include the hostname (although the destination host's IP and the Request-URI are included, such as `192.168.1.1:/foo.htm`) and FTP command data is not logged at all. Caveats exist with HTTP session logging, such as if the HTTP session request is broken across packets, then the hostname data might not be included in the log data.

Debug messages are also preferred for troubleshooting. You can define inspection rules that match on on debug-level keywords, which send notifications to the appropriate group. Refer to PIX debug messages for interesting keywords.

Cisco recommends enabling debug for optimal use of your STM solution. If a Cisco firewall device is unable to sustain debug-level messages due to performance reasons, the informational level should be used. In non-debug mode, the URL information is not available; only the 5 tuple is available for queries and reports.

Step 5 For FWSM, enter the following command:

```
logging rate-limit <eps rate desired> 1
```

Step 6 For Cisco ASA, PIX 7.0, and FWSM, repeat [Step 2](#) through [Step 5](#) for each context defined, admin and security.

Step 7 (Cisco ASA only) If an Advanced Inspection and Prevention (AIP) module is installed, you need to prepare that module as you would any IPS 5.0 module. For more information, see [Cisco IPS Modules, page 6-9](#).

Add and Configure a Cisco Firewall Device in MARS

The process of adding a PIX security appliance, Cisco ASA, or FWSM to MARS involves many of the same steps, regardless of the version of software that is running. The process is exactly the same for PIX software versions 6.0, 6.1, 6.2, and 6.3. However, Cisco ASA, PIX 7.0, and FWSM provide the ability to define multiple security contexts, or virtual firewalls.

Adding a Cisco ASA, PIX 7.0, and FWSM to MARS has two distinct steps. First, you must define the settings for the admin context. Then, if multiple context mode is enabled, you define or discover the settings for its security contexts. These Cisco firewall device have two type of contexts: one admin context, which is used for configuration of the device itself, and one or more security contexts. For Cisco ASA, you can also define or discover any modules that are installed in the appliance.

To be monitored by MARS, the Cisco ASA, PIX 7.0, and FWSM device types have the following additional requirements:

- each context requires a unique routable IP address for sending syslog messages to MARS
- each context must have a unique name (hostname+ domain name)

**Note**

The Cisco ASA, PIX 7.0, and FWSM can run in single context mode, which means that the system context acts as both the admin context and a security context.

To add and configure a Cisco firewall device, follow these steps:

Step 1 Do one of the following:

- If you are adding an FWSM, you must be on the main page of the Cisco switch to which you are adding it. On that page, click **Add Module**.
- If you are adding a PIX security appliance or a Cisco ASA, an Select **Admin > System Setup > Security and Monitor Devices > Add**.

Step 2 Select one of the following options from the Device Type list.

- Cisco PIX 6.0
- Cisco PIX 6.1
- Cisco PIX 6.2
- Cisco PIX 6.3
- Cisco PIX 7.0
- Cisco ASA 7.0
- Cisco FWSM 1.1
- Cisco FWSM 2.2
- Cisco FWSM 2.3

Device Type:

→ *Device Name:	Admin
→ *Access IP:	10 1 1 23
→ *Reporting IP:	10 1 1 23
→ *Access Type:	SSH 3DES
Login:	pix
Password:	*****
Enable Password:	*****
Config Path:	
File Name:	
SNMP RO Community:	public

143178

Step 3 Enter the name of the firewall device in the Device Name field.

MARS maps this name to the reporting IP address. This name is used in topology maps, queries, and in the Security and Monitoring Device list. For devices that support the discovery operation, such as routers and firewalls, MARS renames this field's value to match the name discovered in the device configuration, which typically uses the *hostname.domain* format. For devices that cannot be discovered, such as Windows and Linux hosts and host applications, MARS uses the provided value.

- Step 4** (Optional) To enable MARS to discover settings from this firewall device, enter the administrative IP address in the Access IP field.



Note If the device is running Cisco ASA, PIX 7.0, or FWSM, this address corresponds to IP address from which the syslog messages of the admin context are sent.

To learn more about the access IP address, its role, and dependencies, see [Understanding Access IP, Reporting IP, and Interface Settings, page 2-8](#).

- Step 5** Enter the IP address of the interface that publishes syslog messages or SNMP notifications, or both in the Reporting IP field.



Note If the device is running Cisco ASA, PIX 7.0, or FWSM, this address corresponds to the address from which the admin context syslog messages are published.

To learn more about the reporting IP address, its role, and dependencies, see [Understanding Access IP, Reporting IP, and Interface Settings, page 2-8](#).

- Step 6** If you entered an address in the Access IP field, select **TELNET**, **SSH**, or **FTP** from the Access Type list, and continue with the procedure that matches your selection:

- [Configure Telnet Access for Devices in MARS, page 2-11](#)
- [Configure SSH Access for Devices in MARS, page 2-12](#)
- [Configure FTP Access for Devices in MARS, page 2-12](#)



Note If you select the FTP access type and you are defining a Cisco ASA, PIX 7.0, or FWSM, you cannot discover the non-admin context settings. Therefore, this access type is not recommended.

For more information on determining the access type, see [Selecting the Access Type, page 2-10](#).

- Step 7** (Optional) To enable MARS to retrieve MIB objects for this reporting device, enter the device's read-only community string in the SNMP RO Community field.

Before you can specify the SNMP RO string, you must define an access IP address. MARS uses the SNMP RO string to read MIBs related to a reporting device's CPU usage, network usage, and device anomaly data and to discover device and network settings.

- Step 8** (Optional) To enable MARS to monitor this device for anomalous resource usage, select **Yes** from the Monitor Resource Usage list.

Result: MARS monitors the device for anomalous consumption of resources, such as memory and CPU. If anomalies are detected, MARS generates an incident. Resource utilization statistics are also used to generate reports. For more information, see [Configuring Resource Usage Data, page 2-41](#).

- Step 9** (Cisco ASA, FWSM, and PIX 7.0 Only) do one of the following:

- Click **Discover** to let MARS contact the device and conduct a topology and context configuration discovery. Information about the security contexts is presented in the Context section of the main page. To edit discovered contexts, continue with [Edit Discovered Security Contexts, page 4-11](#).
- Click **Next** to commit your changes and allow for manual definition of security contexts or modules. Continue with [Add Security Contexts Manually, page 4-8](#), [Add Discovered Contexts, page 4-10](#), or [Add an IPS Module to a Cisco Switch or Cisco ASA, page 6-11](#).

For PIX and FWSM, you can add one or more security contexts. For Cisco ASA, you can add one or more security contexts or Advanced Inspection and Prevention (AIP) modules, running the Cisco IPS 5.x software.

Device Type: Cisco PIX 7.0

→ *Device Name:

→ *Access IP:

→ *Reporting IP:

→ *Access Type:

Login:

Password:

Enable Password:

Config Path:

File Name:

SNMP RO Community:

- Step 10** (Optional) If you defined an access IP and selected and configured an access type, click **Discover** to determine the device settings, including any security contexts and their settings.

Result: If the username and password are correct and the MARS Appliance is configured as an administrative host for the device, the “Discovery is done.” dialog box appears when the discovery operation completes. Otherwise, an error message appears. After the initial pull, the MARS Appliance pulls based on the schedule that you define. For more information, see [Scheduling Topology Updates](#), page 2-39.

- Step 11** To add this device to the MARS database, click **Submit**.

Result: The submit operation records the changes in the database tables. However, it does not load the changes into working memory of the MARS Appliance. The activate operation loads submitted changes into working memory.

- Step 12** Click **Activate**.

Result: MARS begins to sessionize events generated by this device and evaluate those events using the defined inspection and drop rules. Any events published by the device to MARS before activation can be queried using the reporting IP address of the device as a match criterion. For more information on the activate action, see [Activate the Reporting and Mitigation Devices](#), page 2-27.

Add Security Contexts Manually

You can manually define security contexts in PIX 7.0, Cisco ASA, or FWSM.

Step 1 Do one of the following:

- (PIX 7.0 and FWSM) Click **Add Context**.
- (Cisco ASA) Click **Add Module**.

Device Type: Cisco PIX 7.0 ▼

→	*Device Name:	firewall1
→	*Context Name:	context1
→	*Reporting IP:	10 1 1 23
	SNMP RO Community:	public

Discover
Cancel
Submit

143179

Step 2 In the Device Type list, do one of the following:

- For Cisco ASA, select **Cisco ASA 7.0**.
- For PIX 7.0, select **Cisco PIX 7.0**.
- For FWSM, select **Cisco FWSM x.y**, where x.y is the version number of the software running on the module.

Step 3 Enter the name of the firewall device in the Device Name field.

MARS maps this name to the reporting IP address. This name is used in topology maps, queries, and in the Security and Monitoring Device list. For devices that support the discovery operation, such as routers and firewalls, MARS renames this field's value to match the name discovered in the device configuration, which typically uses the *hostname.domain* format. For devices that cannot be discovered, such as Windows and Linux hosts and host applications, MARS uses the provided value.

Step 4 Enter the name of the security context in the Context Name field.

This name must exactly match the context name defined on the device.

Step 5 Enter the IP address of the security context from which syslog messages or SNMP notifications, or both are published in the Reporting IP field.

To learn more about the reporting IP address, its role, and dependencies, see [Understanding Access IP, Reporting IP, and Interface Settings](#), page 2-8.

Step 6 (Optional) To enable MARS to retrieve MIB objects for this security context, enter the device's read-only community string in the SNMP RO Community field.

Before you can specify the SNMP RO string, you must define an access IP address. MARS uses the SNMP RO string to read MIBs related to a security context's CPU usage, network usage, and device anomaly data and to discover device and network settings.

Step 7 To discover the settings of the defined context click **Discover**.

This discovery collects all of the route, NAT, and ACL-related information. In addition, the name of the device may change to the *hostname.domain* format if it was not already entered as such.

Step 8 To save your changes, click **Submit**.

Add Discovered Contexts

When you select Discover on a Cisco ASA, PIX 7.0 or FWSM, MARS discovers the contexts that are defined for that firewall device. However, you must still manually add discovered contents.



Note

You cannot discover a module install in a Cisco ASA; you must manually define IPS modules. However, the discovered contexts do appear under the Module area on the main page.

Step 1 Do one of the following:

- (PIX 7.0 and FWSM) Click **Add Available Context**.
- (Cisco ASA) Click **Add Available Module**.

Module Name	Module Type
☐ asa context	Cisco ASA 7.0
☐ ips context	Cisco IPS 5.x

Step 2 Select a security context from the Select list.

Step 3 Click **Add**.

Step 4 Repeat for other contexts.

Step 5 To save your changes, click **Submit**.

After you add discovered contexts, you must edit them to provide the contact information required by MARS. Continue with [Edit Discovered Security Contexts](#), page 4-11.

Edit Discovered Security Contexts


Note

You must edit all discovered contexts to specify the reporting IP address and the SNMP RO community string.

Step 1

From the list of discovered contexts, select the one that you want to edit and select the action appropriate to the device type:

- (PIX 7.0) Click **Edit Context**.
- (Cisco ASA and FWSM) Click **Edit Module**.

Device Type: Cisco ASA 7.0

→ *Device Name:	qa.protegonetworks.co
→ *Context Name:	qa
→ *Reporting IP:	10 4 2 9
SNMP RO Community:	public

143211

Step 2

Enter the IP address from which the syslog messages of the security context are sent in the Reporting IP field.

Step 3

(Optional) To enable MARS to retrieve MIB objects for this context, enter the device's read-only community string in the SNMP RO Community field.

Before you can specify the SNMP RO string, you must define an access IP address. MARS uses the SNMP RO string to read MIBs related to a reporting device's CPU usage, network usage, and device anomaly data and to discover device and network settings.

Step 4

(Optional) To enable MARS to monitor this context for anomalous resource usage, select **Yes** from the Monitor Resource Usage list.

Result: MARS monitors the device for anomalous consumption of resources, such as memory and CPU. If anomalies are detected, MARS generates an incident. Resource utilization statistics are also used to generate reports. For more information, see [Configuring Resource Usage Data, page 2-41](#).

Step 5

To save your changes, click **Submit**.

Step 6

Repeat for each discovered context.

NetScreen ScreenOS Devices

MARS can monitor NetScreen ScreenOS devices, versions 4.0 and 5.0. To enable this monitoring, you must:

1. Provide MARS with SNMP, SSH or Telnet administrative access to NetScreen device.

2. Define the SNMP RO community strings to shared between the NetScreen device and MARS.
3. Specify which syslog messages to published to MARS.
4. Add the Netscreen Device to the MARS web interface.


To accomplish these requirements, you must perform two procedures:

- [Bootstrap the NetScreen Device, page 4-12](#)
- [Add the NetScreen Device to MARS, page 4-17](#)

Bootstrap the NetScreen Device

To prepare the NetScreen device to be monitored by MARS, follow these steps:

- Step 1** Login to the NetScreen with appropriate username and password.
- Step 2** In the main screen, on the left hand column click **Network > Interfaces**.



- Step 3** Click **Edit** next to the appropriate interface to configure for MARS to have access to SNMP and Telnet/SSH.

Network > Interfaces > Edit NS5GT-DI

Interface: trust (IP/Netmask: 192.168.5.200/24) [Back To Interface](#)

Properties: **Basic** MIP DIP Secondary IP IGMP

Interface Name: trust (mac 0010.db5b.8dd2)
Zone Name: Trust

Obtain IP using DHCP ☐ Automatic update DHCP server parameters
Obtain IP using PPPoE: None [Create new pppoe setting](#)
Static IP ☒
IP Address / Netmask: 192.168.5.200 / 24 ☒ Manageable
Manage IP *: 192.168.5.200 (mac 0010.db5b.8dd2)

Interface Mode: ☒ NAT ☐ Route
Block Intra-Subnet Traffic: ☐

Service Options

Management Services	☒ Web UI	☒ Telnet	☒ SSH
	☒ SNMP	☒ SSL	
Other Services	☒ Ping	☐ Ident-reset	

Maximum Transfer Unit (MTU): 1500 Bytes

DNS Proxy: ☐
WebAuth: ☐ IP: 0.0.0.0 ☐ SSL Only

Traffic Bandwidth: 0 kbps

OK Apply Cancel

Step 4 Under Service Options, select one of the following values:

- SNMP
- Telnet
- SCS (4.0 only)
- SSH (5.0 and later)

MARS can only use one of the access methods to perform configuration discovery. This value will also be selected in the Access Type value of [Add the NetScreen Device to MARS](#), page 4-17.

Step 5 Click **Apply** then click **OK**.

Step 6 Configure the SNMP information by selecting **Configure > Report Settings > SNMP**.

143195

Configuration > Report Settings > SNMP NS5GT-DI

Juniper®
TREND MICRO SECURE
NetScreen-5GT

Home
Configuration
Date/Time
Update
Admin
Auth
Report Settings
Log Settings
Email
SNMP
Syslog
WebTrends
Network
Screening
Policies
MCast Policies
VPNs
Objects

SNMP Report Settings

System Name NS5GT-DI

System Contact

Location

Listen Port 161

Trap Port 162

Enable Authentication Fail Trap ☒

Apply Cancel

Communities:

Name	Write	Trap	Traffic	Hosts	Configure
No entry available					

143200

- Step 7** Add the MARS IP address in the **Host List** by clicking **Edit**.
- Step 8** Enter the MARS IP address and verify that this Community Name in this window is the same community string entered in the MARS web interface when adding this device.
- Step 9** (Optional) If the community string does not match, click **New Community** to define one that matches the one defined in MARS.

Configuration > Report Settings > SNMP > Community Edit NS5GT-DI?

Juniper®
TREND MICRO SECURE
NetScreen-5GT

Home
Configuration
Date/Time
Update
Admin
Auth
Report Settings
Log Settings
Email
SNMP
Syslog
WebTrends
Network
Screening
Policies
MCast Policies
VPNs
Objects
Reports
Wizards
Help
Logout

Community Name


Permissions
☒ Write
☒ Trap
☒ Including Traffic Alarms

Version V1


Hosts IP Address	Netmask	Trap Version	Source Interface
0.0.0.0		V1	Not specified
0.0.0.0		V1	Not specified
0.0.0.0		V1	Not specified
0.0.0.0		V1	Not specified
0.0.0.0		V1	Not specified
0.0.0.0		V1	Not specified
0.0.0.0		V1	Not specified
0.0.0.0		V1	Not specified

OK Cancel

- Step 10** Configure the Syslog information by selecting **Configure > Report Settings > Syslog**.
- Step 11** Verify that the **Enable Syslog Messages** and **Include Traffic Log** boxes are checked.
- Step 12** Enter the IP address of the MARS Appliance that will listen for events from this device
- Step 13** Verify that the default syslog port number of 514 is selected.
- Step 14** Select the **AUTH/SEC** for **Security Facility** and **LOCAL0** for **Facility**.
- Step 15** For NetScreen 5.0, select the **Event Log** in addition to **Traffic Log**.
- Step 16** Click **Apply**.



Step 17 Configure logging for each policy that user wants to send the events to the MARS Appliance. Select **Policies** on the left hand area.



Step 18 Click **Edit** then **Advance** and verify that **Logging** box is checked. Repeat for all policies which events need to be sent to MARS.

The screenshot displays the NetScreen ScreenOS configuration interface for creating or editing a policy. The left-hand navigation pane includes links to Home, Configuration, Network, Screening, Policies, MCast Policies, VPNs, Objects, Reports, Wizards, Help, and Logout. The main configuration area is titled 'Policies (From Untrust To Trust)'. It features several input fields and checkboxes: 'Name (optional)' is a text box; 'Source Address' and 'Destination Address' each have radio buttons for 'New Address' and 'Address Book Entry' (currently set to 'Any'), along with a 'Multiple' button; 'Service' is a dropdown menu set to 'HTTP' with a 'Multiple' button; 'Application' is a dropdown menu set to 'None'; 'URL Filtering' is an unchecked checkbox; 'Action' is a dropdown menu set to 'Permit' with a 'Deep Inspection' button; 'Antivirus Objects' includes a list of 'Attached AV Object Names' and 'Available AV Object Names' (containing 'scan-mgr') with left and right arrow buttons; 'Tunnel' is a dropdown menu set to 'VPN' with a 'None' option; 'Modify matching bidirectional VPN policy' is an unchecked checkbox; 'L2TP' is a dropdown menu set to 'None'; and 'Logging' is a checked checkbox. At the bottom are 'OK', 'Cancel', and 'Advanced' buttons. A vertical text '14 3198' is visible on the right edge of the interface.

- Step 19** Verify that all the Syslog event severity levels that need to be sent to MARS are configured. Verify which Syslog severity levels that are enabled by selecting **Configuration > Report Settings > Log Settings**.

Add the NetScreen Device to MARS

- Step 1** Select **Admin > System Setup > Security and Monitor Devices > Add**.
- Step 2** Select the appropriate version of NetScreen ScreenOS from the Device Type list.
- Step 3** Enter the name of the device in the Device Name field.

MARS maps this name to the reporting IP address. This name is used in topology maps, queries, and in the Security and Monitoring Device list. For devices that support the discovery operation, such as routers and firewalls, MARS renames this field's value to match the name discovered in the device configuration, which typically uses the *hostname.domain* format. For devices that cannot be discovered, such as Windows and Linux hosts and host applications, MARS uses the provided value.

- Step 4** (Optional) To enable MARS to discover settings from this device, enter the administrative IP address in the Access IP field.

To learn more about the access IP address, its role, and dependencies, see [Understanding Access IP, Reporting IP, and Interface Settings](#), page 2-8.

- Step 5** Enter the IP address of the interface that publishes syslog messages or SNMP notifications, or both in the Reporting IP field.
- To learn more about the reporting IP address, its role, and dependencies, see [Understanding Access IP, Reporting IP, and Interface Settings, page 2-8](#).
- Step 6** If you entered an address in the Access IP field, select **SNMP**, **TELNET**, or **SSH**, from the Access Type list, and continue with the procedure that matches your selection:
- [Configure SNMP Access for Devices in MARS, page 2-11](#)
 - [Configure Telnet Access for Devices in MARS, page 2-11](#)
 - [Configure SSH Access for Devices in MARS, page 2-12](#)
- For more information on determining the access type, see [Selecting the Access Type, page 2-10](#).
- Step 7** (Optional) To enable MARS to retrieve MIB objects for this reporting device, enter the device's read-only community string in the SNMP RO Community field.
- Before you can specify the SNMP RO string, you must define an access IP address. MARS uses the SNMP RO string to read MIBs related to a reporting device's CPU usage, network usage, and device anomaly data and to discover device and network settings.
- Step 8** (Optional) If you defined an access IP and selected and configured an access type, click **Discover** to determine the device settings.
- Result:* If the username and password are correct and the MARS Appliance is configured as an administrative host for the device, the "Discovery is done." dialog box appears when the discovery operation completes. Otherwise, an error message appears. After the initial pull, the MARS Appliance pulls based on the schedule that you define. For more information, see [Scheduling Topology Updates, page 2-39](#).
- Step 9** To add this device to the MARS database, click **Submit**.
- Result:* The submit operation records the changes in the database tables. However, it does not load the changes into working memory of the MARS Appliance. The activate operation loads submitted changes into working memory.
- Step 10** Click **Activate**.
- Result:* MARS begins to sessionize events generated by this device and evaluate those events using the defined inspection and drop rules. Any events published by the device to MARS before activation can be queried using the reporting IP address of the device as a match criterion. For more information on the activate action, see [Activate the Reporting and Mitigation Devices, page 2-27](#).
-

Check Point Devices

The Check Point security product family can be distributed and tiered. As such, you must understand the deployment method, components, and release versions of this product family, their relationships, and how MARS interacts with them. You must also understand the many acronyms and abbreviations associated with this product family. [Table 4-1](#) lists the abbreviations and acronyms used in the topics that follow.

Table 4-1 *Check Point Abbreviations and Acronyms*

Abbreviation	Expansion	Additional Information
ASYMSSLCA	Secure Sockets Layer Certificate Authority using an asymmetric key cipher	Communications protocol used for establishing secure sessions.
CLM	Customer Log Modules	Standalone log server for collecting log data from the Check Point enforcement modules.
CMA	Customer Management Add-ons	A a virtual instance of SmartCenter and only exists within the context of a Provider-1/SiteManager-1 infrastructure.
CPMI	Check Point Management Interface	Communications protocol used for configuration discovery.
LEA	Log Export API	Communications protocol used for retrieving audit and firewall logs.
MDG	Multi Domain GUI	GUI used for managing Provider-1/SiteManager-1 deployments. The MDG is the parent GUI that can launch specific SmartDashboard GUIs for a CMA.
MDS	Multi Domain Server	Is the umbrella manager for the CMA instances in a Provider-1/SiteManager-1 deployment.
MLM	Multi Domain Log Module	Usually found in Provider-1/SiteManager-1 deployments and provides the ability to create multiple instances of a CLM on a single logging server.
NG AI	Next Generation with Application Intelligence	All current trains of Check Point are released under the NG AI umbrella with specific release numbers, such as NG AI R55 and NG AI R60.
NG FP3	Next Generation Feature Pack 3	—
NGX	Next Generation eXtension	NGX is also NG AI R60
OPSEC	Open Platform for Security	An alliance, certification and integration methodology for products and solutions that integrate into a Check Point infrastructure.
P-1	Check Point Provider-1	—

Table 4-1 Check Point Abbreviations and Acronyms

Abbreviation	Expansion	Additional Information
SSLCA	Secure Sockets Layer Certificate Authority, using a symmetric key cipher (protocol)	—
SIC	Secure Internal Communication	—
SIC DN	SIC Distinguished Name	—
VIPs	Virtual IP Addresses	Usually used in a Provider-1/ SiteManager-1 deployment to assign unique IP addresses for CMA instances.
VPN-1	Check Point VPN-1 Pro and Edge	VPN-1 Pro is the Check Point enforcement gateway that does the inspection, firewalling, VPN encryption and QoS tagging. VPN-1 Edge is treated as a normal enforcement point.

To understand what MARS supports, we must first clarify the product terminology used by Check Point. NG refers to the 5.x product family, and it included three feature packs: FP1, FP2, and FP3. NG is different from NG AI in that NG AI improved upon, and renamed, the SmartDefense feature set that was introduced in NG FP2. NG AI also provides a larger number of application-aware inspections,; hence the name Application Intelligence. NG AI included releases R54 and R55. NGX refers to the 6.x product family and began with the R60 release.

MARS supports and has been tested with the following releases:

- NG FP3
- NG AI (R55)
- NGX (R60)

The different security platforms, Provider-1, SiteManager-1, SmartCenter, and SmartCenter Pro are bundles of the technologies released under the NG, NG AI, and NGX release trains. From this perspective, MARS works with any of the security platforms as long as it belongs to one of the supported release trains.

Check Point Provider-1 is a security management system for the managed security service providers (MSSP) and multi-site enterprises, respectively. Service providers are able to manage the Check Point gateways (firewall and VPN gateways) on their customer sites. The security policies and the system configurations are stored on the MDS. Each per-customer security policy is managed through a CMA, which also reside on the MDS. The Provider-1 system allows the service provider and the end customers to maintain separate log servers, using the MLM and CLM respectively. The user interface for Provider-1 is called the MDG. This system also support a tiered fault-tolerant configuration via redundancy at the gateway, CMA, or MDS level.

The Provider-1 system ensures secure and private communication between its components and Check Point gateways. Each CMA has its own internal certificate authority that issues certificates for secure communication between the CMA, log servers, and its own network. All communication between MDSs is authenticated and secured, and every MDS communicates securely with the CMAs that it houses.

The SiteManager-1 system operates much the same as Provider-1; however, it is targeted toward large enterprise customers. The Check Point components are the same as those found in Provider-1.

SmartCenter and SmartCenter Pro are security management systems also targeted toward enterprise customers. They can support the Provider-1 system, serving as a backup server at the CMA level. However, their primary function is to provide centralized security and VPN policy and security event management through SmartDashboard, which is the user interface for both systems. From the MARS perspective, SmartCenter has the ability to extend the view of the network by managing the policies and events associated with gateway and desktop nodes:

- VPN-1 perimeter security gateways,
- InterSpect internal security gateways
- Connectra Web security gateways
- SecureClient, a personal firewall running on desktops and servers.

MARS monitors the primary management servers, such as the MDS in Provider-1 and SiteManager-1 and the SmartCenter Server in SmartCenter and SmartCenter Pro. These management servers are where the actual security and audit policies are centrally managed and stored. If the Check Point deployment requires, MARS also monitors those components managed by the management stations, such as individual firewalls, VPN gateways, and log servers. Whether you configure MARS to monitor these remote components depends on whether their security event logs are propagated to the centralized management servers (SmartCenter or CMA). If the logs are not forwarded to the primary management server, then you must define where the log repository exists, whether local to the enforcement module, or forwarded to a separate logging module (CLM).

In addition to understanding the components, it is important to understand how Check Point components use Secure Internal Communications (SIC) to securely communicate with each other and with third-party OPSEC applications. SIC is the process by which MARS Appliance authenticates to the SmartCenter Server and other Check Point components. SIC uses a shared secret as the seed for negotiating session keys. This shared secret is referred to as an activation key. The authentication and communication setup works as follows:

1. Using a username and password pair, MARS authenticates to the SmartCenter Server and other Check Point components, such as remote log servers, using TCP port 18210.
2. If authenticated, the peers swap the activation key and each other's SIC using TCP port 18190.
3. If each peer validates the authenticity of the other, the Check Point component establish an encrypted session over TCP port 18184 with the MARS Appliance. It is over this channel that the Check Point components send encrypted log data to MARS.

The following topics support the integration of MARS into a Check Point environment:

- [Determine Devices to Monitor and Restrictions, page 4-21](#)
- [Bootstrap the Check Point Devices, page 4-22](#)
- [Add and Configure Check Point Devices in MARS, page 4-36](#)
- [Troubleshooting MARS and Check Point, page 4-53](#)

Determine Devices to Monitor and Restrictions

To configure Check Point devices, you must identify the central management server and managed components, bootstrap them, and add and configure them in the MARS web interface. The Check Point product line and release, as well as the number of devices managed, determines which tasks you must perform to configure MARS to monitor your Check Point devices.

Representing a Check Point device in MARS involve two steps:

1. **Define a primary management station.** This primary management station represents the central management server that manages remote components, such as firewalls, VPN gateways, and log servers.
2. **Define one or more child enforcement modules.** Child enforcement modules are the remote components managed by the primary management station. They represent firewalls, VPN gateways, and log servers.

When managing SmartCenter and SmartCenter Pro, the primary management station is the SmartCenter server. When managing Provider-1/SiteManager-1 releases NG FP3, NG AI (R55), and NGX (R60), the primary management station is not the MDS, but each CMA defined under the MDS. In other words, you must define each CMA as a separate primary management station. The child enforcement modules are those gateways and logs servers (CLMs) managed as part of that customer or site as defined by the CMA.

Part of what you must determine is where the security event logs are stored. Two options exist:

- **Central Event Correlation.** The MLM or SmartCenter server pulls logs from all remote components.
- **Distributed Event Correlation.** In addition to the MLM or SmartCenter Server, one or more remote log servers exist where aggregation to the central management server does not occur. These servers, the CLMs, must also be represented and configured so that MARS can pull the events from them.

If the security events are stored in a distributed fashion, you must plan to define and establish SIC communication between the MARS Appliance and each Check Point log module. For SmartCenter and SmartCenter Pro, the server SIC DN is the one assigned to the primary management station. However, for Provider-1 and SiteManager-1, the server SIC DN varies based on release. For Provider-1 and SiteManager-1 NG FP3 and NG AI (R55), the server SIC DN is the one associated with the CMA. For Provider-1 and SiteManager-1 NGX (R60), you can use the SIC assigned to the MDS for all CMAs and CLMs that you define.

One other restriction exists with the Provider-1 and SiteManager-1 products. For Provider-1 and SiteManager-1 NG FP3 and NG AI (R55), you must define an OPSEC application representing the MARS Appliance in each CMA (using the CMAs SmartDashboard user interface). For Provider-1 and SiteManager-1 NGX (R60), you can define one OPSEC application representing the MARS Appliance and push that definition to all CMAs and CLMs managed by the MDS.

Bootstrap the Check Point Devices

Bootstrapping the Check Point devices involves preparing those devices to send data to the MARS Appliance, as well as enabling the MARS Appliance to discover the Check Point configuration settings. In addition to preparing the Check Point devices, you must gather the information required to represent the Check Point devices in the MARS web interface.

You bootstrap the central Check Point management server, whether it be a CMA or a SmartCenter server by defining the MARS Appliance as a target log host and OPSEC Application object.

1. Using Check Point SmartDashboard or the Check Point Provider-1/SiteManager-1 MDG, add the MARS Appliance as a host.
2. Create and install an OPSEC Application object for the defined host, import the authorization key, and generate the client SIC DN. This SIC DN is the one used by OPSEC applications, including the management server, to validate the MARS Appliance. You specify this client SIC DN in the MARS web interface. When a session is established between the MARS Appliance and the Check Point management server, the appliance publishes this SIC to the management server to ensure non-repudiation of the appliance.

3. Obtain the server SIC DN of the Check Point management server. You specify this sever SIC in the MARS web interface. The MARS Appliance validates the server SIC DN against the SIC published to the appliance by the management server during session setup. This validation ensures non-repudiation of the server.
4. Create the policies to permit SIC traffic between the defined host (MARS Appliance), the Check Point management server, and any remote servers. After you identify the devices, you must verify that the network services they use for SIC-based management and reporting are permitted on the reporting device. To enable these traffic flows, you must verify or update the policies that enable the SIC traffic to flow between each reporting device and the MARS Appliance. Once you have updated these policies, you must install the policies.
5. Define the log settings to push the correct events to the defined host. You must ensure that all of the security, firewall, user authentication, and audit events are logged and configured to be published to the MARS Appliance.
6. Install the policies. Once the policies are defined, you must update the Check Point components with the policies. Policy installation include an object database push that make the Check Point modules aware of the OPSEC Application representing the MARS Appliance. Without this step, the modules will not forward any log information via LEA.

To perform this task, you need a Check Point user account with administrative privileges. This account must be able to create a new host, define OPSEC application, define and install new policies, and access the settings of each managed Check Point component.

After completing this task, you should have collected the following information:

- The Client and server SIC DNs.
- If you are defining a CMA for Provider-1 or SiteManager-1 NG FP3 or NG AI (R55), then you must have the virtual IP address (VIP) for each CMA and CLM managed by the MDS. Only Provider-1 and SiteManager-1 NGX (R60) requires the physical IP addresses of the MDS and MLM servers.
- Any CLMs, instead of CMAs, to which security logs are being sent. If logs are being sent to CLMs, LEA is only supported using clear text.

To bootstrap the Check Point devices, perform the following procedures:

- [Add the MARS Appliance as a Host in Check Point, page 4-23](#)
- [Define an OPSEC Application that Represents MARS, page 4-24](#)
- [Obtain the Server Entity SIC Name, page 4-27](#)
- [Select the Access Type for LEA and CPMI Traffic, page 4-29](#)
- [Create and Install Policies, page 4-31](#)
- [Verify Communication Path Between MARS Appliance and Check Point Devices, page 4-32](#)

Add the MARS Appliance as a Host in Check Point

Representing the MARS Appliance in Check Point enables the following supporting tasks:

- Generate a client SIC DN for the MARS Appliance.
- Define policies to allow SIC and syslog traffic between the Check Point components and the MARS Appliance.
- Direct log traffic to the MARS Appliance.

To define the MARS Appliance as a host, follow these steps:


Step 1 Log in to the correct Check Point user interface using an account with administrative privileges.
If you are using SmartCenter, use the SmartDashboard for that server. If you are using Provider-1 or SiteManager-1 NG FP3 or NG AI (R55), use the SmartDashboard of the CMA. If you are using Provider-1 or SiteManager-1 NGX, use the MDG.

Step 2 Select **Manage > Network Objects** from the main menu.

Result: The Network Objects dialog box appears.

Step 3 Click the **New** button, and then select **Node > Host** on the menu list.

Result: The Host Node dialog appears, with the General Properties settings selected.



Step 4 Enter the name MARS Appliance in the Name field of the General Properties page

Any Check Point policies defined to enable access or send logs to this appliance will reference the appliance by this name. Cisco best practice recommends using the actual hostname of the MARS Appliance.

Step 5 Enter the IP address of the monitoring interface in the MARS Appliance in the IP Address field

Typically, the monitoring interface is eth0. However, if one or more intermediate gateways are applying NAT rules to the physical IP address, enter the IP address that is exposed to the Check Point central management server.

Step 6 Click **OK** to close the Host Node dialog box.

Step 7 Click **Close** to close the Network Objects dialog box.

Result: The host representing the MARS Appliance is defined. You can now use this host when defining new policies in the Check Point user interface.


Step 8 Continue with [Define an OPSEC Application that Represents MARS, page 4-24](#).

Define an OPSEC Application that Represents MARS

To integrate a third-party OPSEC application with Check Point components, you must define the application and associate it with the host on which the application is running. In addition to identifying this OPSEC application to the Check Point system, this procedure results in the generation of the client SIC DN for the MARS Appliance. Both the client SIC DN and the server SIC DN, obtained in [Obtain the Server Entity SIC Name, page 4-27](#), are required to enable secure communications between the appliance and Check Point components.

This procedure also involves selecting an activation key, or shared secret, that is also required to enable the secure communications. You must record both the activation key and the client SIC DN for use when defining the Check Point devices in the MARS web interface.


- Step 1** Log in to the correct Check Point user interface using an account with administrative privileges.
If you are using SmartCenter, use the SmartDashboard for that server. If you are using Provider-1 or SiteManager-1 NG FP3 or NG AI (R55), use the SmartDashboard of the CMA. If you are using Provider-1 or SiteManager-1 NGX, use the MDG.
- Step 2** Select **Manage > Servers and OPSEC Applications** from the main menu.
Result: The Servers and OPSEC Application dialog box appears.
- Step 3** Click the **New** button, and then click **OPSEC Application** on the menu list.
Result: The OPSEC Application Properties dialog box appears.



- Step 4** Specify the name for this object in the Name field.
This value must be different from the name specified in [Step 4 of Add the MARS Appliance as a Host in Check Point, page 4-23](#). Best practice recommends using the actual hostname of the host object plus some other descriptor, which combines for a unique name.
- Step 5** In the Host list, select the host that you specified in [Step 4 of Add the MARS Appliance as a Host in Check Point, page 4-23](#).
- Step 6** *Result:* This OPSEC application definition is associated with the host that represents the MARS Appliance.
- Step 7** Verify that **User defined** is selected in the Vendor field.
- Step 8** Select the **LEA** and **CPMI** check boxes under Client Entities.
These values identify the OPSEC services required by the MARS Appliance.

- Step 9** Click the **Communication** button under Secure Internal Communication.

Result: The Communication dialog box appears.



- Step 10** Enter the activation key in the Activation Key and Confirm Activation Key fields of the Communication dialog box.




Note

Remember this key for future use with MARS.

- Step 11** Click **Initialize** to generate the client SIC DN.

Result: The client SIC DN is generated and the Communication dialog box closes, returning to the OPSEC Application Properties dialog box. The new SIC appears in the DN field.



- Step 12** Click **Close** to close the Communication dialog box.

- Step 13** Record the contents of the DN field that appears under Secure Internal Communication.

This value is used to populate the Client Entity SIC Name field of MARS in [Add a Check Point Primary Management Station to MARS](#), page 4-37.



Tip

If possible, you should cut and paste the Secure Internal Communication DN field value into an application, such as Notepad, for later use. Transcribing this field is error prone. Use a mouse to select the contents of read-only field, and then use Ctrl+Insert to copy the field to memory. You can paste the value using Shift+Insert. Be careful to avoid trailing spaces when you paste the value into MARS.

- Step 14** Select the **CPMI Permissions** tab and verify that either **Administrator's credentials** or a permissions profile with administrative credentials is selected under Login to SmartCenter with.

Step 15 Click **OK** to close the OPSEC Application Properties dialog box.

Step 16 Click **Close** to close the Servers and OPSEC Application dialog box.

Result: The OPSEC Application that represents MARS is defined and associated to the correct host. You also have obtained the activation key and client SIC DN for later use in [Add a Check Point Primary Management Station to MARS, page 4-37](#).

Step 17 Select **Policy > Install Database** on the main menu.

Result: This operation updates the remote Check Point components (child enforcement modules), such as CMAs, CLMs, log servers, and firewalls. It provides them with the authorization and credentials of the MARS Appliance, as an OPSEC component and SIC client.



Tip

Using the Check Point log viewer, you can verify that the OPSEC object was pushed successfully.

Step 18 Continue with [Obtain the Server Entity SIC Name, page 4-27](#).

Obtain the Server Entity SIC Name

The server SIC DN is one of the shared secrets used to provide non-repudiation during a secure communication between a Check Point component and the MARS Appliance. This value is used when defining a primary management station in MARS as defined in [Add a Check Point Primary Management Station to MARS, page 4-37](#).

Step 1 Log in to the correct Check Point user interface using an account with administrative privileges.

If you are using SmartCenter, use the SmartDashboard for that server. If you are using Provider-1 or SiteManager-1 NG FP3 or NG AI (R55), use the SmartDashboard of the CMA. If you are using Provider-1 or SiteManager-1 NGX (R60), use the MDG.

Step 2 Select **Manage > Network Objects** on the main menu.

Step 3 Select **Check Points** in the Show list.

Step 4 Select the correct Check Point component in the Network objects list.

Which Check Point component you select depends on which SIC you need and what Check Point system you are using. Specifically, you want to obtain SICs for:

- Each management server to discover configuration settings via CPML.
- Each management server to which logs are forwarded by remote components.
- Each remote log server that does not forward logs to a central management server, either the MDS or a SmartCenter.


Management servers are the following devices:

- SmartCenter server for standalone SmartCenter and SmartCenter Pro installations.
- Each CMA of a Provider-1 or SiteManager-1 NG FP3 or NG AI (R55) installation.
- The MDS of a Provider-1 or SiteManager-1 NGX (R60) installation.

Log servers are the following devices:


- SmartCenter server for standalone SmartCenter and SmartCenter Pro installations.
- Each CLM of a Provider-1 or SiteManager-1 NG FP3 or NG AI (R55) installation.

- The MLM of a Provider-1 or SiteManager-1 NGX (R60) installation.



Step 5 Click **Edit**.

The Check Point Host - Management dialog box appears, with the General Properties page selected.



Step 6 Record the value defined in the DN field under Secure Internal Communication.

This value is used to populate the Server Entity SIC Name field of MARS in either [Add a Check Point Primary Management Station to MARS](#), page 4-37, [Manually Add a Child Enforcement Module or Log Server to a Check Point Primary Management Station](#), page 4-41, or [Edit Discovered Firewall on a Check Point Primary Management Station](#), page 4-47.

Step 7 Click **OK** to close the Check Point Host dialog box.

- Step 8** For each additional management or log server in this Check Point installation, select that device in the Network Objects list, and repeat [Step 5](#) through [Step 7](#).
- Step 9** Click **Close** to close the Network Objects dialog box.
- Step 10** Continue with [Select the Access Type for LEA and CPMI Traffic, page 4-29](#).

Select the Access Type for LEA and CPMI Traffic

Check Point devices use special access types for configuration discovery and event log queries. For configuration discovery, the protocol is CPMI. For event log queries, the protocol is LEA. Each of these protocols has specific configurable attributes, including whether to use bulk encryption, what cipher to use, and what port to use for communications.

You must understand what the supported settings are so that you can verify the Check Point devices are configured correctly. MARS supports only three of the available Check Point authentication mode:

- **CLEAR.** Indicates that the traffic is neither authenticated nor encrypted.
- **SSLCA.** Indicates that the communications need to be authenticated and encrypted using an symmetric key cipher
- **ASYMSSLCA.** Indicates that the communications need to be authenticated and encrypted using an asymmetric key cipher.

These access protocols are configured as follows:



Note

Typically, the default values should be used unless your Check Point deployment includes CLMs.

- `<service> auth_port <port_number>`

This line is required in the `fwopsec.conf` file. The *service* value is either **LEA_SERVER** or **CPMI_SERVER**. Two possible values exist for *port_number*: **0**, which indicates an the server is not listening for authenticated session requests, and the port number of an authenticated and/or encrypted protocol. If the *port_number* value is 0, you must configure the server to listen for session requests in CLEAR mode on a valid port using the `<service> port <port_number>` settings.

- `<service> auth_type <cipher>`

The *service* value is either **LEA_SERVER** or **CPMI_SERVER**. Two possible values are supported for *cipher*: **sslca** for authentication and encryption using a symmetric key cipher, or **asym_sslca** for authentication and encryption using an asymmetric key cipher. If the **auth_port** setting is set to 0 (zero) for this service, then you do not need to specify the *auth_type* in the `fwopsec.conf` file. You can comment out this line.

- `<service> port <port_number>`

This line is required in the `fwopsec.conf` file. The *service* value is either **LEA_SERVER** or **CPMI_SERVER**. The value for *port_number* must match the port number on which the desired network service listens. A *port_number* of **0** (zero) indicates that log server is not listening in CLEAR mode.

If it is some other number, then any service can come pull the logs without authenticating. For **LEA_SERVER**, you cannot use port 18184, as it is used for encrypted log communications. For **CPMI_SERVER**, you cannot use port 18190. When CLEAR is enabled, authentication is disabled

for this port. Any host with access to the Check Point component at this port can pull logs. If you chose to enable CLEAR, which is less expensive in terms of overall transaction costs, you define policies that restrict access to the MARS Appliance and other known management hosts.

**Note**

Prior to MARS 4.1 and when using Provider-1 or SiteManager-1 NG FP3 or NG AI (R55), you could not use SSLCA mode for log retrieval by the MARS Appliance. Instead, you were required to configure each CMA and CLM to accept LEA session requests using CLEAR mode. It was unnecessary to configure the LEA settings for the MLM.

The following example indicates that LEA is using ASYMSSLCA-based authentication connecting over port 18184 (default), the traffic is encrypted via SSL, and the log server is not listening for requests in cleartext.

```
LEA_SERVER auth_port 18184
LEA_SERVER auth_type asym_sslca
LEA_SERVER port 0
```

The following example indicates that the log server is listening for requests in cleartext at port 18187. Such requests will be serviced and the sessions will be neither authenticated nor encrypted.

```
LEA_SERVER port 18187
```

Check Point uses the following default settings:

- For LEA, SSLCA is the authentication method and communications occur over TCP 18184.
- For CPML, SSLCA is the authentication method and communications occur over TCP 18190.

To review or change the access type settings, follow these steps:

Step 1 Log on to the Check Point server.

For Provider-1 and SiteManager-1, this server is the MDS, MLM, or CLM. Otherwise, it is the SmartCenter server.

Step 2 Open the `fwopsec.conf` file found in the subdirectory for each CMA and CLM.

The following example uses the `find` command to locate the file. Customer1 identifies the CLM.

```
[Expert@logger]# find . -name "fwopsec.conf" -print
./var/opt/CPfw1-R55/conf/fwopsec.conf
./var/opt/CPmds-R55/customers/Cust1Log/CPfw1-R55/conf/fwopsec.conf
[Expert@logger]# cd /var/opt/CPmds-R55/customers/Cust1Log/CPfw1-R55/conf
```

Step 3 Using a text editor, such as `vi` or Notepad, edit the `fwopsec.conf` file and modify the LEA and CPML communication settings as needed.

Step 4 Save your changes to the file.

Step 5 Repeat [Step 2](#) through [Step 4](#) for each CLM and CMA.

Step 6 Restart the Check Point server after the changes are made.

Result: The CPML and LEA servers are restarted, which reloads their configuration information, and ensures they are listening to the correct ports for session requests.

Step 7 Continue with [Create and Install Policies](#), page 4-31.

Create and Install Policies


You must create firewall policies that permit the MARS Appliance to access the relevant ports of the Check Point central management server and any remote log servers. The default ports are as follows:

- **TCP port 18190.** Used by CPMI to discover configuration settings.
- **TCP port 18210.** Used to retrieve the certificate from the Certificate Authority on the SmartCenter, MDS, MLM, CMA, or CLM.
- **TCP port 18184.** Used to pull security event logs from the log servers, such as the MLM or CLM.

However, you must use the CPMI and LEA servers settings specified in [Select the Access Type for LEA and CPMI Traffic, page 4-29](#). When the policies are defined, you must install them on any firewall modules that inspect traffic between the Check Point components and the MARS Appliance.

If the management server has a Check Point firewall installed, follow these steps:

-
Step 1	Log in to the correct Check Point user interface using an account with administrative privileges. If you are using SmartCenter, use the SmartDashboard for that server. If you are using Provider-1 or SiteManager-1 NG FP3 or NG AI (R55), use the SmartDashboard of the CMA. If you are using Provider-1 or SiteManager-1 NGX, use the MDG.
Step 2	If Check Point firewall components reside between the Check Point components (central management and log server) and the MARS Appliance monitoring those components, define the security policies that allow management and log traffic between those devices. If you have enabled CPMI discovery, the service condition must include CMPI. To enable the log access, the service list must include FW1_lea.
Step 3	Verify that the security policies are set to log. The Track column of each rule should display the Log action. To enable logging, right-click the Track field of a rule and select Log on the shortcut menu.
Step 4	Once you have defined the security policies that enable traffic flows between the Check Point and MARS components, select Policy > Install on the main menu.



- Step 5** In the Install Policy dialog box, verify the Advanced Security check box is selected for each selected installation target.

The target devices should be those firewalls that reside between the Check Point components and the MARS Appliance.

- Step 6** Click **OK** to install the policies on the selected devices.

Result: The security policies on the target firewall devices are updated, enabling CPMI and LEA traffic flows between the Check Point components and the MARS Appliance.



Tip

Using the Check Point log viewer, you can verify that the policies were installed successfully.

Verify Communication Path Between MARS Appliance and Check Point Devices

You should verify that the MARS Appliance can reach the Check Point devices, including the SmartCenter server and the remote log servers. Use the telnet command at CLI of the MARS Appliance to verify access to the SmartCenter server and log servers. The ports to check are defined in [Select the Access Type for LEA and CPMI Traffic, page 4-29](#). For more information on accessing the CLI, see [Log In to the Appliance via the Console, page 6-1](#) of the *Install and Setup Guide for Cisco Security Monitoring, Analysis, and Response System*.

The command syntax is as follows

```
telnet <ip_address> <port_number>
```


If you are unsuccessful, verify the settings of the ports for each Check Point component and verify that no firewalls are blocking the traffic. For more information on **telnet**, see [telnet](#), page A-32 in the *Install and Setup Guide for Cisco Security Monitoring, Analysis, and Response System*.

Reset the OPSEC Application Certificate of the MARS Appliance

If you encounter an error when pulling the certificate as part of defining the Check Point devices in the MARS web interface, you must reset the certificate before you can attempt to pull it again. This procedure details how to reset the certificate, or SIC, associated with the OPSEC Application that is associated with the host that represents the MARS Appliance.


To reset the OPSEC application certificate, follow these steps:

-
- Step 1** Log in to the correct Check Point user interface using an account with administrative privileges.
- If you are using SmartCenter, use the SmartDashboard for that server. If you are using Provider-1 or SiteManager-1 NG FP3 or NG AI (R55), use the SmartDashboard of the CMA. If you are using Provider-1 or SiteManager-1 NGX, use the MDG.
- Step 2** Select **Manage > Servers and OPSEC Applications** from the main menu.
- Result:* The Servers and OPSEC Application dialog box appears.
- Step 3** Select **OPSEC Applications** in the Show list.
- Step 4** Select the OPSEC application that represents the MARS Appliance in the Servers and OPSEC Applications list, and click **Edit**.
- Result:* The OPSEC Application Properties dialog box appears.



Step 5 Click the **Communication** button under Secure Internal Communication.

Result: The Communication dialog box appears.



Step 6 Click **Reset** to reset the certificate.

Step 7 Click **Close** to close the Communication dialog box.

Result: The client SIC DN is generated and the Communication dialog box closes, returning to the OPSEC Application Properties dialog box. The new SIC appears in the DN field.

Step 8 Click **OK** to close the OPSEC Application Properties dialog box.

Step 9 Click **Close** to close the Servers and OPSEC Application dialog box.

Result: The OPSEC Application that represents MARS is defined and associated to the correct host. You also have obtained the activation key and client SIC DN for later use in [Add a Check Point Primary Management Station to MARS, page 4-37](#).

Add and Configure Check Point Devices in MARS

After you identify and bootstrap the Check Point reporting devices and install the policies that enable the required traffic flows, you must represent those devices in MARS, which uses this information to communicate with the devices. When adding a Check Point device, you add two types of devices:

- **Primary management station.** The primary management station represents the SmartCenter server or CMA that manages other Check Point components. In the web interface, the bases module is defined as a software application (Check Point Management Console application) running on a host.
- **Child enforcement module.** A child enforcement module is a Check Point component, a firewall or log server, that is managed by a primary management station. When viewing the Security and Monitoring Devices list, child enforcement modules appear as children of the hosts that are running the primary management station.

With these definitions in mind, adding and configuring the Check Point device involves the following:

1. Define a host that represents the Check Point primary management station, specifying the hostname and management and reporting IP addresses.
2. Define all of the interfaces of the host.
3. Add the correct Check Point software application to the host. This application represents the primary management station.
4. Specify the communication settings for the primary management station. These settings include identifying which access types are allowed (CPMI, LEA or both) and the authentication type and port to use for each supported access type.
5. (Optional) Define the settings for secure communications. If the access communication are not conducted in CLEAR, then you must specify the client and server SIC DN's and identify the certificate authority.
6. (Optional) Define the routes used by the firewall running on the primary management station. If a firewall is running on the primary management station, the route information is required to enable the path analysis and mitigation features of MARS.
7. Discover the child enforcement modules and the configuration settings of the primary management station. Discovery of child enforcement modules includes any log servers and firewalls managed by the primary management station. MARS discovers configuration settings, such as policies, NAT, modules, and clusters, as well as event information, such as traffic logs, SmartDefense events, and user authentication events.
8. Configure the discovered log servers. To configure these log servers, select the Self option from the Log Info page associated with each server, and specify the access type settings.
9. Define any log servers not managed by the primary management station. These servers are used by one or more of the firewalls that were discovered or by the primary management station.
10. Edit each firewall child enforcement module to select a log server.
11. (Optional) Specify an SNMP RO community string for each firewall child enforcement module for which resource utilization monitoring is desired.
12. (Optional) Define the routes used by each firewall child enforcement module. Route information is required to enable the path analysis and mitigation features of MARS.
13. Click Activate in MARS.

To add a Check Point device in MARS, you must perform the following procedures:

- [Add a Check Point Primary Management Station to MARS, page 4-37](#)

- [Manually Add a Child Enforcement Module or Log Server to a Check Point Primary Management Station, page 4-41](#)
- [Edit Discovered Log Servers on a Check Point Primary Management Station, page 4-45](#)
- [Edit Discovered Firewall on a Check Point Primary Management Station, page 4-47](#)
- [Verify Connectivity Between MARS and Check Point Devices, page 4-52](#)

If discovery of Check Point configuration settings is not enabled for MARS, you must perform the following manual configuration procedures:

- [Manually Add a Child Enforcement Module or Log Server to a Check Point Primary Management Station, page 4-41](#)
- [Specify Log Info Settings for a Child Enforcement Module or Log Server, page 4-49](#)

Before You Begin

To perform this procedure, you need the following information:

- A MARS account with Administrative privileges.
- A Check Point CMA or SmartCenter username and password that has READ access (minimum requirement).
- The client and server SIC DNs.
- If you are defining a CMA for Provider-1 or SiteManager-1, you must have the virtual IP address (VIP) for each CMA and CLM managed by the MDS.

Add a Check Point Primary Management Station to MARS

The primary management station represents one of the following:

- The SmartCenter server in a SmartCenter or SmartCenter Pro installation.
- A CMA of a Provider-1 or SiteManager-1 installation.



Note

Check Point 4.1, NG FP1, and NG FP2 devices are not officially supported. They cannot be configured to retrieve configuration information using CPML. However, they can be configured to retrieve logs using LEA. To configure one of these devices to work with the MARS, leave the Access IP field blank on the host that represents the base platform.

You must define each individual CMA of a Provider-1 or SiteManager installation, regardless of the release and version.

-
- Step 1** Select **Admin > System Setup > Security and Monitor Devices > Add**.
- Step 2** Do one of the following:
- Select **Add SW Security apps on a new host** from the Device Type list, and continue with [Step 3](#)
 - Select **Add SW security apps on existing host** from the Device Type list. Select the device to which you want to add the software application and click Add. Continue with [Step 7](#).
- Step 3** Specify values for the following fields:

- **Device Name** — Enter the name of the device. This name must exactly match the hostname shown in the Check Point user interface. MARS maps this name to the reporting IP address. This name is used in topology maps, queries, and as the primary management station in the Security and Monitoring Device list.
- **Access IP** — (Optional) This address is used to pull from a Check Point device using CPMI, enabling MARS to discover settings from this device. This address represents either a virtual IP address associated with a CMA or the physical IP address of the SmartCenter server. To learn more about the access IP address, its role, and dependencies, see [Understanding Access IP, Reporting IP, and Interface Settings, page 2-8](#).
- **Reporting IP** — Enter the IP address of the interface in the Check Point server from which MARS will pull traffic and audit logs. Check Point audit logs save information regarding user interaction with Check Point devices, such as log in and out of the Check Point user interface, initialize or revoke certificate, install or uninstall policy, create, modify, and delete objects, etc. No additional configuration is needed to turn on audit log on Check Point device.

This address represents either a virtual IP address associated with a CMA or the physical IP address of the SmartCenter server. To learn more about the reporting IP address, its role, and dependencies, see [Understanding Access IP, Reporting IP, and Interface Settings, page 2-8](#).

- Step 4** Under Enter interface information, enter the interface name, IP address, and netmask value of each interface in the Check Point server from which configuration information will be discovered and from which security event logs will be pulled.
- This address represents either a virtual IP address associated with a CMA or the physical IP address of the SmartCenter server. To learn more about the interface settings, its role, and dependencies, see [Understanding Access IP, Reporting IP, and Interface Settings, page 2-8](#).
- Step 5** (Optional) To enable MARS to monitor this device for anomalous resource usage, select **Yes** from the Monitor Resource Usage list.
- Result:* MARS monitors the device for anomalous consumption of resources, such as memory and CPU. If anomalies are detected, MARS generates an incident. Resource utilization statistics are also used to generate reports. For more information, see [Configuring Resource Usage Data, page 2-41](#).
- Step 6** Click **Apply** to save these settings.
- Step 7** Click **Next** to access the Reporting Applications tab.
- Step 8** Select the appropriate version of Check Point Opsec from the Select Application list, and click **Add**.
- The following options are available:
- **CheckPoint Opsec NG FP3**. Select this option for Check Point NG FP3 devices.
 - **CheckPoint Opsec NG AI**. Select this option for Check Point NG AI (R55) and Check Point NGX (R60) devices.

↓

General	Reporting Applications	Vulnerability Assessment
Enter reporting application:		
→ Device Name: Softie II → Select application: Select one Edit Remove Device Type Select one CheckPoint Opsec NG AI CheckPoint Opsec NG FP3 Cisco ACS 3.x Cisco CSA 4.x Cisco ICS 1.x Enterasys Dragon 6.x Entercept Entercept 2.5 Entercept Entercept 4.0 Foundstone FoundScan 3.0 Generic Web Server Generic ISS RealSecure 6.5 ISS RealSecure 7.0 IntruVert IntruShield 1.5 McAfee ePO 3.5 NetScreen IDP 2.1 Oracle Database Server Generic Snort Snort 2.0 Symantec Anti Virus 9.x Symantec ManHunt 3.x		
right © 2003, 2005 Cisco Systems, Inc. All rights reserved. Summary :: Incidents :: Query / Reports :: Rules :: Management		

Step 9 If you entered an address in the Access IP field on the host that represents this primary management station, specify values for the following fields:

- **Access Type** — This value identifies the authentication method to use for CPMI traffic, which is the protocol used to discover configuration information. Select **ASYMSSLC**, **CLEAR**, or **SSLCA**. The discovery operation identifies any child enforcement modules managed by this primary management station. It also discovers the NAT and ACL information necessary for NAT-based correlation, attack path calculation, and mitigation analysis. For more information on the access type, see [Select the Access Type for LEA and CPMI Traffic, page 4-29](#).

Access Information [Optional: for NAT-related session correlation, attack path calculation, and mitigation enter access information]		Secure Internal Communication Information	
→ *Access Type: SSLCA → *Access Port: 18190 (Default:18190) *Login: *Password: 	*Certificate: Select Certificate Add Edit *Client Entity SIC Name: *Server Entity SIC Name: SNMP RD Community: 		
Reporting Information → *LEA Access Type: SSLCA *LEA Port: 18184 (Default:18184)		Route Info	
Info Discover Cancel Submit			

- **Access Port** — Verify that the port number corresponds to the value specified in the CPMI_SERVER auth_port line of the fwopsec.conf file. The default authentication method for configuration discovery is SSLCA and data is passed on port 18190. For more information on this setting, see [Select the Access Type for LEA and CPMI Traffic, page 4-29](#).

- **Login** — Identifies the Check Point administrative account to be used to discover configuration settings.
- **Password** — Identifies the password associated with the Login account.

Step 10 Specify values for the following fields:

- **LEA Access Type** — If a log server is running on this primary management station select **ASYMSSLCA**, **CLEAR**, or **SSLCA**. You must have entered an address in the Reporting IP field on the host that represents this primary management station. This value identifies the authentication method to use for LEA traffic, which is the protocol used to pull security logs from the log server. For more information on the access type, see [Select the Access Type for LEA and CPMI Traffic, page 4-29](#).
- **LEA Port** — Verify that the port number corresponds to the value specified in the LEA_SERVER auth_port line of the `fwopsec.conf` file. The default authentication method for configuration discovery is SSLCA and data is passed on port 18184. For more information on this setting, see [Select the Access Type for LEA and CPMI Traffic, page 4-29](#).

Step 11 If this device uses SSLCA or ASYMSSLCA as an authentication method, specify values for the following fields (Otherwise, the authentication method is CLEAR. Skip to [Step 12](#).):

- **Certificate** — Either select the previously defined server from the list or click **Add** to define a new certificate authority and continue with [Add a Check Point Certificate Server, page 4-44](#).
- **Client SIC Name** — Enter the SIC DN of the OPSEC application for the MARS Appliance. This value was obtained in [Define an OPSEC Application that Represents MARS, page 4-24](#).
- **Server SIC Name** — Enter the SIC DN for this primary management station. This value was obtained in [Obtain the Server Entity SIC Name, page 4-27](#). Typically, this value is the SIC DN of the SmartCenter server or of the CMA. In the case of Provider-1 and SiteManager-1 NGX (R60), this value is the SIC DN of the MDS that manages the CMA.

Step 12 (Optional) To enable MARS to retrieve MIB objects for this reporting device, enter the device's read-only community string in the SNMP RO Community field.

Before you can specify the SNMP RO string, you must define an access IP address on host that represents the primary management station and you must configure the Access Information settings on the primary management station. MARS uses the SNMP RO string to perform resource utilization monitoring. Currently, it is not used for configuration or log discovery.

Step 13 (Optional) To enable MARS to monitor this device for anomalous resource usage, select **Yes** from the Monitor Resource Usage list.

Before you can enable this feature, you must provide a SNMP RO Community string.

Result: MARS monitors the device for anomalous consumption of resources, such as memory and CPU. If anomalies are detected, MARS generates an incident. Resource utilization statistics are also used to generate reports. For more information, see [Configuring Resource Usage Data, page 2-41](#).

Step 14 (Optional) To specify the route information for a firewall running on this primary management station, continue with [Define Route Information for Check Point Firewall Modules, page 4-47](#).

Step 15 (Optional) If you defined an access IP and selected and configured an access type, click **Discover** to determine the device settings.

Result: If the username and password are correct and the MARS Appliance is configured as an administrative host for the device, the "Discovery is done." dialog box appears when the discovery operation completes. Otherwise, an error message appears. After the initial pull, the MARS Appliance pulls based on the schedule that you define. For more information, see [Scheduling Topology Updates, page 2-39](#).

**Note**

Sometimes, the discovery operation times out, in which case you should try again. At other times, a message appears that states the discovery is taking a long time and that you should proceed to performing other tasks in MARS.

Step 16 To add this device to the MARS database and continue adding firewall modules manually, click **Submit**.

Result: The submit operation records the changes in the database tables. However, it does not load the changes into working memory of the MARS Appliance. The activate operation loads submitted changes into working memory.

Step 17 Do one of the following:

- To manually define the child enforcement modules that are managed by this primary management station, continue with [Manually Add a Child Enforcement Module or Log Server to a Check Point Primary Management Station, page 4-41](#).
- To edit the settings of the discovered child enforcement modules, continue with [Edit Discovered Firewall on a Check Point Primary Management Station, page 4-47](#).

Step 18 Click **Activate**.

Result: Once the MARS Appliance is activated, it connects to the Check Point log modules and retrieves the traffic and audit logs. MARS also begins to sessionize events generated by this device and evaluate those events using the defined inspection and drop rules. Any events published by the device to MARS before activation can be queried using the reporting IP address of the device as a match criterion. For more information on the activate action, see [Activate the Reporting and Mitigation Devices, page 2-27](#).

Manually Add a Child Enforcement Module or Log Server to a Check Point Primary Management Station

If you have not enabled configuration discovery on the primary management station or if one or more of the managed firewalls uses a log server that is not managed by the primary management station, you can manually define firewalls or log servers. Your goal should be to represent all of the firewalls managed by this primary management station and all log servers used by those firewalls and the primary management station. While MARS does not discover configuration settings of the firewalls, it uses the defined information to discover topology, calculate attack paths, and identify preferred mitigation points in the network.

For example, if you are defining a primary management station that represents a CMA, you must define the CLM associated with that CMA. Any firewalls managed under that CMA may either act as their own log servers, publish information to the CLM, or publish information to a MLM. In the case of the later, you must define that relationship by defining the firewalls and then specifying which log servers pull their traffic and audit logs. First, however, must also define the MLM settings, as it is a log server that external to the perspective of the CMA, and it cannot be referred by a firewall until it has been defined. The CLM, however, would be considered part of the CMA (assuming the reporting IP and LEA settings are specified), so you would not define a separate child enforcement module to represent it. Instead, you would select the Management option in the Log Info dialog for firewalls that use the CLM as their log server. For more information on selecting the log server option, see [Specify Log Info Settings for a Child Enforcement Module or Log Server, page 4-49](#).

To manually define a child enforcement module that is managed by the primary management station or a log server to which either the primary management station or a child enforcement module publishes its audit and security logs, follow these steps:

- Step 1** Select **Admin > System Setup > Security and Monitor Devices**.
- Step 2** From the Security and Monitor Devices list, select the host that represents the primary management station and click **Edit**.
- Such devices have CheckPoint Management Console as an entry in the Device Type column.
- Step 3** Click **Next** to access the Reporting Applications tab.

↓

General	Reporting Applications	Vulnerability Assessment Info

Enter reporting application:

→ Device Name: DEV-CMA

→ Select application: Select one Add

Edit
Remove

Device Type
☐ CheckPoint Management Console

Done

143632

- Step 4** Select the **CheckPoint Management Console** check box in the Device Type list and click **Edit**.
- The Access Information page appears.

Access Information
 [Optional: for NAT-related session correlation, attack path calculation, and mitigation enter access information]

→ *Access Type: SSLCA

→ *Access Port: 18190 (Default:18190)

*Login: eng

*Password: ●●●●●●

Reporting Information

→ *LEA Access Type: SSLCA

*LEA Port: 18184 (Default:18184)

Secure Internal Communication Information

*Certificate: testServer Add Edit

*Client Entity SIC Name: testServer

*Server Entity SIC Name: testServer

SNMP RO Community:

Route Info

Firewall & Log Server Settings

Add
Edit
Delete
Log Info
Route Info

Info
Discover
Cancel
Submit

143627

- Step 5** Click **Add** under Firewall & Log Server Settings.
- Result:* The list of available hosts appears.
- Step 6** Do one of the following:

- Select the host on which the child enforcement module is running, click **Change Existing**, and continue with [Step 7](#)

Result: A page with a read-only device name appears, prompting you to specify the SNMP RO Community string.

→ *Device Name: Test host

→ SNMP RO Community:

Add Interface **Remove Interface**

Name: IP Address: Network Mask:

☐ ether0 192 168 3 13 255 255 255 255

Cancel Submit

143623

- Click **Add New** to define a new host, and continue with [Step 7](#)

Result: A page appears, prompting you to specify device name and SNMP RO Community string.

→ *Device Name:

→ SNMP RO Community:

Add Interface **Remove Interface**

Name: IP Address: Network Mask:

☐ ether0

Cancel Submit

143787

- Step 7** Enter the name of the child enforcement module or log server in the Device Name field. MARS maps this name to the IP address specified in the interfaces. This name is used in topology maps, queries, and appears in the Children column of the base Check Point module in the Security and Monitoring Device list.
- Step 8** (Optional) To enable MARS to retrieve MIB objects for this reporting device, enter the child enforcement module's read-only community string in the SNMP RO Community field.

Before you can specify the SNMP RO string, you must define an access IP address on host that represents the primary management station. MARS uses the SNMP RO string to read MIBs related to a reporting device's CPU usage, network usage, and device anomaly data and to discover device and network settings.

- Step 9** Under Enter interface information, enter the interface name, IP address, and netmask value of each interface installed in the child enforcement module or log server.
- These interfaces include the ones from which the configuration information will be discovered and security event logs will be pulled. To learn more about the interface settings, its role, and dependencies, see [Understanding Access IP, Reporting IP, and Interface Settings, page 2-8](#).
- Step 10** Click **Submit** to add this module to the primary management station.
- Step 11** (Optional) To specify the route information for a firewall child enforcement module, continue with [Define Route Information for Check Point Firewall Modules, page 4-47](#).
- Step 12** If the child enforcement module does not propagate its logs to the primary management station or if you are defining a log server that is not managed by this primary management station, you must specify where its logs are stored. Continue with [Specify Log Info Settings for a Child Enforcement Module or Log Server, page 4-49](#).
- Step 13** Repeat [Step 5](#) through [Step 12](#) for each child enforcement module that is managed by this primary management station and each log server that is used by the primary management station or child enforcement modules.
- Step 14** To add this device to the MARS database, click **Submit**.
- Result:* The submit operation records the changes in the database tables. However, it does not load the changes into working memory of the MARS Appliance. The activate operation loads submitted changes into working memory.
- Step 15** Click **Done** to close the Reporting Applications tab and return to the Security and Monitoring Devices list.
- Step 16** Click **Activate**.
- Result:* Once the MARS Appliance is activated, it connects to the Check Point log modules and retrieves the traffic and audit logs. MARS also begins to sessionize events generated by this device and its modules and evaluate those events using the defined inspection and drop rules. Any events published by the device to MARS before activation can be queried using the reporting IP address of the device as a match criterion. For more information on the activate action, see [Activate the Reporting and Mitigation Devices, page 2-27](#).
-

Add a Check Point Certificate Server

When defining a Check Point module that uses secured communications, you must identify the certificate sever that authenticates the SICs provided by the client and the server. Typically, a SmartCenter server or the CMA has its own certificate server, however, your configuration may use a central server. If that is the case, you must define the certificate server as part of a defining a base or child enforcement module.




Note

This procedure assumes you have been refer to it, and that you are in the middle of defining a primary management station or child enforcement module.

To define a certificate server, follow these steps:

Step 1 Click **Add** to define the settings for the certificate authority.



Step 2 Specify values for the following fields:

- **Certificate Authority IP Address** — Typically, this IP address is the physical IP address of the SmartCenter server or the virtual IP address of the CMA. In the case of Provider-1 and SiteManager-1 NGX (R60), this IP address represents the physical IP address of the MDS that manages the CMA.
- **Client SIC Name** — Enter the SIC DN of the OPSEC application for the MARS Appliance. This value was obtained in [Define an OPSEC Application that Represents MARS](#), page 4-24.
- **Activation Key** — This value was also provided in [Define an OPSEC Application that Represents MARS](#), page 4-24.

Step 3 Click **Pull Certificate**.

Result: A message box appears stating “Discovery is done.”

A certificate can be pulled only once for an OPSEC Application. If for any reason the pull operation fails, you must reset the certificate using the CheckPoint SmartDashboard. For more information, see [Reset the OPSEC Application Certificate of the MARS Appliance](#), page 4-33.

Step 4 Click **Close**.

Edit Discovered Log Servers on a Check Point Primary Management Station

After performing a discovery operation, you must edit each discovered log servers. The purpose of editing this log server is to identify that it is its own log server and to provide the SIC communication settings.

To edit a discovered log server, follow these steps:

Step 1 Under Firewall & Log Server Settings, select the check box next to the desired log server, and click **Log Info**.

Step 2 Select **Self**.

Management ☐ *Reporting IP: [][][][]

Log Server ☐ Certificate: [Select Certificate] [Add] [Edit]

Self ☒ Client SIC Name: []

Server SIC Name: []

*Logging Access Type: [SSLCA]

*Logging Access Port: [18184] (Default:18184)

[Cancel] [Submit]

Step 3 Specify values for the following fields:

- **Reporting IP** — Enter the IP address of the interface in the log server from which MARS will pull security event logs. This address represents either a virtual IP address associated with a CLM, an MLM, or another log server. To learn more about the reporting IP address, its role, and dependencies, see [Understanding Access IP, Reporting IP, and Interface Settings, page 2-8](#).
- **Logging Access Type** — This value identifies the authentication method to use for LEA traffic, which is the protocol used to pull security logs from the log server. Select **ASYMSSLCA**, **CLEAR**, or **SSLCA**. For more information on the access type and port, see [Select the Access Type for LEA and CPMI Traffic, page 4-29](#).
- **Logging Port** — Verify that the port number in the corresponds to the value specified in the LEA_SERVER auth_port line of the fwopsec.conf file on this log server. The default authentication method for configuration discovery is SSLCA and data is passed on port 18184.

Step 4 If this log server uses SSLCA or ASYMSSLCA as an authentication method, specify values for the following fields (Otherwise, the authentication method is CLEAR. Skip to [Step 5](#)):

- **Certificate** — Either select the previously defined server from the list or click **Add** to define a new certificate authority and continue with [Add a Check Point Certificate Server, page 4-44](#).
- **Client SIC Name** — Enter the SIC DN of the OPSEC application for the MARS Appliance. This value was obtained in [Define an OPSEC Application that Represents MARS, page 4-24](#).
- **Server SIC Name** — Enter the SIC DN for the child enforcement module. This value was obtained in [Obtain the Server Entity SIC Name, page 4-27](#). Typically, this value is the SIC DN of the SmartCenter server or of the CMA. In the case of Provider-1 and SiteManager-1 NGX (R60), this value is the SIC DN of the MDS that manages the CMA.

Step 5 Click **Submit** to save your changes to this log server.

Step 6 Repeat [Step 1](#) through [Step 5](#) for each discovered log server.

Edit Discovered Firewall on a Check Point Primary Management Station

After performing a discovery operation, you must edit any discovered firewalls. You must specify which log server the firewall uses, define the route information, and if you want to monitor resource utilization, you must specify the SNMP RO community string.



Note

When editing a Check Point Firewall, never select a Check Point Firewall from the Security and Monitoring Devices list. Instead, select the Check Point Management Console that acts as the primary management station for that firewall.



Note

You must configure the discovered log servers and define any log servers not managed by the primary management station before editing the discovered firewalls. To configure the discovered log servers, see [Edit Discovered Log Servers on a Check Point Primary Management Station, page 4-45](#). To manually define log servers, see [Manually Add a Child Enforcement Module or Log Server to a Check Point Primary Management Station, page 4-41](#).

To edit a discovered firewall, follow these steps:

- Step 1** Under Firewall & Log Server Settings, select the check box next to the desired firewall.
- Step 2** Click **Edit**.
- Step 3** (Optional) To enable MARS to retrieve MIB objects for this reporting device, enter the child enforcement module's read-only community string in the SNMP RO Community field.

Before you can specify the SNMP RO string, you must define an access IP address on host that represents the primary management station. MARS uses the SNMP RO string to read MIBs related to a reporting device's CPU usage, network usage, and device anomaly data and to discover device and network settings.
- Step 4** Click **Submit**.
- Step 5** To define the route settings for this firewall, continue with [Define Route Information for Check Point Firewall Modules, page 4-47](#).
- Step 6** To select the log server used by this firewall, continue with [Specify Log Info Settings for a Child Enforcement Module or Log Server, page 4-49](#).
- Step 7** Repeat [Step 1](#) through [Step 6](#) for each discovered firewall.

Define Route Information for Check Point Firewall Modules

To perform attack path analysis and to provide suggested mitigation configurations, MARS must understand the static routes that are defined on a firewall module. This requirement is true for firewalls running on the primary management station as well as for each firewall child enforcement module managed by the primary management station. To provide this information, you must define the routes manually in the MARS web interface. You will need a list of the routes for all interfaces in the firewall before you attempt to enter this information.

**Note**

You do not need to specify which interface the route is associated with. MARS derives this information based on the interface settings you have specified for the host.

To define the static routes used by a firewall, follow these steps:

Step 1 Do one of the following:

- To specify the route information for the primary management station, click **Route Info** on the primary management station page.
- To specify the route information for a firewall child enforcement module, select the server under Device Type, click **Route Info**.

Result: The Route Information dialog box appears.

Step 2 Specify values for the following fields:

- **Destination Address** — Enter the internal or external destination network address
- **Destination Mask** — Enter the corresponding network mask value.
- **Next Hop Address** — Enter the IP address of the default gateway.
- **Metric** — Identifies the priority for using a specific route. When routing network packets, a gateway device uses the rule with the most specific network within the rule's definition. Only in cases where two routing rules have the same network is the metric used to determine which rule is applied. If they are the same, the lowest metric value takes priority. If no routing rule exists, the network packet is dropped, and if the gateway is not detected (dead), the network packet is dropped.

A *metric* is a measurement of the cost of a route based on the number of hops (hop count) to the network on which a specific host resides. Hop count refers to the number of networks that a network packet must traverse, including the destination network, before it reaches its final destination.

Because the hop count includes the destination network, all directly connected networks have a metric of 1. For the metric value, specify a number between 1 and 15.

Step 3 Click **Submit** to add the route to the list of routes**Step 4** Repeat • through [Step 3](#) for each route defined on the firewall.**Step 5** Click **Close** to return to the Access Information page.

Specify Log Info Settings for a Child Enforcement Module or Log Server

There are two occasions when you must define the log settings manually:

- If you do not discover the settings of the primary management station, which does discover the log settings.
- If the child enforcement module does not propagate its logs up to the primary management station.

Three options exist for manually specifying the log settings:

- **Management.** Identifies that the child enforcement module propagates its logs up to the primary management station, the MLM or the SmartCenter server. You do not specify these settings; they are derived from the settings on the primary management station. However, the option is available if the configuration of a child enforcement module changes. If the primary management station is the log server for a child enforcement module, the log server information is populated when you perform the test connectivity operation.

Figure 4-1 Log Information Published to Primary Management Station

☒ Management	Reporting IP	10.1.1.17
	Certificate:	testServer
	Client SIC Name:	testServer
☐ Log Server	Server SIC Name:	testServer
	Logging Access Type:	SSLCA
	Logging Access Port:	18184
☐ Self		

- **Log Server.** Identifies that another log server, such as a CLM, is acting as the log server for this child enforcement module. You must either select a pre-defined log server or define the settings for a new one and select it.
- **Self.** Identifies that the child enforcement module is acting as its own log server. In this case, you must specify the communication settings required to pull the logs from that module or server.

To specify the log server settings of a child enforcement module manually, follow these steps:

- Step 1** (Firewall only) If a child enforcement module does not propagate its log information to the primary management station, then select that child enforcement module under Device Type, click **Log Info**, and do one of the following:
- To specify that the child enforcement module is acting as its own log server, select **Self** and continue with [Step 3](#), omitting the Device Name field.

Figure 4-2 Log Information Published to Self

The screenshot shows the 'Log Information' dialog box with the 'Self' radio button selected. The fields are as follows:

- Management:** *Reporting IP: [][][][]
- Log Server:** Certificate: [Select Certificate] [Add] [Edit]
- Self (selected):**
 - Client SIC Name: []
 - Server SIC Name: []
 - *Logging Access Type: [SSLCA]
 - *Logging Access Port: [18184] (Default:18184)

Buttons at the bottom: [Cancel] [Submit]

- To specify an alternate log server, select **Log Server**, and continue with [Step 2](#).

Result: The Log Information dialog box appears, and the desired option is selected.

Step 2 Do one of the following:

- Select a predefined log server from the Select list, click **Submit**, and continue with [Step 5](#).

The screenshot shows the 'Log Information' dialog box with the 'Log Server' radio button selected. The fields are as follows:

- Management:** [Select] [Add] [Edit]
- Log Server (selected):**
- Self:**

Buttons at the bottom: [Cancel] [Submit]

- Click **Add** to define a new log server.

The screenshot shows the 'Add' dialog box for defining a new log server. The fields are as follows:

- *Device Name: []
- *Reporting IP: [][][][]
- Certificate: [Select Certificate] [Add] [Edit]
- Client SIC Name: []
- Server SIC Name: []
- *Logging Access Type: [SSLCA]
- *Logging Access Port: [18184] (Default:18184)

Buttons at the bottom: [Back] [Cancel] [Submit]

Step 3 Specify values for the following fields:

- Device Name** — Enter the name of the log server. MARS maps this name to the reporting IP address. This name is used in topology maps, queries, and as the primary management station in the Security and Monitoring Device list. For devices that support the discovery operation, such as routers and

firewalls, MARS renames this field's value to match the name discovered in the device configuration, which typically uses the *hostname.domain* format. For devices that cannot be discovered, such as Windows and Linux hosts and host applications, MARS uses the provided value.

- **Reporting IP** — Enter the IP address of the interface in the log server from which MARS will pull security event logs. This address represents either a virtual IP address associated with a CLM, an MLM, or another log server. To learn more about the reporting IP address, its role, and dependencies, see [Understanding Access IP, Reporting IP, and Interface Settings, page 2-8](#).
- **Logging Access Type** — This value identifies the authentication method to use for LEA traffic, which is the protocol used to pull security logs from the log server. Select **ASYMSSLCA**, **CLEAR**, or **SSLCA**. For more information on the access type, see [Select the Access Type for LEA and CPML Traffic, page 4-29](#).
- **Logging Port** — Verify that the port number in the corresponds to the value specified in the LEA_SERVER auth_port line of the `fwopsec.conf` file on this log server. The default authentication method for configuration discovery is SSLCA and data is passed on port 18184. For more information on this setting, see [Select the Access Type for LEA and CPML Traffic, page 4-29](#).

Step 4 If this log server uses SSLCA or ASYMSSLCA as an authentication method specify values for the following fields (Otherwise, CLEAR is the authentication method for Access Type and LEA Access Type, and you should skip to [Step 5](#)):

- **Certificate** — Either select the previously defined server from the list or click **Add** to define a new certificate authority and continue with [Add a Check Point Certificate Server, page 4-44](#).
- **Client SIC Name** — Enter the SIC DN of the OPSEC application for the MARS Appliance. This value was obtained in [Define an OPSEC Application that Represents MARS, page 4-24](#).
- **Server SIC Name** — Enter the SIC DN for the child enforcement module. This value was obtained in [Obtain the Server Entity SIC Name, page 4-27](#). Typically, this value is the SIC DN of the SmartCenter server or of the CMA. In the case of Provider-1 and SiteManager-1 NGX (R60), this value is the SIC DN of the MDS that manages the CMA.

Step 5 To add this child enforcement module to the primary management station, click **Submit**.

Step 6 To add the primary management station to the MARS database, click **Submit**.

Result: The submit operation records the changes in the database tables. However, it does not load the changes into working memory of the MARS Appliance. The activate operation loads submitted changes into working memory.

Step 7 Click **Done** to close the Reporting Applications tab and return to the Security and Monitoring Devices list.

Step 8 Click **Activate**.

Result: Once the MARS Appliance is activated, it connects to the Check Point log modules and retrieves the traffic and audit logs. MARS also begins to sessionize events generated by this device and its modules and evaluate those events using the defined inspection and drop rules. Any events published by the device to MARS before activation can be queried using the reporting IP address of the device as a match criterion. For more information on the activate action, see [Activate the Reporting and Mitigation Devices, page 2-27](#).

Verify Connectivity Between MARS and Check Point Devices

After defining the Check Point device and clicking **Activate** in the MARS web interface, the MARS Appliance connects to the log servers and pulls the traffic and audit logs stored on them. You can verify that these transactions are successful using the following method:

- Perform an ad hoc query for Event Types/Sessions specify to the Check Point primary management station.

Result: The netstat command should display two connections per log server.

Remove a Firewall or Log Server from a Check Point Primary Management Station

If the configuration of your network changes so that a firewall or log server is no longer managed by the primary management station under which it is defined, you must remove the child enforcement module.

To remove a child enforcement module from the primary management station, follow these steps:

-
- Step 1** Select **Admin > System Setup > Security and Monitor Devices**.
- Step 2** From the Security and Monitor Devices list, select the host that represents the primary management station of the Check Point server and click **Edit**.
- Such devices have CheckPoint Management Console as an entry in the Device Type column.
- Step 3** Click **Next** to access the Reporting Applications tab.

↓

General	Reporting Applications	Vulnerability Assessment Info

Enter reporting application:

→ Device Name: DEV-CMA

→ Select application: Select one ▼ Add

Edit Remove

Device Type

☐ CheckPoint Management Console

Done

143632

- Step 4** Select CheckPoint Management Console from the Device Type list and click **Edit**.
- The Access Information page appears.
- Step 5** Under Firewall & Log Server Settings, check the box next to the child enforcement module that you want to remove.
- Step 6** Click **Remove**.
- Result:* The Confirmation screen appears.

Step 7 Click **Submit** to remove the child enforcement module from the primary management station.

Troubleshooting MARS and Check Point

The following information can be used to troubleshoot communication issues between the MARS Appliance and Check Point components.

- To view attack information by user, run a query where the device is a Check Point device.
- If you attempt to discover the certificate and it returns to the CheckPoint Certificate screen instead of displaying the “Discovery done.” message box, then the discover operation failed. The likely cause is an incorrect SIC value.



Note

A certificate can be pulled only once for an OPSEC Application. If for any reason the pull operation fails, you must reset the certificate using the CheckPoint SmartDashboard. For more information, see [Reset the OPSEC Application Certificate of the MARS Appliance, page 4-33](#).

- If the device discovery operation fails, click the **View Error** button for a detailed error message.

Common reasons for failure of device discovery are as follows:

- client SIC DN name or server SIC DN name is incorrect. Use copy and paste from SmartDashboard to avoid erroneous entry.
- Invalid Certificate used.
- Invalid user name, password, or both used. Verify that the credentials provided for the Access IP match an Check Point account with administrative privileges.
- Unsupported version of Check Point. (Discovery works only with NG FP3 and above. Internally we have tested up to Version R60)
- Invalid authentication method used. The default method is SSLCA. Check the `fwopsec.conf` file to determine which method is used. CS-MARS currently support only three authentication methods for CPMI communication: SSLCA, ASYM_SSLCA and CLEAR. For more information on specifying these settings, see [Select the Access Type for LEA and CPMI Traffic, page 4-29](#).
- Invalid access port. Default port for secured CPMI-based communication is TCP 18180. Check the `fwopsec.conf` to verify the configured port.
- The MARS Appliance does not have access to port 18190, or an alternate specified in `fwopsec.conf` for CPMI. At the CLI of the MARS Appliance, use the **telnet** command to test the access port. For more information on **telnet**, see [Verify Communication Path Between MARS Appliance and Check Point Devices, page 4-32](#).
- The policy database was not installed after creating OPSEC Application in the SmartDashboard.
- Firewall policies were not created and installed that permitted the MARS Appliance to connect to the Check Point primary management station. For information, see [Create and Install Policies, page 4-31](#).

For additional Check Point discovery-related debug information, use the **pnlog** command at the CLI of the MARS Appliance. You can use the `cpdebug` attribute to specify appropriate debug level. Level 9 presents all debug messages. You can view the debug messages using the **pnlog showlog cpdebug** command at the CLI. For more information on **pnlog**, see [pnlog, page A-18](#) in the *Install and Setup Guide for Cisco Security Monitoring, Analysis, and Response System*.



Configuring VPN Devices

VPN devices provide MARS with information about remote hosts, login in requests and denials, and access times. With this data, MARS can provide end-to-end attack path analysis and identify the VPN device through which attacks are launched.

This chapter explains how to bootstrap and add the following VPN device to MARS:

- [Cisco VPN 3000 Concentrator, page 5-1](#)

Cisco VPN 3000 Concentrator

MARS can receive and process events from the Cisco VPN 3000 Concentrator, versions 4.0.1 and 4.7. To enable communications, you must perform two tasks:

- [Bootstrap the VPN 3000 Concentrator, page 5-1](#)
- [Add the VPN 3000 Concentrator to MARS, page 5-2](#)

Bootstrap the VPN 3000 Concentrator

To configure a Cisco VPN 3000 Concentrator to generate and publish events to the MARS Appliance, you must verify that the correct events are generated in the correct format, and you must direct the Cisco VPN 3000 Concentrator to publish syslog events to the MARS Appliance.

To configure Cisco VPN 3000 Concentrator to send syslog events to MARS, follow these steps:

-
- Step 1** Open your browser and log in to the Cisco VPN 3000 Concentrator Series Manager.
- Step 2** From the tree on the left, select **Configuration > System > Events > General**.

Configuration | System | Events | General

This section lets you configure default event handling.

Save Log on Wrap	☐	Check to save the event log to a file on wrap.
Save Log Format	Multiline	Select the format of the saved log files.
FTP Saved Log on Wrap	☐	Check to automatically FTP the saved log to a remote destination.
E-mail Source Address		Enter the e-mail address that appears in the From: field.
Syslog Format	Original	Select the format of Syslog messages.
Events to Log	Severities 1-5	Select the events to enter in the log.
Events to Console	Severities 1-3	Select the events to display on the console.
Events to Syslog	Severities 1-5	Select the events to send to a Syslog Server.
Events to E-mail	None	Select the events to send to an E-mail Recipient.
Events to Trap	Severities 1-3	Select the events to send to an SNMP Trap Destination.

143210

- Step 3** Verify that the Syslog Format is Original.
- Step 4** Select **Severities 1-5** in the Events to Syslog field.
- Step 5** From the tree on the left, select **Configuration > System > Events > Syslog Servers**.
- Step 6** Click **Add** to define a target syslog server.

Configuration | System | Events | Syslog Servers | Add

Add a syslog server.

Syslog Server	cs-mars	Enter the IP address or hostname of the syslog server.
Port	514	Enter the port used by the syslog server.
Facility	Local 7	Select the syslog facility tag for events sent to this server.
Add Cancel		

143209

- Step 7** In the Syslog Server field, enter the IP address or hostname of the MARS Appliance.
- Step 8** Click **Add** to save the syslog server settings.
- Step 9** Click **Save** in the top-right corner to save all changes.

Add the VPN 3000 Concentrator to MARS

To add the VPN 3000 Concentrator to MARS, follow these steps:

- Step 1** Select **Admin > Security and Monitor Devices > Add**.
- Step 2** Select either **Cisco VPN Concentrator 4.0.1** or **Cisco VPN Concentrator 4.7** from the Device Type list.

Device Type: Cisco VPN Concentrator 4.0.3

→ *Device Name:

→ Access IP: ...

→ Reporting IP: ...

→ *Access Type: Select

SNMP RO Community:

→ Monitor Resource Usage: NO

143208

- Step 3** Enter the name of the VPN Concentrator in the Device Name field.
- Step 4** Enter the IP address used to administer the VPN Concentrator in the Access IP field.
- Step 5** Enter the IP address from which the syslog messages are sent to MARS in the Reporting IP field.
- Step 6** Select **SNMP** from the Access Type list.
- Step 7** (Optional) To enable MARS to retrieve MIB objects for this Concentrator, enter the device's read-only community string in the SNMP RO Community field.
- MARS uses the SNMP RO string to read MIBs related to the reporting device's CPU usage and other device anomaly data.
- Step 8** Click **Discover**.
- Step 9** Click **Submit**.
-



Configuring Network-based IDS and IPS Devices

Network intrusion detection and intrusion prevention systems are a critical source for identifying active attacks to MARS.

This chapter explains how to bootstrap and add the following network-based IDS and IPS devices to MARS:

- [Cisco IDS 3.1 Sensors, page 6-1](#)
- [Cisco IDS 4.0 and IPS 5.x Sensors, page 6-5](#)
- [Cisco IPS Modules, page 6-9](#)
- [ISS Site Protector, page 6-13](#)
- [ISS RealSecure 6.5 and 7.0, page 6-17](#)
- [IntruVert IntruShield, page 6-22](#)
- [Snort 2.0, page 6-28](#)
- [Symantec ManHunt, page 6-29](#)
- [NetScreen IDP 2.1, page 6-31](#)
- [Enterasys Dragon 6.x, page 6-33](#)

Cisco IDS 3.1 Sensors

Before you add the Cisco IDS 3.1 device, make sure that you have configured the Cisco IDS device for the MARS to retrieve the device configuration. The device configuration would be used for mapping of the logs received by MARS.

When configuring the IDS device to send logs to the MARS, you must use the exact name of the MARS Appliance. To determine the name of the appliance, select **Admin > System Setup > Configuration Information** and review the value in the Name field.

Configure Sensors Running IDS 3.1

Step 1 Log in to the Cisco IDS device.

Step 2 Change to directory that has all the configurations files that need to be edited:


```
cd /usr/nr/etc
```

Step 3 You need to edit 4 files (**organizations**, **hosts**, **routes** and **destinations**) that are in this directory. In the **organizations** file add a line indicating your organization name or grouping;

e.g., 1 protego

where 1 is the item number followed by the organization name `protego`. If there is already item in this file, simply increase the item number (has to be unique).

Figure 6-1 Add MARS Information to Cisco IDS 3.1 Organizations File




In the **hosts** file add a line indicating your MARS appliances' name associated to the organization that was previously added in the organizations file;

e.g., 2001.1 pnmars.protego

where 2001.1 is a unique item number followed by the MARS appliances' name and organization name `protego`. If there is already items in this file, simply increase the item number (has to be unique).

Figure 6-2 Add MARS Information to Cisco IDS 3.1 Hosts File




In the **routes** file add a line indicating your MARS appliances' name and its IP address;

e.g., `pnmars.protego 1 10.1.1.10 45000 1 5`

where `pnmars.protego` is the MARS's name (with organizations' name) followed by 1 then the MARS appliances' IP address.

The 45000 is the port number that the IDS will use to send its logs to MARS. Add a 1 follows by a 5 at the end of this line (these numbers are not used by MARS).

Figure 6-3 Add MARS Information to Cisco IDS 3.1 Routes File



In the **destinations** file add a line indicating your MARS appliances' name (as defined in the routes file) the client process that the appliance is using to listen for events from the sensor (in this case `smid`), and the list of log types you want sent to the appliance as a comma separated list:

e.g., `pnmars.protego smid ERRORS, EVENTS, COMMANDS`

where `pnmars.protego` is the MARS's name (with organizations' name) followed by `smid` and the list of log types that the loggerd daemon will publish to the appliance.

Figure 6-4 Add MARS Information to Cisco IDS 3.1 Destinations File

```

10.1.1.5 - default - SSH Secure Shell
File Edit View Window Help
Quick Connect Profiles
Generated: Wed Sep 18 20:17:16 2002
# Template: C:\Program Files\Cisco Systems\Cisco Secure Policy
Manager\bin\templates\common\etc\destinations.template
# Sensor Version: 3.1(2)S25
# Sensor OS: SunOS

1 Sensor3.protego loggerd 1 ERRORS,EVENTS,COMMANDS
2 NISVR1.protego smid 1 ERRORS,EVENTS,COMMANDS
3 sunshine.protego smid 1 ERRORS,EVENTS,COMMANDS
4 pnguard.protego smid 1 ERRORS,EVENTS,COMMANDS
5 pnmars50.protego smid 1 ERRORS,EVENTS,COMMANDS
6 pluto.protego smid 1 ERRORS,EVENTS,COMMANDS
7 athens.protego smid 1 ERRORS,EVENTS,COMMANDS
8 pnmars50_1.protego smid 1 ERRORS,EVENTS,COMMANDS
9 pnmars50_2.protego smid 1 ERRORS,EVENTS,COMMANDS
10 pnmars100_33.protego smid 1 ERRORS,EVENTS,COMMANDS
11 pnmars50_4.protego smid 1 ERRORS,EVENTS,COMMANDS
12 pnlister.protego smid 1 ERRORS,EVENTS,COMMANDS

Connected to 10.1.1.5 SSH2 - 3des-cbc - hmac-md5 - none 143228

```

Step 4 Once you've edited these four files (organizations, hosts, routes, and destinations), reboot the sensor using the following commands:

- a. `nrstop`
- b. `nstart`

Add and Configure a Cisco IDS 3.1 Device in MARS

To add and configure a Cisco IDS device in MARS, follow these steps:

- Step 1** Click **Admin > System Setup > Security and Monitor Devices > Add**.
- Step 2** Select **Cisco IDS 3.1** from the Device Type list.
- Step 3** Enter the hostname of the sensor in the Device Name field.
The Device Name value must be identical to the configured sensor name.
- Step 4** Enter the administrative IP address in the Access IP field.
- Step 5** Enter the administrative IP address in the Reporting IP field.
The Reporting IP address is the same address as the administrative IP address.
- Step 6** Select either **SSH** or **TELNET**.
- Step 7** Enter "**netrangr**" as the **Login** and its **Password**.

When adding a Cisco IDS 3.1 device, use the `netrangr` username or some other username that is not the root login for the sensor. Using the root login causes MARS to fail to parse the login prompt correctly, which in turn, cause the Test Connectivity to fail.

Figure 6-5 *Configure Cisco IDS 3.1*

Device Type: Cisco IDS 3.1

→ *Device Name: HQ-NIDS1

→ *Access IP: 10 1 1 100

→ *Reporting IP: 10 1 1 100

→ *Access Type: SSH

Login: netrangr

Password:

143212

Step 8 For attack path calculation and mitigation, add networks into the Monitored Networks field.

- a. Click the **Select a Network** or **Define a Network** radio button.
 - In the Select a Network list, click a network.
 - In the Define a Network field, enter its network IP and network mask information.
- b. Click **Add** to move the selected networks into the Monitored Networks field.

Step 9 (Optional) To discover the device settings, click **Discover**.

Step 10 Click **Submit**.

Cisco IDS 4.0 and IPS 5.x Sensors

Adding a Cisco IDS or IPS network sensor to MARS involves two parts:

1. [Bootstrap the Sensor, page 6-5](#)
2. [Add and Configure a Cisco IDS or IPS Device in MARS, page 6-6](#)

The following topic supports Cisco IDS and IPS devices:

- [View Detailed Event Data for Cisco IPS Devices, page 6-9](#)



Note

If you've imported your sensor definitions using the seed file format, as specified in [Load Devices From the Seed File, page 2-24](#), you must define the networks monitored by the sensor.

Bootstrap the Sensor

Preparing a sensor to be monitored by MARS involves two steps:

- [Enable the Access Protocol on the Sensor, page 6-6](#)
- [Enable the Correct Signatures and Actions, page 6-6](#)

Enable the Access Protocol on the Sensor

The configuration of the sensor depends on the version of the software that is running on the sensor. The following topics identify the requirements of each version:

- [Cisco IDS 4.x Software, page 6-6](#)
- [Cisco IPS 5.x Software, page 6-6](#)

Cisco IDS 4.x Software

For Cisco IDS 4.x devices, MARS pulls the logs using RDEP over SSL. Therefore, MARS must have HTTPS access to the sensor. To prepare the sensor, you must enable the HTTP server on the sensor, enable TLS to allow HTTPS access, and make sure that the IP address of MARS is defined as an allowed host, one that can access the sensor and pull events. If the sensors have been configured to allow access from limited hosts or subnets on the network, you can use the `accessList ipAddress ip_address netmask` command to enable this access.

Cisco IPS 5.x Software

For Cisco IPS 5.x devices, MARS pulls the logs using SDEE over SSL. Therefore, MARS must have HTTPS access to the sensor. To prepare the sensor, you must enable the HTTP server on the sensor, enable TLS to allow HTTPS access, and make sure that the IP address of MARS is defined as an allowed host, one that can access the sensor and pull events. If the sensors have been configured to allow access from limited hosts or subnets on the network, you can use the `access-list ip_address/netmask` command to enable this access.

Enable the Correct Signatures and Actions

If the signature actions are correctly configured, MARS can display the trigger packet information for the first event that fires a signature on a Cisco IDS or IPS device. MARS is also able to pull the IP log data from Cisco IDS and IPS devices, however, this operation is system intensive. Therefore, you should select the set of signatures that generate IP log data carefully.

When configuring the active signatures on a Cisco IDS or IPS device, you must specify the alert action and the action that generates the desired data:

- To view trigger packets, you must enable the “produce-verbose-alert” action.
- To view IP logs, you must enable the alert or “produce-verbose-alert” action and the “log-pair-packets” action.



Caution

Configuring IP logging and verbose alerts on the sensor is system intensive and does affect the performance of your sensor. In addition, it affects the performance of your MARS Appliance. Because of these effects, you be cautious in configuring signatures to generate IP logs.

Add and Configure a Cisco IDS or IPS Device in MARS

To add and configure a Cisco IDS or IPS device in MARS, follow these steps:

-
- Step 1** Click **Admin > System Setup > Security and Monitor Devices > Add**.
- Step 2** Do one of the following:

- Select **Cisco IDS 4.0** from the Device Type list.

Figure 6-6 *Configure Cisco IDS 4.0*

Device Type: Cisco IDS 4.0

- Select **Cisco IPS 5.x** from the Device Type list.

Figure 6-7 *Configure Cisco IPS 5.x*

Device Type: Cisco IPS 5.x

- Step 3** Enter the hostname of the sensor in the Device Name field.
The Device Name value must be identical to the configured sensor name.
- Step 4** Enter the administrative IP address in the Access IP field.

- Step 5** Enter the administrative IP address in the Reporting IP field.
The Reporting IP address is the same address as the administrative IP address.
- Step 6** In the Login field, enter the username associated with the administrative account that will be used to access the reporting device.
- Step 7** In the Password field, enter the password associated with the username specified in the Login field.
- Step 8** In the Port field, enter the TCP port on which the webserver running on the sensor listens. The default HTTPS port is 443.



Note While it is possible to configure HTTP only, MARS requires HTTPS.

- Step 9** To pull the IP logs from the sensor, select **Yes** in the Pull IP Logs box.
- Step 10** For attack path calculation and mitigation, specify the networks being monitored by the sensor. Do one of the following:
To manually define the networks, select the **Define a Network** radio button.
- Enter the network address in the Network IP field.
 - Enter the corresponding network mask value in the Mask field.
 - Click **Add** to move the specified network into the Monitored Networks field.
 - Repeat as needed.
- To select the networks that are attached to the device, click the **Select a Network** radio button.
- Select a network from in the Select a Network list.
 - Click **Add** to move the selected network into the Monitored Networks field.
 - Repeat as needed.
- Step 11** To verify the configuration, click **Test Connectivity**.
- Step 12** Click **Submit**.

Specify the Monitored Networks for Cisco IPS or IDS Device Imported from a Seed File

After you import a Cisco IPS or IDS device into MARS using a seed file, you must define the networks that are monitored by that sensor.

To define the networks monitored by a sensor, follow these steps:

- Step 1** Click **Admin > System Setup > Security and Monitor Devices**.
- Step 2** Select the check box next to the Cisco IPS or IDS device that was imported using a seed file. and click **Edit**.
- Step 3** To specify the networks being monitored by the sensor, do one of the following:
To manually define the networks, select the **Define a Network** radio button.
- Enter the network address in the Network IP field.
 - Enter the corresponding network mask value in the Mask field.

- c. Click **Add** to move the specified network into the Monitored Networks field.
- d. Repeat as needed.

To select the networks that are attached to the device, click the **Select a Network** radio button.

- a. Select a network from in the Select a Network list.
- b. Click **Add** to move the selected network into the Monitored Networks field.
- c. Repeat as needed.

Step 4 To save your changes, click **Submit**.

Step 5 To enable MARS to start sessionizing events from this module, click **Activate**.

View Detailed Event Data for Cisco IPS Devices

You can view the trigger packets and IP log data associated with incidents reported by Cisco IDS 4.x and Cisco IPS 5.x devices, whether they are sensor appliances or modules. This information is useful when an in-depth understanding of the attack method is desired. MARS includes two event types that focus on the these two data types:

- **Trigger packet data.** Identifies the data that was being transmitted on the network the instant an alarm was detected. You can use this information to help diagnose the nature of an attack. The trigger packet provides a single data packet—the data packet that caused the alarm to fire.
- **Packet data.** Identifies the data that was being transmitted on the network the instant an alarm was detected. You can use this information to help diagnose the nature of an attack. Although the amount of data contained in an IP log varies based on sensor configuration, by default an IP log contains 30 seconds of packet data. To view this data, you must enable the Pull IP Logs option on the Cisco IPS device under Admin > System Setup > Security and Monitor Devices.



Note

MARS does not collect this data for Cisco IDS 3.x devices.

For the correct signature settings required to generate this data, see [Enable the Correct Signatures and Actions, page 6-6](#).

If the IP log feature is enable for the reporting Cisco IPS device, these event types are combined as part of the incident data. You can view this data by drilling down in an incident, expanding the desired event type (either Packet Data or Trigger Packet Data), selecting an event, and clicking on the RAW Events for this Session icon under the Reporting Device column of that event. The source, destination, and other data displayed for these events matches that of the original alert. In addition, this data appears hexadecimal and binary format.



Note

The trigger packet and IP log data is stored using a base64-encoded format in the MARS database. Therefore, keyword search does not work on it if you just provide the search string.

Cisco IPS Modules

MARS can monitor Cisco IPS modules installed in Cisco switches and Cisco ASA appliances. To prepare these modules, you must perform the following tasks:

- Define the base module, either the router, switch, or Cisco ASA, as defined in [Cisco Router Devices, page 3-1](#), [Cisco Switch Devices, page 3-9](#), and [Cisco Firewall Devices \(PIX, ASA, and FWSM\), page 4-1](#).
- Bootstrap the base module to enable SDEE traffic on the Cisco IPS module, to forward events to the MARS Appliance, and to enable MARS to access the SDEE events stored on the modules. Module access enables MARS to retrieve trigger packets and IP log information.
- Add the IPS feature set to the base module previously defined in the web interface.

This section contains the following topics:

- [Enable DTM Support, page 6-10](#)
- [Enable SDEE on the Cisco IOS Device with an IPS Module, page 6-10](#)
- [Add an IPS Module to a Cisco Switch or Cisco ASA, page 6-11](#)

Enable DTM Support

To support DTM, you must configure your IPS module as follows:

- Purchase or enable the IOS IPS feature set.
- Enable HTTPS for SDEE.
- Enable SSH to discover settings, which is the method recommended over Telnet.

Enable SDEE on the Cisco IOS Device with an IPS Module

In addition to enabling either Telnet or SSH for configuration discovery on a Cisco IOS device, you must also enable SDEE on the device that supports IPS module. SDEE is used to publish events to MARS about signatures that have fired.

To enable SDEE protocol on the Cisco IOS device that supports IPS module, perform the following steps:

-
- Step 1** Log in to the Cisco IOS device using the enable password.
- Step 2** Enter the following commands to enable MARS to retrieve the events from the IPS module:

```
Router(config)#ip http secure-server
Router(config)#ip ips notify sdee
Router(config)#ip sdee subscriptions 3
Router(config)#ip sdee events 1000
Router(config)#no ip ips notify log
```

**Note**

The “no ips notify log” causes the IPS modules to stop sending IPS events over syslog.

Add an IPS Module to a Cisco Switch or Cisco ASA

You can enable in-line IPS functionality and signature detection in multi-purpose Cisco platforms. You can identify an IDS-M2 running in a Cisco Switch or an ASA-SSM running in a Cisco ASA. To represent either of these modules, you must define the settings for the module as part of the base platform, which must be previously defined under Admin > System Setup > Security and Monitor Devices.

To add an IPS module to a Cisco Switch or Cisco ASA, follow these steps:

-
- Step 1** Click **Admin > System Setup > Security and Monitor Devices**.
- Step 2** From the list of devices, select the Cisco switch or Cisco ASA to which you want to add the IPS module and click **Edit**.
- Step 3** Click **Add Module**.

Device Type: Cisco ASA 7.0
Cisco ASA 7.0
Cisco IPS 5.x

→ *Device Name:	
→ *Context Name:	
→ *Reporting IP:	
SNMP RO Community:	

143172

- Step 4** Select **Cisco IPS 5.x** in the Device Type list.

For Cisco switches, you can also add a Cisco IPS 4.0 module or an IDS 3.1 module. You configure these modules just as you would a standalone sensor. For instructions on configuring these modules, refer to [Cisco IDS 3.1 Sensors, page 6-1](#) and [Cisco IDS 4.0 and IPS 5.x Sensors, page 6-5](#).

Figure 6-8 *Configure Cisco IPS 5.x*

Device Type: Cisco IPS 5.x

→ ***Device Name:**

→ **Reporting IP:**

→ ***Access Type:** **SSL**

Login:

Password:

Port:

→ **Monitor Resource Usage:** NO

Pull IP Logs: NO

[Optional: for attack path calculation and mitigation enter monitoring networks information]

→ **Monitored Networks:**

☐ **Select a Network:**

10.0.0.0/255.0.0.0(n-10.0.0.0/8)

☐ **Define a Network:**

Network IP:

Mask:

- Step 5** Enter the hostname of the sensor in the Device Name field.
- Step 6** Enter the administrative IP address in the Reporting IP field.
- Step 7** The Reporting IP address is the same address as the administrative IP address.
- Step 8** In the Login field, enter the username associated with the administrative account that will be used to access the reporting device.
- Step 9** In the Password field, enter the password associated with the username specified in the Login field.
- Step 10** In the Port field, enter the TCP port on which the webserver running on the sensor listens. The default HTTPS port is 443.

**Note**

While it is possible to configure HTTP only, MARS requires HTTPS.

- Step 11** For attack path calculation and mitigation, specify the networks being monitored by the sensor. Do one of the following:
- To manually define the networks, select the Define a Network radio button.
- a. Enter the network address in the Network IP field.
 - b. Enter the corresponding network mask value in the Mask field.
 - c. Click **Add** to move the specified network into the Monitored Networks field.

d. Repeat as needed.

To select the networks that are attached to the device, click the Select a Network radio button.

a. Select a network from in the Select a Network list.

b. Click **Add** to move the selected network into the Monitored Networks field.

c. Repeat as needed.

Step 12 Click **Test Connectivity** to verify the configuration.

Step 13 To save your changes, click **Submit**.

Step 14 To enable MARS to start sessionizing events from this module, click **Activate**.

ISS Site Protector



Note

This topic describes how to use Site Protector to configure the ISS NIDS and HIDS; Site Protector is not a device type that can be monitored or used as an aggregation point for ISS event data from the perspective of MARS. MARS cannot parse event data from Site Protector, unless you develop a custom event parser for each event type as described in [Adding User Defined Log Parser Templates, page 15-1](#).

MARS supports ISS NIDS and HIDS event retrieval via SNMP. However, when configuring ISS RealSecure sensors (NIDS) and hosts (HIDS), you must configure each active signature to send an alert to the MARS Appliance. This task can be very tedious as it must be done for each sensor and after each signature upgrade, as it resets the redirect configuration. One approach to simplifying this task is to use the ISS Site Protector management console to define these changes globally and apply them to each sensor.

ISS Site Protector 2.0 allows you to centrally manage SNMP alert destinations, such as the MARS Appliance, for group policies. You can then push these group policies to all desired host and network sensors. For each ISS signature update, you must specify the MARS Appliance as an SNMP alert destination before you apply the downloaded signatures to sensors using Site Protector.



Note


By default, the group policy response configuration is supported only on Proventia G400 and G2000 models. For all other models, including the G100 mentioned, a firmware upgrade is required. See the documentation that came with ISS Site Protector for more information.

To perform the major configuration steps required to use Site Protector to forward the SNMP alerts generated by sensors to MARS Appliance, follow these steps:


Step 1 Using the Add Sensor Wizard, register the sensor to Site Protector Console.

Other methods exist for registering sensors in Site Protector. For more information on using the Wizard as well as these other methods, see Chapter 9, Registering Software Managed by SiteProtector, on page 105 at the following URL:


<http://documents.iss.net/literature/SiteProtector/SPUserGuideforSecurityManagers20SP52.pdf>



Step 2 Right-click the sensor to edit, and click **Edit Settings** on the shortcut menu.




The Edit Settings dialog appears.




Step 3 Create a new SNMP response that sends messages to the IP address of the MARS Appliance:

- Select **Response Objects** from the settings tree.
- Select the **SNMP** tab.
- Click **Add** to create a new SNMP response object using the IP address of the MARS Appliance.



Step 4 Select the Security Events to configure new SNMP destination.



- a. Select **Security Events** under the sensor folder.
- b. Select the required security events from the Security Events tab.


The Group By button allows you to group policies using any number of parameters.



Note You can also select policies and edit them at the group level.

- c. Click **Edit** to configure SNMP response of all the selected policies.

Step 5 Select the MARS Appliance on SNMP tab.



- a. Enable all the security events by selecting the **Enabled** checkbox located at the top of the Edit Security Events dialog box.
- a. Select the **SNMP** tab under Responses, and then select the **Enabled** checkbox next to the name of MARS Appliance created in [Step 3](#).
- a. Click **OK**.

The security events and updated response target are applied to the selected sensor during the next synchronization.

ISS RealSecure 6.5 and 7.0

To configure ISS RealSecure, you must perform the following four tasks:

1. Prepare each ISS sensor as follows:
 - Edit the `common.policy` files to point to the MARS Appliance as an SNMP target.
 - Modify the `current.policy` files to configure each signature so that the SNMP notification is a default response when triggered.
 - Edit the `response.policy` files to specify the IP of the SNMP manager (MARS Appliance) and the community string.

- Restart the ISS daemon for the changes to take effect.

For more information, see [Configure ISS RealSecure to Send SNMP Traps to MARS, page 6-18](#).

2. Add the ISS sensor to MARS as a network-based IDS device. For more information, see [Add an ISS RealSecure Device as a NIDS, page 6-19](#).
3. Click **Activate** to enable proper processing of received events.

Configure ISS RealSecure to Send SNMP Traps to MARS

To configure an ISS RealSecure sensor, follow these steps:

Step 1 Log into the sensor.

Step 2 Locate the `common.policy` files in these directories:

```
Microsoft Windows
Program Files\ISS\issSensors\server_sensor_1
Program Files\ISS\issSensors\network_sensor_1

Linux
/opt/ISS/issSensors/server_sensor_1
/opt/ISS/issSensors/network_sensor_1
```

Step 3 Open the `common.policy` files in a text editor.

Step 4 Change the line that reads:

```
Manager =S
```

to:

```
Manager =S <MARS's IP address>
```

If MARS Appliance's IP address is NATed, you may need to use the NATed address. If you use the MARS Appliance's IP address as the destination IP address, make sure the SNMP trap can reach MARS Appliance.

Step 5 Save these edited files and exit the editor.

Step 6 Locate the `current.policy` files in these directories:

```
Microsoft Windows
Program Files\ISS\issSensors\server_sensor_1
Program Files\ISS\issSensors\network_sensor_1

Linux
/opt/ISS/issSensors/server_sensor_1
/opt/ISS/issSensors/network_sensor_1
```

Step 7 Open the `current.policy` files in a text editor.

Edit each signature to have SNMP as one of its responses, and set the choice for SNMP trap as default. For example, in this original signature:

```
[\template\features\AOLIM_File_Xfer\Response\];
[\template\features\AOLIM_File_Xfer\Response\DISPLAY\];
Choice =S Default;
[\template\features\AOLIM_File_Xfer\Response\LOGDB\];
Choice =S LogWithoutRaw;
```

Insert the following bolded lines to make it look similar to the following:

```
[\template\features\AOLIM_File_Xfer\Response\];
[\template\features\AOLIM_File_Xfer\Response\DISPLAY\];
Choice =S Default;
[\template\features\AOLIM_File_Xfer\Response\SNMP\];
Choice =S Default;
[\template\features\AOLIM_File_Xfer\Response\LOGDB\];
Choice =S LogWithoutRaw;
```

Step 8 Save these edited files and exit the editor.

Step 9 Locate the `response.policy` files in these directories:

```
Microsoft Windows
Program Files\ISS\RealSecure SiteProtector\Console
```

```
Linux
/opt/ISS/RealSecure SiteProtector/Console
```

Step 10 Edit the `response.policy` files to specify the IP of the SNMP manager (MARS Appliance) and the community string:

```
SMTP_HOST =S ;
addr_1 =S ;
[\Response\SNMP\];
[\Response\SNMP\Default\];
Manager =S ;
Community =S public;
```

to:

```
Manager =S <MARS's IP address> ;
Community = S <string> public;
```

If MARS Appliance's IP address is NATed, you may need to use the NATed address. If you use the MARS Appliance's IP address as the destination IP address, make sure the SNMP trap can reach MARS Appliance.

Step 11 Save these edited files and exit the editor.

Step 12 Restart the ISS daemon.

- For sensors installed on Microsoft Windows, restart it in the Services menu.
- For sensors installed on Linux, run:

```
/etc/init.d/RealSecure stop
/etc/init.d/RealSecure start
```

Add an ISS RealSecure Device as a NIDS

Step 1 Click **Admin > System Setup > Security and Monitor Devices > Add**.

Step 2 From the **Device Type** list, select **Add SW Security apps on a new host** or **Add SW security apps on existing host**.

Step 3 Enter the **Device Name**.

Step 4 Click **Apply**.

Step 5 Click on **Reporting Applications** tab.

- Step 6** From the **Select Application** list, select **RealSecure (6.5 or 7.0)**.
- Step 7** Click **Add**.
- Step 8** Click the **NIDS** radio button, if it is not already selected.

Figure 6-9 *Configure ISS Real Secure NIDS*

- Step 9** For attack path calculation and mitigation, specify the networks being monitored by the sensor. Do one of the following:
- To manually define the networks, select the **Define a Network** radio button.
- Enter the network address in the Network IP field.
 - Enter the corresponding network mask value in the Mask field.
 - Click **Add** to move the specified network into the Monitored Networks field.
 - Repeat as needed.
- To select the networks that are attached to the device, click the **Select a Network** radio button.
- Select a network from in the Select a Network list.
 - Click **Add** to move the selected network into the Monitored Networks field.
 - Repeat as needed.
- Step 10** To save your changes, click **Submit**.
- Step 11** To enable MARS to start sessionizing events from this module, click **Activate**.

Add an ISS RealSecure Device as a HIDS

- Step 1** Click **Admin > System Setup > Security and Monitor Devices > Add**.
- Step 2** From the **Device Type** list, select **Add SW Security apps on a new host** or **Add SW security apps on existing host**.
- Step 3** Enter the **Device Name**.

- Step 4** Click **Apply**.
- Step 5** Click on **Reporting Applications** tab.
- Step 6** From the **Select Application** list, select **RealSecure (6.5 or 7.0)**.
- Step 7** Click **Add**.
- Step 8** Click the **HIDS** radio button.

Figure 6-10 *Configure ISS Real Secure HIDS*

→ ☐ NIDS ☒ HIDS

To add HIDS RealSecure, select the radio button and submit, then add interfaces in the General Tab.

Cancel Submit

- Step 9** Click **Submit**.
- Step 10** For multiple interfaces, click on **General Tab**, and add the new interfaces' name, IP address, and network mask.

Figure 6-11 *Adding Multiple Interfaces*

Device Type: Edit host with security applications

General Reporting Applications Vulnerability Assessment Info

→ *Device Name: H-10.1.1.103

→ Access IP: [][][][]

→ Reporting IP: [][][][]

→ Operating System: Generic [v] Logging Info

Enter interface information:

Add Interface Remove Interface/IP

Name: IP Address: Network Mask:

☐ eth0 10 1 1 103 255 255 0 0 Add IP/Network Mask

- Step 11** Click **Apply**.

IntruVert IntruShield

To configure IntruVert IntruShield in MARS, you must perform the following tasks:

1. Generate CSV file that identifies each of the IntruShield sensor hosts by logging into the database to which IntruShield Manager writes and performing and saving a database query.
2. Configure the IntruShield Manager to send SNMP traps to the MARS Appliance
3. Define a host that represents the management console (IntruVert Manger) in MARS web interface.
4. From that host in the MARS web interface, import the IntruShield sensor seed file to identify the IntruVert sensors running on other hosts.

The following sections provide details on performing each of these tasks:

- [Extracting Intruvert Sensor Information from the IntruShield Manager, page 6-22](#)
- [Configure IntruShield Version 1.5 to Send SNMP traps to MARS, page 6-23](#)
- [Configure IntruShield Version 1.8 to Send SNMP Traps to MARS, page 6-23](#)
- [Add and Configure an IntruShield Manager and its Sensors in MARS, page 6-25](#)

Extracting Intruvert Sensor Information from the IntruShield Manager

IntruVert sensor information is saved in a database on the IntruShield Manager host. When you configure the MARS to add Intruvert sensors, you can manually add the mapping of each Intruvert sensor name or you can extract them as a seed file from the database on the Intruvert Manager.



Note

The instructions apply for Intruvert IntruShield version 1.5. IntruVert supports both MySQL and Oracle.

To create a CSV file for IntruVert IntruShield 1.5, follow these steps:

- Step 1** Log in to the database.
- Step 2** Perform the query:

```
use lf; select name, ip_address from iv_sensor where ip_address is not NULL;
```
- Step 3** Store the query result into a file, remove the header, trailer, and separator lines, and edit the result to a CSV format.

For example, the query result could be:

```
+-----+-----+
| name      | ip_address |
+-----+-----+
| intruvert | 0A010134  |
| intruvert1 | 0A010135  |
+-----+-----+
```

2 row in set (0.00 sec)

You would then edit the above file to appear as:

```
intruvert,0A010134
intruvert1,0A010135
```

- Step 4** Save the edited CSV file, move the file to an FTP server from which you can load the seed file using the MARS web interface.
-

Configure IntruShield Version 1.5 to Send SNMP traps to MARS

- Step 1** Log in to the IntruShield Manager version 1.5.
- Step 2** Click **Configure**.
- Step 3** In the Resource Tree, click **My Company**.
- Step 4** Click the **Forwarding** tab.
- Step 5** In the **Add SNMP Server** field, enter:
- a. **Target Server IP Address:** Enter MARS's IP address as it appears to IntruShield.
 - b. **Target Server Port Number:** Enter MARS's port number 162.
 - c. **SNMP Version:** 1
 - d. Check the **Forward Alerts** box.
 - e. Select the **For this and child admin domains** radio button.
 - f. Select the severity from the list. Cisco recommends selecting **High and Medium** severity.
 - g. Check the **Forward Faults** box.
 - h. Select the severity from the list. Cisco recommends selecting **Error and above** severity.
- Step 6** Click **Save** and exit the program.
-

Configure IntruShield Version 1.8 to Send SNMP Traps to MARS

- Step 1** Log in to the IntruShield Manager version 1.8.
- Step 2** Click **Configure**.
- Step 3** In the Resource Tree, click **My Company**.
- Step 4** Click the **Alert Notification** tab.
- Step 5** Click the **SNMP Forwarder** sub-tab.

Figure 6-12 IntruShield SNMP Forwarder Configuration

The screenshot shows the IntruShield Manager interface. On the left is a tree view with 'My Company' expanded, showing 'Manager', 'Policies', and 'Sensors'. Under 'Sensors', 'intruvert' is expanded, showing sensors 1A, 1B, 2A-2B, 3A, 3B, 4A, and 4B. The main window is titled 'SYSTEM CONFIGURATION' and has tabs for 'Admin Domain', 'Alert Notification', 'Fault Notification', 'User', and 'Log'. The 'Alert Notification' tab is active, and within it, the 'SNMP Forwarder' sub-tab is selected. Below the tabs are buttons for 'View Details', 'SNMP Forwarder', 'Syslog Forwarder', 'Email', 'Pager', and 'Script'. The main area contains a table with the following data:

	Target Server IP Address	Target Server Port Number	SNMP Version	Alert Forwarding
☐	10.1.1.132	162	1	For this and children admin domain for ALL alerts.
☐	10.1.1.12	162	1	For this admin domain only for ALL alerts.
☐	10.1.2.4	162	1	For this admin domain only for ALL alerts.
☐	10.1.1.134	162	1	For this admin domain only for ALL alerts.
☐	10.1.4.102	162	1	For this admin domain only for ALL alerts.
☐	10.1.1.55	162	1	For this and children admin domain for ALL alerts.
☐	10.1.1.114	162	1	For this and children admin domain for ALL alerts.
☐	10.2.3.42	162	1	For this and children admin domain for ALL alerts.
☒	10.1.1.212	162	1	For this and children admin domain for ALL alerts.

At the bottom of the table are three buttons: 'ADD', 'EDIT', and 'DELETE'.

Step 6 Click the **Add** button.

Figure 6-13 IntruShield Target SNMP Server

143363

- Step 7** On the SNMP Forwarder page, enter:
- Enable SNMP Forwarder:** Select the **Yes** radio button.
 - Target Server (IP Address):** Enter MARS's IP address as it appears to IntruShield.
 - Target Server Port Number:** Enter MARS's port number 162.
 - SNMP Version:** 1
 - Forward Alerts
 - Select the severity from the list. Cisco recommends selecting **Informational and above** severity.
 - Customize Community:** Enter the community string that you want to use.
- Step 8** Click **Apply** and exit the program.

Add and Configure an IntruShield Manager and its Sensors in MARS

Adding an IntruVert device has two distinct steps. First, you add configuration information for the for the IntruShield Manager host. Second, you add the sensors managed by that host.


- [Add the IntruShield Manager Host to MARS, page 6-26](#)
- [Add IntruShield Sensors Manually, page 6-26](#)
- [Add IntruShield Sensors Using a Seed File, page 6-27](#)

Add the IntruShield Manager Host to MARS

To define the host and represent the management console for IntruShield, follow these steps:

-
- Step 1** Click **Admin > System Setup > Security and Monitor Devices > Add**.
 - Step 2** From the **Device Type** list, select **Add SW Security apps on a new host** or **Add SW security apps on existing host**.
 - Step 3** Enter the **Device Name** and **IP addresses** if adding a new host.
 - Step 4** Click **Apply**.
 - Step 5** Click **Reporting Applications** tab.
 - Step 6** Select **IntruVert IntruShield 1.5** from the Select Application list.
 - Step 7** To complete the definition of this console, click **Add**.

Figure 6-14 Add IntruShield Sensors



- Step 8** Continue defining the sensors that the console manages using one of two methods:
 - [Add IntruShield Sensors Manually, page 6-26](#)
 - [Add IntruShield Sensors Using a Seed File, page 6-27](#)
-

Add IntruShield Sensors Manually

To add sensors manually, follow these steps:

-
- Step 1** Click **Add Sensor**.
 - Step 2** Enter the **Device Name**, **Sensor Name**, and its **Reporting IP** address.
 - **Device Name** – the DNS entry for this device
 - **Sensor Name** – the name as it appears in the console
 - **Reporting IP** – the IP address that the agent uses to send logs to the console
 - Step 3** Add the interface information.
 - Step 4** For attack path calculation and mitigation, specify the networks being monitored by the sensor. Do one of the following:

To manually define the networks, select the **Define a Network** radio button.

 - a. Enter the network address in the Network IP field.

- b. Enter the corresponding network mask value in the Mask field.
- c. Click **Add** to move the specified network into the Monitored Networks field.
- d. Repeat as needed.

To select the networks that are attached to the device, click the **Select a Network** radio button.

- a. Select a network from in the Select a Network list.
- b. Click **Add** to move the selected network into the Monitored Networks field.
- c. Repeat as needed.

Step 5 To save your changes, click **Submit**.

Step 6 To enable MARS to start sessionizing events from this module, click **Activate**.

Add IntruShield Sensors Using a Seed File

To add sensors using a seed file, follow these steps:

Step 1 Click **Load From CSV**.

Step 2 Enter the FTP server information and location of the CSV (comma separated values) file.

- If you need to generate the IntruShield sensors CSV file, [Extracting Intruvert Sensor Information from the IntruShield Manager, page 6-22](#).

Step 3 Click **Submit**.

The list of sensors appears on the management console page.

Step 4 For each sensor that appears in the management console page, select the check box next to the sensor and click **Edit Sensor**.

Step 5 For attack path calculation and mitigation, specify the networks being monitored by the sensor. Do one of the following:

To manually define the networks, select the **Define a Network** radio button.

- a. Enter the network address in the Network IP field.
- b. Enter the corresponding network mask value in the Mask field.
- c. Click **Add** to move the specified network into the Monitored Networks field.
- d. Repeat as needed.

To select the networks that are attached to the device, click the **Select a Network** radio button.

- a. Select a network from in the Select a Network list.
- b. Click **Add** to move the selected network into the Monitored Networks field.
- c. Repeat as needed.

Step 6 To save your changes, click **Submit**.

Step 7 To save the changes made to this management console and the sensors it manages, click **Submit**.

Step 8 To enable MARS to start sessionizing events from this module, click **Activate**.

Snort 2.0

Configure Snort to Send Syslogs to MARS

For Snort, use the syslog as your output plugin. Configure your syslogd to send copies to another host. On most older-style systems (Solaris/Linux), you need to edit `/etc/syslog.conf`. (Assuming that the system is based on syslogd, and not any of the newer system logging facilities. The newer logging facilities are not supported by Snort.)

To configure Snort to send syslog messages to the MARS Appliance, follow these steps:

-
- Step 1** Make Snort's output go to syslog with log facility local4 in `snort.conf` (you can pick any local facility that's unused.)
- ```
output alert_syslog: LOG_LOCAL4 LOG_ALERT
```
- `snort.conf` is normally in `/etc/snort`.
- Step 2** Add a redirector in your `/etc/syslog.conf` on your Snort box to send syslog to MARS.
- ```
local4.alert @IPAddrOffMarsbox
```
- Step 3** Restart the Snort daemon and the syslogd daemon on your Snort box.
-

Add the Snort Device to MARS

To add the Snort device to MARS, follow these steps:

-
- Step 1** Click **Admin > System Setup > Security and Monitor Devices > Add**
- Step 2** From the **Device Type** list, select **Add SW Security apps on a new host** or **Add SW security apps on existing host**
- Step 3** Enter the **Device Name** and **IP addresses** if adding a new host.
- Step 4** Click **Apply**
- Step 5** Click **Reporting Applications** tab
- Step 6** From the **Select Application** list, select **Snort Snort 2.0**
- Step 7** Click **Add**
- Step 8** For attack path calculation and mitigation, specify the networks being monitored by the sensor. Do one of the following:
- To manually define the networks, select the **Define a Network** radio button.
- Enter the network address in the Network IP field.
 - Enter the corresponding network mask value in the Mask field.
 - Click **Add** to move the specified network into the Monitored Networks field.
 - Repeat as needed.

To select the networks that are attached to the device, click the **Select a Network** radio button.

- a. Select a network from in the Select a Network list.
- b. Click **Add** to move the selected network into the Monitored Networks field.
- c. Repeat as needed.

Step 9 To save your changes, click **Submit**.

Step 10 To enable MARS to start sessionizing events from this module, click **Activate**.


Symantec ManHunt

Symantec ManHunt Side Configuration

Step 1 Login to the Symantec ManHunt with appropriate username and password.


Step 2 In the main screen, click **Setup > Policy > Response Rules**, then Response Rules window will appear.

Figure 6-15 ManHunt Configuration



Step 3 In the Response Rules window, click **Action > Add response Rules**.

Step 4 Click in the field of **Response Action**

Figure 6-16 ManHunt Response Rule Config

Step 5 In the left menu, click SNMP Notification and enter the following information:

- a. **SNMP Manager IP address:** Reporting IP address of MARS
- b. **Maximum number of SNMP notification:** (Example: 100000).
- c. **Delay between SNMP notification (mins):** (Example: 1 min)

Step 6 Click **OK** to return to main screen.

MARS Side Configuration

Add Configuration Information for Symantec ManHunt 3.x

- Step 1** Click **Admin > System Setup > Security and Monitor Devices > Add**
- Step 2** From the **Device Type** list, select **Add SW Security apps on a new host** or **Add SW security apps on existing host**
- Step 3** Enter the **Device Name** and **IP addresses** if adding a new host.
- Step 4** Click **Apply**
- Step 5** Click **Reporting Applications** tab

- Step 6** From the **Select Application** list, select **Symantec ManHunt 3.x**
- Step 7** Click **Add**
- Step 8** For attack path calculation and mitigation, specify the networks being monitored by the sensor. Do one of the following:
- To manually define the networks, select the **Define a Network** radio button.
- Enter the network address in the Network IP field.
 - Enter the corresponding network mask value in the Mask field.
 - Click **Add** to move the specified network into the Monitored Networks field.
 - Repeat as needed.
- To select the networks that are attached to the device, click the **Select a Network** radio button.
- Select a network from in the Select a Network list.
 - Click **Add** to move the selected network into the Monitored Networks field.
 - Repeat as needed.
- Step 9** To save your changes, click **Submit**.
- Step 10** To enable MARS to start sessionizing events from this module, click **Activate**.
-

NetScreen IDP 2.1

IDP-side Configuration

-
- Step 1** Click **NetScreen-Global Pro > IDP Manager > IDP**.
- Step 2** Log in to the IDP Manager.
- Step 3** From the main menu, click **Tools > Preferences**.
- Step 4** In the tree on the left, click **Management Server**, enter the Local Controller's address in the Syslog host field, and click **OK**.
- Step 5** Click **Security Policies**, and the name of your policy.
- Step 6** In the **Notification** column, right-click anywhere in the cell in the field and select **Configure**.
- Step 7** Check **enable logging** and **syslog** for each policy, and click **OK**. Repeat for all of your policies.
- Step 8** From the main menu, click **Policy > Install**.
-


MARS-side Configuration

Add Configuration Information for the IDP

-
- Step 1** Click **Admin > System Setup > Security and Monitor Devices > Add**.

- Step 2** From the **Device Type** list, select **Add SW Security apps on a new host** or **Add SW security apps on existing host**.
- Step 3** Enter the **Device Name** and **IP Addresses** if adding a new host.
- Step 4** Click **Apply**.
- Step 5** Click **Reporting Applications** tab.
- Step 6** From the **Select Application** list, select **NetScreen IDP 2.1**.
- Step 7** Click **Add**.

Figure 6-17 Add NetScreen IDP 2.1 Sensors



Add NetScreen IDP 2.1 Sensors Manually

- Step 1** Click **Add Sensor**.
- Step 2** Select existing device or **Add New** device.
- Step 3** Enter the **Device Name**, **Sensor Name**, and its **Reporting IP** address.
- **Device Name** – the DNS entry for this device
 - **Sensor Name** – the name as it appears in the console
 - **Reporting IP** – the IP address that the agent uses to send logs to the console
- Step 4** Add the interfaces, which important information for attack path calculation.
- For multiple interfaces, click **Add Interface**, and add the new interfaces's name, IP address and mask.
- Step 5** For attack path calculation and mitigation, specify the networks being monitored by the sensor. Do one of the following:
- To manually define the networks, select the **Define a Network** radio button.
- a. Enter the network address in the Network IP field.
 - b. Enter the corresponding network mask value in the Mask field.
 - c. Click **Add** to move the specified network into the Monitored Networks field.
 - d. Repeat as needed.
- To select the networks that are attached to the device, click the **Select a Network** radio button.
- a. Select a network from in the Select a Network list.
 - b. Click **Add** to move the selected network into the Monitored Networks field.

c. Repeat as needed.

Step 6 To save your changes, click **Submit**.

Step 7 To enable MARS to start sessionizing events from this module, click **Activate**.

Enterasys Dragon 6.x

To configure the Enterasys Dragon devices, you must:

- Configure the Dragon Policy Manager (DPM) or Event Flow Processor (EFP).
- Configure the syslog daemon running on the same system as the DPM or EFP.
- Configure the MARS.

DPM/EFP Configuration

Before you configure the DPM or EFP, you must install and enable the Alarmtool.

Configure the DPM or EFP

Step 1 Log into the DPM or EFP.

Step 2 Click **Alarmtool**.

Step 3 In the left menu, click **Notification Rules**.

Step 4 In the right window, select syslog if it exists. If not, you need to create it:

- Click **New Notification Rules** and select **syslog**.
- Facility** - Make sure the **localn** you select is not in use by the syslog daemon
- Level** - Select **Debug**
- Message** - Make sure its in such format:

```
%TIME% %DATE% SigName=%NAME% from Sensor=%SENSOR%
SrcIP=%SIP% DstIP=%DIP% SrcPort=%SPORT% DstPort=%DPORT%
Protocol=%PROTO%
```

Step 5 Click **Save**.

Step 6 In the left menu, click **Alarm**.

Step 7 Set the **Type** to **Real-time** and the **Notification Rule** to **syslog**.

Step 8 Click **Save**.

Step 9 In the left menu, click **Deployment**.

Step 10 In the main screen, click **View Configuration**. Make sure the **localn** set in both notify syslog and alarm syslog match.

Step 11 In the main screen, click **Deploy and Reset** to confirm the configuration change.

Host-side Configuration

Configure the syslog on the UNIX host

-
- Step 1** Log into the host as the root user.
- Step 2** On the same system running the DPM or EFP, edit the file `/etc/syslog.conf`.
- Step 3** Make sure `n` in `localn` matches the syslog entry you used on the DPM or EFP.
- Step 4** Restart the syslog daemon by entering:
- ```
/etc/init.d/syslog restart
```
- 

## MARS-side Configuration

### Add Configuration Information for the Enterasys Dragon

- 
- Step 1** Click **Admin > System Setup > Security and Monitor Devices > Add**.
- Step 2** From the **Device Type** list, select **Add SW Security apps on a new host** or **Add SW security apps on existing host**
- Step 3** Enter the **Device Name** and **IP Addresses** if adding a new host.
- Step 4** Click **Apply**
- Step 5** Click **Reporting Applications** tab
- Step 6** From the **Select Application** list, select **Enterasys Dragon 6.x**
- Step 7** Click **Add**.
- 

### Add a Dragon NIDS Device

- 
- Step 1** Click **Add Sensor**.
- Step 2** Select existing device or **Add New** device.
- Step 3** Enter the **Device Name**, **Sensor Name**, and its **Reporting IP** address.
- **Device Name** – the DNS entry for this device
  - **Sensor Name** – the name as it appears in the console
  - **Reporting IP** – the IP address that the agent uses to send logs to the console
- Step 4** Add the interfaces, which important information for attack path calculation.
- For multiple interfaces, click **Add Interface**, and add the new interfaces's name, IP address and mask.
- Step 5** For attack path calculation and mitigation, specify the networks being monitored by the sensor. Do one of the following:

To manually define the networks, select the **Define a Network** radio button.

- a. Enter the network address in the Network IP field.
- b. Enter the corresponding network mask value in the Mask field.
- c. Click **Add** to move the specified network into the Monitored Networks field.
- d. Repeat as needed.

To select the networks that are attached to the device, click the **Select a Network** radio button.

- a. Select a network from in the Select a Network list.
- b. Click **Add** to move the selected network into the Monitored Networks field.
- c. Repeat as needed.

**Step 6** To save your changes, click **Submit**.

**Step 7** Click **Done** when you are done adding the sensor.

**Step 8** To enable MARS to start sessionizing events from this module, click **Activate**.

---







## Configuring Host-Based IDS and IPS Devices

Host-based intrusion detection and prevention devices provide MARS with detailed information about attacks seen at the host level, rather than the network level. They also provide information about the host operating system and successful prevention of attacks, both of which provide more targeted data for false positive analysis.

This chapter explains how to bootstrap and add the following host-based IDS and IPS devices to MARS:

- [Entercept Entercept 2.5 and 4.0, page 7-1](#)
- [Cisco Security Agent 4.x Device, page 7-5](#)

### Entercept Entercept 2.5 and 4.0

To configure Entercept in MARS, you must perform the following tasks:

1. Generate CSV file that identifies each of the Entercept hosts by logging into the host running the Entercept console and copying the data out of the database table.
2. Configure the Entercept console to send SNMP traps to the MARS Appliance
3. Identify the events that should be generated as SNMP traps.
4. Define a host that represents the management console (Entercept console) in MARS web interface.
5. From that host in the MARS web interface, import the CSV seed file to identify the Entercept agents running on other hosts.

The following sections provide details on performing each of these tasks:

- [Extracting Entercept Agent Information into a CSV file \(for Entercept Version 2.5\), page 7-1](#)
- [Define the MARS Appliance as an SNMP Trap Target, page 7-2](#)
- [Specify the Events to Generate SNMP Traps for MARS, page 7-2](#)

### Extracting Entercept Agent Information into a CSV file (for Entercept Version 2.5)



**Note**

Entercept agent information is saved in a database file on the Entercept console.

When you configure the MARS box to add Entercept agents, you can extract them from the database file on the Entercept console, instead of typing the mapping for each agent.

## Create a CSV file for Entercept Agents in Version 2.5

- 
- Step 1** Go to the directory Program Files\Cisco IDS\Console\Database and copy the file CoreShield.mdb to another directory, e.g.: C:\temp.
  - Step 2** Open the copied CoreShield.mdb with Microsoft Access, and go to the “Agents” table.
  - Step 3** Export the table to a file named: Agents.txt and choose the exported file format to be CSV.
  - Step 4** Copy Agents.txt to a specific directory that is ready for the MARS box to load.

A sample agents.txt file could be:  
 1,3,"entercept1",6,1,1,1,438,1,"127.0.0.1",0,,1051055867,2086

where the fields are: AgentID, AgentTypeID, ComputerName, ComputerType, NewFlag, StatusID, OperatingModeID, VersionID, VersionModeID, IP, License, Note, NoConnection, and UpTime.

---

## Define the MARS Appliance as an SNMP Trap Target

- 
- Step 1** Log in to the Entercept Console.
  - Step 2** Click **Configuration**.
  - Step 3** Click the **Address Book** tab.
  - Step 4** In the All Contacts tree, click **SNMP Trap**.
  - Step 5** Click the Plus (+) button.
  - Step 6** In the New SNMP Trap page:
    - a. Enter an **Alias** for the MARS Appliance.
    - b. Set **Privilege** Level to Global.
    - c. Set **Status** to Enabled.
    - d. Enter the MARS Appliance’s name if the DNS server can resolve the name. Otherwise, use its IP address.
    - e. Enter a community string name in the **Community** field.
    - f. Enter a **Port** number.
    - g. Select a **Protocol**.
- 

## Specific the Events to Generate SNMP Traps for MARS

- 
- Step 1** Click the **Notifications** tab.
  - Step 2** Click the Plus (+) button.

- Step 3** On the General tab, in the name field, enter a name for the notification.
- Step 4** Click the **Agent Groups** tab and select the **All Agents** radio button.
- Step 5** Click the **Security Events** tab and select the **Events by Severity Levels** radio button. Select the events that you want (**High**, **Medium**, **Low**, and **Information**).
- Step 6** Click the **System Events** tab and select the **Events by Severity Levels** radio button. Select the events that you want (**Error**, **Warning**, and **Information**).
- Step 7** Click the **Address Book** tab and click a destination in the Available Destinations field. Click the **Down** arrow to move it into the Selected Destinations field.
- Step 8** Click **OK** and exit the program.
- 

## Add and Configure an Entercept Console and its Agents in MARS

Adding an Entercept device has two distinct steps. First, you add configuration information for the for the Entercept Console host. Second, you add the agents managed by that console.

- [Add and Configure an Entercept Console and its Agents in MARS, page 7-3](#)
- [Add Entercept Agents Manually, page 7-4](#)
- [Add Entercept Agents Using a Seed File, page 7-4](#)

### Add the Entercept Console Host to MARS

---

- Step 1** Click **Admin > Security and Monitor Devices > Add**.
- Step 2** From the **Device Type** list, select **Add SW Security apps on a new host** or **Add SW security apps on existing host**.
- Step 3** Enter the **Device Name** and **IP addresses** if adding a new host.
- Step 4** Click **Apply**.
- Step 5** Click on **Reporting Applications** tab.
- Step 6** From the **Select Application** list, select **Entercept 2.5 or 4.0**
- Step 7** Click **Add**.
- Step 8** Enter the **Console Name**.
- Step 9** Check the “**Is Sensor**” check box—which is asking if it is a sensor or not.
- Step 10** Enter the sensor’s **Agent Name**, which is the agent name for the console if it is an agent.

Management Console

→ \*Console Name:   
 → ☐ Is Sensor  
 \*Agent Name:

143220

- Step 11** Click **Submit**.  
 You could now add the agents.
- 

## Add Entercept Agents Manually

---

- Step 1** Click **Add Agent**.  
**Step 2** Select the device that already has agent running or **Add New**.  
**Step 3** Enter the **Device Name**, **Agent Name**, and its **Reporting IP** address if  
 Adding new device
- For the first interface, enter an IP address and mask.
  - For multiple interfaces, click **Add Interface**, and add the new interfaces' IP address and mask.
- Step 4** Click **Submit**.
- 

## Add Entercept Agents Using a Seed File

---

- Step 1** Click **Load From CSV**.  
**Step 2** Enter the FTP server information and location of the CSV (comma separated values) file.
- If you need to generate the Entercept Agent CSV file, see [Extracting Entercept Agent Information into a CSV file \(for Entercept Version 2.5\)](#), page 7-1.
- Step 3** Click **Submit**.
-

# Cisco Security Agent 4.x Device

To enable Cisco Security Agent (CSA) as a reporting device in MARS, you must identify the CSA Management Console (CSA MC) as the reporting device. The CSA MC receives alerts from the CSA agents that it monitors, and it forwards those alerts to MARS as SNMP notifications.

When MARS receives the SNMP notification, the source IP address in the notification is that of the CSA agent that originally triggered the event, rather than the CSA MC that forwarded it. Therefore, MARS requires host definitions for each of the CSA agents that can potentially trigger an event. These definitions are added as sub-components under the device definition of the CSA MC.

As of MARS, release 4.1.1, the MARS Appliance discovers CSA agents as they generate alerts, eliminating the need to manually define them. MARS parses the alert to identify the CSA agent hostname and to discover the host operating system (OS). MARS uses this information to add any undefined agents as children of the CSA MC as a host with either the Generic Windows (all Windows) or Generic (Unix or Linux) operating system value. You are still required to define the CSA MC; however, you are not required to define each agent. The default topology presentation for discovered CSA agents is within a cloud.

**Note**

The first SNMP notification from an unknown CSA agent appears to originate from the CSA MC. MARS parses this notification and defines a child agent of the CSA MC using the discovered settings. Once the agent is defined, all subsequent messages appear to originate from the CSA agent.

Prior to 4.1.1., you were required to manually add each agent or by using an exported hosts file, as defined in [Export CSA Agent Information to File, page 7-6](#).

**Note**

Prior to the 4.1.1 release, CSA was identified by the device type name *Cisco CSA 4.0*. As part of an upgrade, any Cisco CSA 4.0 devices were renamed as *Cisco CSA 4.x*. This new name includes support for Cisco CSA 4.0 and 4.5.

This section contains the following topics:

- [Configure CSA Management Center to Generate Required Data, page 7-5](#)
- [Add and Configure a CSA MC Device in MARS, page 7-7](#)
- [Troubleshooting CSA Agent Installs, page 7-10](#)

## Configure CSA Management Center to Generate Required Data

To bootstrap CSA, you must configure the CSA MC to forward SNMP notifications to the MARS Appliance. In addition, you can export the list of CSA agents in a format that MARS can import. However, this export operation is not necessary, as MARS discovers the agents as they generate notifications.

This section contains the following topics:

- [Configure CSA MC to Forward SNMP Notifications to MARS, page 7-6](#)
- [Export CSA Agent Information to File, page 7-6](#)

## Configure CSA MC to Forward SNMP Notifications to MARS

The only required configuration is to ensure that CSA MC forwards the SNMP notifications that it receives from agents to MARS. From these notifications, MARS is able to discover the agent and its relevant settings. It is also from these events that MARS learns about the host-level activities transpiring on your network.

To forward all notifications to the MARS Appliance, follow these steps:

- 
- Step 1** Log in to the CiscoWorks Server desktop.
  - Step 2** From the navigation tree, select **VPN/Security Management Solution >Management Center > Security Agents**.
  - Step 3** In the Management Center screen, click the **Alerts** link.
  - Step 4** Click **New**.
  - Step 5** In the Name and Description fields, enter a name and description for the SNMP notification.
  - Step 6** Scroll down and select the **SNMP** check box.
  - Step 7** In the Community name field, enter the SNMP notification's community name.
  - Step 8** In the Manager IP address field, enter the MARS's IP address.
  - Step 9** Click **Save** and exit the program.
- 

## Export CSA Agent Information to File

With the release of MARS 4.1.1, you are no longer required to define each Cisco CSA agent, as they are discovered as a device sends an SNMP notification to the CSA Management Console (CSA MC).

**Note**

The following instructions apply to Cisco CSA 4.x when Microsoft Internet Explorer is used to access the CSA MC web interface.

---

To export the all hosts report as a tab-delimited file, follow these steps:

- 
- Step 1** Log in to the CSA MC by accessing the console using the fully qualified domain name in the URL.  
When accessing the CSA MC, you must use a fully qualified domain name in the URL. If you use the CiscoWorks Desktop to launch CSA MC, the ActiveX reports do not display.
  - Step 2** Click **Reports > Host Details**.
  - Step 3** Click **New**.
  - Step 4** In **Groups**, choose **<All Hosts>**, in **Viewer Type**, choose **ActiveX (IE only)**.
  - Step 5** Click **View report**.  
A window appears that contains the host details.
  - Step 6** Click **Export**, and select export to an **Excel 5.0 Document** type.
  - Step 7** In the **Name** box, identifies the name for the file that you are exporting, for example, csahosts.xls.
  - Step 8** Open the exported file in Excel, and click **File > Save As...**
  - Step 9** In the Save as type box, click **Text (Tab delimited) (\*.txt)**.

**Step 10** In the File name box, enter the name for this file, for example, csahosts.txt, and click **Save**.

**Step 11** Upload the generated file to an FTP server where the MARS Appliance can access it.

You will return to this file when adding the CSA device in the MARS web interface, as defined in [Add and Configure a CSA MC Device in MARS, page 7-7](#).

## Add and Configure a CSA MC Device in MARS

Before you can identify the agents, you must add the CSA MC to MARS. All CSA agents forward notifications to the CSA MC, and the CSA MC forwards SNMP notifications to MARS. Once you define the CSA MC and activate the device, MARS can discover the agents that are managed by that CSA MC. However, you can also choose to manually add the agents.

To add a CSA MC to MARS, follow these steps:

**Step 1** Click **Admin > Security and Monitor Devices > Add**.

**Step 2** From the **Device Type** list, select **Add SW security apps on a new host** or **Add SW security apps on existing host**.

**Step 3** Enter the **Device Name** and **IP addresses** if adding a new host.

**Step 4** Click **Apply**.

**Step 5** Click **Reporting Applications** tab.

**Step 6** From the **Select Application** list, select **Cisco CSA 4.x**.

**Step 7** Click **Add**.

**Step 8** The Management Console page appears.

### Management Console

Add or edit agents for this csa management console.

Add Agent

Edit Agent

Delete Agent

Load From File

Cancel

Submit

143194

**Step 9** Click **Submit**, and then click **Done**.

**Step 10** Do one of the following:

- To save your changes and allow the CSA agents to be discovered automatically, click **Submit**, and then click **Done**.
- To add agents using an exported hosts report, continue with [Add CSA Agents From File, page 7-9](#).
- To add a single agent manually, continue with [Add a CSA Agent Manually, page 7-8](#)

## Add a CSA Agent Manually

You can manually add a CSA Agent as a child of the CSA MC. This feature allows you to represent all of your agents, even if they have not generated any notifications.

To add CSA agents manually, follow these steps:

- Step 1** Click **Admin > Security and Monitoring Devices**.
- Step 2** From the list of devices, select the host running Cisco CSA Management Center, and click **Edit**.
- Step 3** Click the **Reporting Applications** tab, select **Cisco CSA Management Center** in the Device Type list, and click **Edit**.
- Step 4** Click the **Add Agent**.
- Step 5** Do one of the following:
  - Select the existing device, click **Edit Existing**, and continue with [Step 8](#).  
A page displays with the values pre-populated for hostname, reporting IP address, and at least one interface.
  - Click **Add New**, and continue with [Step 6](#).

- Step 6** In the Device Name field, enter the hostname on which this CSA agent resides.  
This value should reflect the DNS entry for this device.
- Step 7** In the Reporting IP field, enter the IP address that the agent uses to send logs to the CSA MC.
- Step 8** Define each interface that is configured for this host by specifying the interface name, IP address, and network mask. To add a new interface, click **Add Interface**.  
The interface settings are used for attack path calculation. It is very important that you identify any dual-homed hosts by defining each interface.
- Step 9** Click **Submit**, and then click **Done**.
- Step 10** To activate this device, click **Activate**.



## Add CSA Agents From File

You can add the complete list of hosts on which CSA Agents are installed by exporting the all hosts report from CSA MC and importing that file into MARS. The only advantage to adding agents using an export file is that the first notification received that originates from the agent is not attributed to the CSA MC.

To add CSA agents from a file, follow these steps:

- Step 1** Click **Admin > Security and Monitoring Devices**.
- Step 2** From the list of devices, select the host running Cisco CSA Management Center, and click **Edit**.
- Step 3** Click the **Reporting Applications** tab, select **Cisco CSA Management Center** in the Device Type list, and click **Edit**.
- Step 4** Click **Load From File**.

Remote File Location:



### Caution

The file should be formatted as a tab delimited file. You cannot use a CSV file. To generate a tab delimited file of the CSA agents managed by the CSA MC, see [Export CSA Agent Information to File](#), page 7-6.

- Step 5** In the IP Address field, enter the address of the FTP server where you stored the exported hosts file, as described in [Export CSA Agent Information to File](#), page 7-6.
- Step 6** In the User Name field, enter the name of the account used to authenticate to the FTP server.
- Step 7** In the Password field, enter the password that corresponds to the account specified in [Step 6](#).
- Step 8** In the Path field, enter the path to the folder where the file is stored. If this file is stored in the root folder, you must specify a backslash (\) in this field. The format of this value is \<path\_here>.
- Step 9** In the File Name field, enter the name of the tab delimited file.
- Step 10** Click **Submit**.

The following message displays and the hosts are added as agents of the CSA MC:

Success:  
Status: OK

- Step 11** Click **Done**.

## Troubleshooting CSA Agent Installs

When importing CSA agents from a file, the following messages can occur.

**Table 7-1** *Error and Status Messages when Importing CSA Agents from File*

| Message                                                                  | Description/Issue                                                                                                                                                                        |
|--------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Status: NumberFormatException occurred parsing the file at line <i>X</i> | Occurs when you have a CSV file rather than a tab delimited file. The line number varies.                                                                                                |
| Error Occurred:<br>Status: DbDevice occurred parsing the file at line -1 | Occurs when duplicate files are imported, even if you have deleted all of the agents and the CSA MC.                                                                                     |
| Success:<br>Status: OK                                                   | Indicates a successful import of CSA agents using the tab-delimited file.                                                                                                                |
| Error Occurred:<br>Status: FileNotFoundException                         | Indicates that the file does not exist at the specified path. If the path is at the root of your FTP server, verify that you have included \ as the path value.                          |
| Error Occurred:<br>Status: NoRouteToHostException                        | Indicates that the identified FTP server is not reachable from the MARS Appliance. You may need to define additional routes or enable traffic flows to ensure the connection is allowed. |



# Configuring Antivirus Devices

---

Antivirus (AV) devices provide detection and prevention against known viruses and anomalies.

This chapter describes how to configure and add the following devices and systems:

- [Symantec AntiVirus Configuration, page 8-1](#)
- [McAfee ePolicy Orchestrator Devices, page 8-8](#)
- [Cisco Incident Control Server, page 8-13](#)

## Symantec AntiVirus Configuration

Configuring the Symantec AV requires performing two tasks:

- [Configure the AV Server to Publish Events to MARS Appliance, page 8-1](#)
- [Add the Device to MARS, page 8-7](#)

In addition, you can perform the following task to expedite populating the Agent list in MARS:


- [Export the AntiVirus Agent List, page 8-7](#)

## Configure the AV Server to Publish Events to MARS Appliance

To configure the AV server to publish events to MARS, follow these steps:


- 
- |               |                                                                                                                                                                          |
|---------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | Log in to the Windows server running Symantec AV.                                                                                                                        |
| <b>Step 2</b> | To identify the Local Controller as a valid SNMP trap destination, click <b>Administrative Tools &gt; Services &gt; SNMP Service &gt; Traps &gt; Trap destinations</b> . |
| <b>Step 3</b> | Enter the IP address of the Local Controller in the Trap Destination page, and click <b>OK</b> to close all open windows.                                                |
| <b>Step 4</b> | Select <b>Start &gt; All Programs &gt; Symantec System Center Console</b> .                                                                                              |
| <b>Step 5</b> | In the Symantec System Center window, click <b>System Hierarchy</b> .                                                                                                    |
| <b>Step 6</b> | Under System Hierarchy, right-click the appropriate server group name and unlock the server group by supplying the configured password.                                  |

Unlocking the server enables you to configure it.

**Figure 8-1 Symantec Unlock Server**


143280

- Step 7** Configure Symantec server (AMS-Alert Management System) to send SNMP traps to MARS. Right click the unlocked server group name, then select **All Tasks > AMS > Configure**.


**Figure 8-2 Symantec AV AMS**

143286

- Step 8** Select **Send SNMP Trap** under each Alert Action, then click **Configure**.


**Figure 8-3 Symantec AV Trap**

**Step 9** Click **Send SNMP trap**, and then click **Next**.

**Figure 8-4 Symantec AV Send SNMP Trap**

**Step 10** Select the Local Controller to send the SNMP trap to as defined in [Step 3](#), and then click **Next** to view the Action Message window.

**Step 11** Add alert parameters to the Alert message list according to the following information:

**Figure 8-5 Symantec AV Action Msg**

The following mandatory fields are required for MARS to parse AV traps. If these fields are among those possible, you must define these fields in order before defining any of the optional fields.

- Alert: <Alert Name>
- Computer: <Computer Name>
- Date: <Date>
- Time: <Time>
- Action: <Actual Action>
- Description: <Description>

**Note**

This ordering is required because some optional fields can be so long as to prevent Mars from correctly parsing the mandatory fields if they do not appear first in the list of attributes.

The following optional fields can be defined after all mandatory fields are defined:

- User: <User>
- Virus Name: <Virus Name>
- File Path: <File Path>
- Severity: <Severity>
- Source: <Source>

The following list identifies the trap type and the full list of possible fields:

**Alert: Virus Found**

- Alert: <Alert Name>

- Computer: <Computer Name>
- Date: <Date>
- Time: <Time>
- Action: <Actual Action>
- Severity: <Severity>
- Source: <Source>
- File Path: <File Path>
- Logger: <Logger>
- Requested Action: <Requested Action>
- User: <User>
- Virus Name: <Virus Name>

**Alert: Virus Definition File Update**

- Alert: <Alert Name>
- Computer: <Computer Name>
- Date: <Date>
- Time: <Time>
- Description: <Description>
- Severity: <Severity>
- Source: <Source>

**Alert: Symantec AntiVirus Startup/Shutdown**

- Alert: <Alert Name>
- Computer: <Computer Name>
- Date: <Date>
- Time: <Time>
- Description: <Description>
- Severity: <Severity>
- Source: <Source>

**Alert: Scan Start/Stop**

- Alert: <Alert Name>
- Computer: <Computer Name>
- Date: <Date>
- Time: <Time>
- Severity: <Severity>
- Source: <Source>
- Source: <Source>
- Logger: <Logger>
- User: <User>

**Alert: Scan Start/Stop**

- Alert: *<Alert Name>*
- Computer: *<Computer Name>*
- Date: *<Date>*
- Time: *<Time>*
- Description: *<Description>*
- Severity: *<Severity>*
- Source: *<Source>*
- Logger: *<Logger>*

**Alert: Default Alert**

- Alert: *<Alert Name>*
- Computer: *<Computer Name>*
- Date: *<Date>*
- Time: *<Time>*
- Severity: *<Severity>*
- Source: *<Source>*
- Failed Alert: *<Failed Alert>*

**Alert: Configuration Change**

- Alert: *<Alert Name>*
- Computer: *<Computer Name>*
- Date: *<Date>*
- Time: *<Time>*
- Severity: *<Severity>*
- Source: *<Source>*
- Failed Alert: *<Failed Alert>*

**Alert: Configuration Change**

- Alert: *<Alert Name>*
- Computer: *<Computer Name>*
- Date: *<Date>*
- Time: *<Time>*
- Description: *<Description>*
- Severity: *<Severity>*
- Source: *<Source>*

**Step 12** Repeat [Step 8](#) through [Step 11](#) for each alert event.

---



## Export the AntiVirus Agent List

You can export the list of Symantec AntiVirus Clients and Agents as a CSV file (\*.csv), which enables you to use the CSV file to load the agents into MARS. For more information on adding agents from the file, [Add Agents from a CSV File, page 8-8](#). This approach is much faster than if you had to identify the agents manually.

To generate the CSV file, follow these steps:

- 
- Step 1** Select **View > Default Console View** to ensure the generated CSV file will be based on the Console Default View.
  - Step 2** Right-click the name of the server that you want to export, choose **Export List**, and save it as Text (Comma Delimited) (\*.csv) file.
  - Step 3** Copy the file to an FTP server that the MARS Appliance can access.
- You will use this file when you add the AntiVirus agents within the web interface.
- 

## Add the Device to MARS

Adding a device to MARS has two distinct steps. First, add the host configuration information. Then, add its agents, either manually or from the seed file.

**Tip**

For Symantec AntiVirus, the Symantec agent hostname (AV client computer name) appears in the "Reported User" column of the event data. Therefore, you can define a query, report or rule related to this agent based on the "Reported User" value.

---

To add the host and application configuration information, follow these steps:

- 
- Step 1** Select **Admin > Security and Monitor Devices > Add**.
  - Step 2** Select **Add SW Security apps on a new host** or **Add SW security apps on existing host** from the Device Type list.
  - Step 3** To add a new host, enter the device name and IP addresses.
  - Step 4** Click **Apply**.
  - Step 5** Click the **Reporting Applications** tab.
  - Step 6** From the Select Application list, select **Symantec AntiVirus 9.x**.
  - Step 7** Click **Add**, then add the agents. Continue with [Add Agent Manually, page 8-7](#).
- 

## Add Agent Manually

To add an agent manually, follow these steps:

- 
- Step 1** Click **Add Agent**.

- Step 2** Select the existing device or click **Add New**.
- Step 3** Enter the following information for new device.
- **Device Name**—The DNS entry for this device.
  - **Reporting IP**—The IP address that the agent uses to send logs to the console.
- Step 4** Under the Interfaces list, specify the IP address and netmask values associated with each interface installed in the host on which the agent is running.
- MARS uses interface information to calculate attack paths.
- Step 5** Click **Submit**.
- 

## Add Agents from a CSV File

You can generate a CSV file that contains the list of agents managed by the Symantec AV server. Once the file is generated, you can use the file to import the list of agents into the MARS web interface as child modules of the Symantec AV server.

To import the list of AV agents into MARS, follow these steps:

- 
- Step 1** Click **Load From CSV**.
- Step 2** Enter the FTP server information and location of the CSV (comma-separated values) file.
- Step 3** Click **Submit**.
- 

## McAfee ePolicy Orchestrator Devices


Configuring MARS to receive and process the data generated by a McAfee ePolicy Orchestrator server requires you to perform two procedures:

- [Configure ePolicy Orchestrator to Generate Required Data, page 8-8](#)
- [Add and Configure ePolicy Orchestrator Server in MARS, page 8-12](#)


## Configure ePolicy Orchestrator to Generate Required Data

To prepare the ePolicy Orchestrator server to forward SNMP events to MARS, follow these steps:

- 
- Step 1** Select **Start > Program Files > Network Associates > ePolicy Orchestrator 3.x Console**.
- Step 2** In the tree, select **McAfee Security > ePolicy Orchestrator**, and click the **Log on to server** link under Global Task List.



- Step 3** In the Log On to Server dialog box, enter the username and password required to access the ePolicy Orchestrator server, and click **OK**.
- Step 4** In the tree, select **McAfee Security > ePolicy Orchestrator > <Server\_Name> > Notifications** and click the **Configuration** tab and click the **SNMP Servers** link.
- Step 5** Click **Add**.




- Step 6** In the Name field, enter the hostname of the MARS Appliance.
- Step 7** In the Server address field, enter the IP address of the eth0 interface, the monitoring interface for the MARS Appliance, and click **OK**.
- The SNMP server is added to represent the MARS Appliance.
- Step 8** Click the **Rules** tab.

You can access the Rules tab by selecting **McAfee Security > ePolicy Orchestrator > <Server\_Name> > Notifications >** and then clicking the **Rules** tab.

**Step 9** Edit each rule in the list so that all notifications are sent to the SNMP server that represents the MARS Appliance. To edit a rule, follow these steps:

a. Click the rule.

The Describe Rule wizard page appears.



b. Click **Next** to proceed to Set Filters page.

c. Under Add or Edit Notification Rule, click the **3. Set Thresholds** link.

Figure 8-6 Set Threshold Values

**Add or Edit Notification Rule** Back

1. Describe Rule | 2. Set Filters | 3. Set Thresholds | 4. Create Notifications | 5. View Summary

For notification rule: **Daily unknown category notification**

You can use aggregation and throttling to limit the number of notifications you receive. Each sends a single notification that summarizes multiple events.

**Aggregation:** ☒ Send a notification for every event

☐ Send a notification for multiple events within:  Minutes

☐ When the number of affected computers is at least:

or

☐ When the number of events is at least:

**Throttling:** ☒ At most, send notification every:  Days

< Back Next > Finish Cancel

- d. Verify the Aggregation and Throttling values are set as shown in Figure 8-6 on page 8-11.
- e. Click **Next** to proceed to the Create Notifications page.

**Add or Edit Notification Rule** Back

1. Describe Rule | 2. Set Filters | 3. Set Thresholds | 4. Create Notifications | 5. View Summary

For notification rule: **Daily unknown category notification**

| Notification Type | Detail          | Recipients    | Test              | Delete         |
|-------------------|-----------------|---------------|-------------------|----------------|
| E-mail            | Standard E-mail | Administrator | <span>Test</span> | <span>✖</span> |
| SNMP Trap         | Warning         | tucson        | <span>Test</span> | <span>✖</span> |

Add E-mail Message Add SNMP Trap Add External Command

< Back Next > Finish Cancel

- f. Click **Add SNMP Trap**.

**Figure 8-7** SNMP Trap Settings

**Add or Edit SNMP Trap**

For notification rule: Daily unknown category notification

SNMP server:

Replace variables with their values in:

Variables to include:

|                                                                    |                                                                  |
|--------------------------------------------------------------------|------------------------------------------------------------------|
| <input checked="" type="checkbox"/> Actual categories              | <input checked="" type="checkbox"/> Actual number of computers   |
| <input checked="" type="checkbox"/> Actual number of events        | <input checked="" type="checkbox"/> Actual products              |
| <input checked="" type="checkbox"/> Actual threat or rule names    | <input checked="" type="checkbox"/> Additional information       |
| <input checked="" type="checkbox"/> Affected computer IP addresses | <input checked="" type="checkbox"/> Affected computer names      |
| <input checked="" type="checkbox"/> Affected objects               | <input checked="" type="checkbox"/> Event descriptions           |
| <input checked="" type="checkbox"/> Event IDs                      | <input checked="" type="checkbox"/> First event time             |
| <input checked="" type="checkbox"/> Notification rule name         | <input checked="" type="checkbox"/> Rule defined at              |
| <input checked="" type="checkbox"/> Rule site                      | <input checked="" type="checkbox"/> Selected categories          |
| <input checked="" type="checkbox"/> Selected products              | <input checked="" type="checkbox"/> Selected threat or rule name |
| <input checked="" type="checkbox"/> Source computers               | <input checked="" type="checkbox"/> Time notification sent       |

- g. In the SNMP server list, select the SNMP server that represents the MARS Appliance.
- h. Verify that all the variables are selected as shown in [Figure 8-7 on page 8-12](#).
- i. Click **Save** to add the SNMP trap to the list of notifications for the selected rule.
- j. Click **Finish** to save the changes to the selected rule.
- k. Repeat Steps [a.](#) through [j.](#) for each rule.

## Add and Configure ePolicy Orchestrator Server in MARS

Before MARS can begin processing SNMP traps from ePolicy Orchestrator, you must define the ePolicy Orchestrator server as software running on a host. When ePolicy Orchestrator is defined as a reporting device, MARS can process any inspection rules that you have defined using ePolicy Orchestrator event types.

After you add the ePolicy Orchestrator server to MARS, the appliance can discover the agents that are managed by the ePolicy Orchestrator server as events are generated by those agents. You do not need to manually define the agents associated with this server.

To add an ePolicy Orchestrator server to MARS, follow these steps:

- Step 1** Select **Admin > Security and Monitor Devices > Add**.
- Step 2** From the Device Type list, select **Add SW Security apps on a new host**.

- Step 3** In the Device Name field, enter the hostname of the server.
- Step 4** In the Reporting IP field, enter the IP address of the interface in the ePolicy Orchestrator server from which SNMP traps will originate.
- Step 5** Under Enter interface information, enter the interface name, IP address, and netmask value of the interface in the ePolicy Orchestrator server from which syslog messages will originate.  
This address is the same value as the Reporting IP address.
- Step 6** Click **Apply**.
- Step 7** Click **Next** to move to the Reporting Applications tab.
- Step 8** In the Select Application field, select **McAfee ePO 3.5**, and then click **Add**.

Management Console

Add or edit agents for this McAfee epo server.

143284

- Step 9** Click **Done** to save the changes.
- Step 10** Click **Submit**.
- Step 11** To activate the device, click **Activate**.

## Cisco Incident Control Server

The Cisco Incident Control Server (Cisco ICS) enables extended protection across Cisco IOS routers, switches, and IPS devices. In coordination with Trend Micro's incident control solutions, Cisco ICS prevents the spread of day-zero outbreaks in three ways:

- First, Cisco ICS issues temporary ACLs to those Cisco mitigation devices that can block such traffic, typically using a protocol and port pair block. This temporary block is referred to as an Outbreak Prevention ACL (OPACL).
- Second, as soon as a signature is available, Cisco ICS updates all Cisco IPS and IDS devices running on your network with the signature required to detect and prevent the specific threat. This signature is referred to as an Outbreak Prevention Signature (OPSig).
- Third, Cisco ICS can manage supporting products (sold separately), such as Trend Micro's Damage Cleanup Services (DCS), which cleans infected hosts by removing trojans and other malware.

To complete the Cisco ICS communication settings, you must perform two tasks: configure Cisco ICS to send syslog messages to the MARS Appliance, and add the Cisco ICS management server to the MARS web interface as a reporting device.

This section contains the following topics:

- [Configure Cisco ICS to Send Syslogs to MARS, page 8-14](#)


- [Add the Cisco ICS Device to MARS, page 8-15](#)
- [Define Rules and Reports for Cisco ICS Events, page 8-15](#)

## Configure Cisco ICS to Send Syslogs to MARS


Cisco ICS publishes syslog messages to MARS. To configure Cisco ICS, you simply define a syslog server with the IP address of the MARS Appliance. You do not need to enable any special logs, and you cannot tune the messages that are sent to MARS. The Cisco ICS events for which syslog messages are generated have been selected to provide the most benefit to your Security Threat Mitigation (STM) system.

To prepare Cisco ICS to publish events to MARS, follow these steps:

- Step 1** Log in to the Cisco ICS Management Console.
- Step 2** Click **Global Settings > Syslog Servers**.



- Step 3** Click **Add**.



- Step 4** In the IP Address field, enter the address of the MARS Appliance to which the Cisco ICS will publish syslog messages.
- Step 5** Click **Save**.

Cisco ICS now publishes syslog message to MARS. For MARS to be aware of this device, you must add the Cisco ICS device as a software application running on a host and you must click Activate in the web interface.



## Add the Cisco ICS Device to MARS

Before MARS can be processing the syslog messages as Cisco ICS messages, you must define the Cisco ICS management server as a software application running on a host. After Cisco ICS is defined as a reporting device, MARS can process any inspection rules that you have defined using Cisco ICS event types.

To add a Cisco ICS server to MARS, follow these steps:

- 
- Step 1** Click **Admin > Security and Monitor Devices > Add**.
  - Step 2** From the Device Type list, select **Add SW Security apps on a new host**.  
You can also select Add SW Security apps on an existing host if you have already defined the host within MARS, perhaps as part of the Management > IP Management settings or if you are running another application on the host, such as Microsoft Internet Information Services.
  - Step 3** In the Device Name field, enter the hostname of the server.
  - Step 4** In the Reporting IP field, enter the IP address of the interface in Cisco ICS server from which the syslog messages will originate.
  - Step 5** Under Enter interface information, enter the interface name, IP address, and netmask value of the interface in Cisco ICS server from which the syslog messages will originate.  
This address is the same value as the Reporting IP address.
  - Step 6** Click **Apply**.
  - Step 7** Click **Next** to move the Reporting Applications tab.
  - Step 8** In the Select Application field, select **Cisco ICS 1.x**, then click **Add**.

### Cisco ICS

Submit if you want to add Cisco ICS to this host



143191

- Step 9** Click **Select** to add the Cisco ICS application to this host.
  - Step 10** Click **Done** to save the changes.
  - Step 11** To activate the device, click **Activate**.
- 

## Define Rules and Reports for Cisco ICS Events

From Cisco ICS, MARS receives syslog messages that allow it to identify outbreaks, successful OPACL and OPSig deployments, and failed attempts to deploy. MARS stays abreast of when the OPACLs and OPSigs fire on Cisco IPS devices. MARS also monitors the Cisco ICS server for system issues, such as database failures.

These events assist MARS in providing an accurate, holistic assessment of your network. OPACL and OPSig matching events provide five-tuple correlation, which MARS uses to perform attack path analysis and verify the containment of threats. You can use the events to define inspection rules that help you perform manual mitigation on devices that cannot use OPACLs and OPSigs.

For example, an inspection rule could be written to match the OPACL event. Your mitigation team can respond by investigating the OPACL that was pushed to the reporting device, from which they can determine the five tuple (source address and port, destination address and port and network service). Using that information, they could push equivalent ACLs to devices not managed by Cisco ICS.

When defining inspection rules or reports, you can access the list of Cisco ICS-specific events by entering *Cisco ICS* in the Description / CVE: field and clicking Search on the Management > Event Management page of the web interface.

There are four predefined system inspection rules for Cisco ICS:

- New Malware Discovered
- New Malware Prevention Deployed
- New Malware Prevention Deployment Failed
- New Malware Traffic Match

In addition, there are five predefined reports:

- Activity: New Malware Discovered - All Events
- Activity: New Malware Prevention Deployment Failure - All Events
- Activity: New Malware Prevention Deployment Success - All Events
- Activity: New Malware Traffic Match - All Events
- Activity: New Malware Traffic Match - Top Sources



## Configuring Vulnerability Assessment Devices

---

Vulnerability assessment (VA) devices provide MARS with valuable information about many of the possible targets of attacks and threats. They provide information useful for accurately assessing false positives. This information includes the operating system (OS) running on a host, the patch level of the OS, the type of applications running on the host, as well as detailed logs about the activities occurring on that host.

This chapter explains how to bootstrap and add the following VA devices to MARS:

- [Foundstone FoundScan 3.0, page 9-1](#)
- [eEye REM 1.0, page 9-3](#)
- [Qualys QualysGuard Devices, page 9-5](#)

### Foundstone FoundScan 3.0

To configure MARS to pull data from FoundScan, you must perform three tasks:

- Configure Foundstone FoundScan to correlate the required data, ensuring that the data is current.
- Add the Foundstone FoundScan server to MARS using the web interface.
- Schedule the interval at which the Foundstone FoundScan server data is pulled by MARS.


This section contains the following topics:

- [Configure FoundScan to Generate Required Data, page 9-1](#)
- [Add and Configure a FoundScan Device in MARS, page 9-2](#)


### Configure FoundScan to Generate Required Data

To configure FoundScan to provide data to MARS, follow these steps:

- 
- Step 1** Run command **svrnetcn** at the DOS prompt on the host where FoundScan is installed.



- Step 2** In the SQL Server Network Utility dialog box, enable TCP/IP by moving **TCP/IP** from the Disabled Protocols list to Enabled Protocols list.



- Step 3** Click **Apply**.

## Add and Configure a FoundScan Device in MARS

To add a FoundScan device in MARS, follow these steps:

- Step 1** Select **Admin > Security and Monitor Devices > Add**.
- Step 2** Select **Add SW Security apps on a new host** or **Add SW security apps on existing host** from the Device Type list.
- Step 3** Enter the device name and IP addresses if adding a new host.
- Step 4** Click **Apply**.
- Step 5** Click the **Reporting Application** tab

**Step 6** From the Select Application list, select **Foundstone FoundScan 3.0**

**Step 7** Click **Add**.

**Step 8** Enter the following information:

- **Database Name**—The name for this database.
- **Access Port**—The default access port is 1433.
- **Access Type**—Verify the value is MS SQL.
- **Login**—The login information for the database.
- **Password**—The password for the database.

**Step 9** Click **Submit**.

**Step 10** Click **Apply**.

Once you activate this device (click Activate in the web interface), you must define the schedule at which MARS should pull data from it. For more information, see [Scheduling Topology Updates, page 2-39](#).

## eEye REM 1.0

To configure MARS to pull this REM data, you must perform three tasks:

- Configure eEye REM to correlate the required data, ensuring that the data is current.
- Add the eEye REM server to MARS using the web interface.
- Schedule the interval at which the eEye REM server data is pulled by MARS.


This section contains the following topics:

- [Configure eEye REM to Generate Required Data, page 9-3](#)
- [Add and Configure the eEye REM Device in MARS, page 9-4](#)


## Configure eEye REM to Generate Required Data

To configure eEye REM to provide the correct data to MARS, follow these steps:

**Step 1** Run command **svrnetcn** at the DOS prompt on the host where eEye REM 1.0 is installed.



- Step 2** In the SQL Server Network Utility dialog box, enable TCP/IP by moving **TCP/IP** from the Disabled Protocols list to Enabled Protocols list.



- Step 3** Click **Apply**.

## Add and Configure the eEye REM Device in MARS

To add the eEye REM device in MARS, follow these steps:

- Step 1** Select **Admin > Security and Monitor Devices > Add**.
- Step 2** Select **Add SW Security apps on a new host** or **Add SW security apps on existing host** from the Device Type list.
- Step 3** Enter the device name and IP addresses if adding a new host.
- Step 4** Click **Apply**.
- Step 5** Click the **Reporting Applications** tab.

**Step 6** From the Select Application list, select **eEye REM 1.0**.

**Step 7** Click **Add**.

**Step 8** Enter the following information:

- **Database Name**—The name for this database.
- **Access Port**—The default access port is 1433.
- **Login**—The login information for the database.
- **Password**—The password information for the database.

**Step 9** Click **Submit**.

**Step 10** Click **Apply**.

Once you activate this device (click **Activate** in the web interface), you must define the schedule at which MARS should pull data from it. For more information, see [Scheduling Topology Updates](#), page 2-39.

## Qualys QualysGuard Devices

In MARS, a QualysGuard device represents a specific report query to the QualysGuard API Server, which is the central API server hosted by Qualys. The only one that you configure to work with MARS is the QualysGuard API Server. You want to ensure that the QualysGuard API Server can provide reports about the devices on the network segments that you are monitoring with the MARS Appliance, as each MARS Appliance is responsible for identifying false positives for the network segments it monitors.

If you have a subscription to the QualysGuard service, MARS can pull VA data from the QualysGuard database using the QualysGuard XML API, version 3.3. To configure MARS to pull this data, you must perform three tasks:

- Configure QualysGuard to collect the required data, ensuring that the data is current.
- Add the QualysGuard device that represents a report query to MARS using the web interface.
- Schedule the interval at which the QualysGuard device data is pulled by MARS.



### Note

If a proxy server resides between the QualysGuard server and the MARS Appliance, the settings defined on the **Admin > System Parameters > Proxy Settings** page are used. For more information, see [Specify the Proxy Settings for the Global Controller or Local Controller](#), page 6-12.

This section contains the following topics:

- [Configure QualysGuard to Scan the Network, page 9-6](#)
- [Add and Configure a QualysGuard Device in MARS, page 9-6](#)
- [Schedule the Interval at Which Data is Pulled, page 9-8](#)
- [Troubleshooting QualysGuard Integration, page 9-8](#)

## Configure QualysGuard to Scan the Network

MARS uses the QualysGuard XML API and password-based authentication over SSL (TCP port 443) to retrieve scan reports from the QualysGuard API Server. As such, you do not need to configure the QualysGuard server to accept connections from MARS. The only required configuration is that you have an active account and Qualys subscription that is configured correctly to scan your network.

By default, MARS assumes that you want to retrieve the most recent scan report saved on the QualysGuard server. Depending on the number of IP addresses analyzed, the QualysGuard scan takes from a few seconds to several minutes. You need to estimate this time so that you can schedule automated scans of your network with a frequency that ensures a recent saved scan report is available. Using the QualysGuard administrative interface, you can determine how long a scan takes and set the schedule accordingly.

## Add and Configure a QualysGuard Device in MARS

Adding an internal QualysGuard API Server as a reporting device entails identifying the server or appliance from which the reports are pulled and providing credentials that MARS can use to log in to the device to pull the reports. You can specify whether you want to pull saved scan reports that are run on a schedule or whether you want to initiate and retrieve an on-demand scan report.

To add a QualysGuard device, follow these steps:

**Step 1** Select **Admin > Security and Monitor Devices > Add**.

**Step 2** Select **QualysGuard 3.x** from the Device Type list.

Note:

1. \* denotes a required field.

Device Type:

→ \*Device Name:

→ Access IP:

→ \*URL:

Login:

Password:

143206

**Step 3** Enter the name of the Qualys device in the Device Name field.



This name is used to identify the Qualys device uniquely within MARS. It is used in reports and query results to identify this device.

The IP address field is read only. The value is also fixed at 165.193.18.12, which is significant because you can only define one schedule for pulling all report queries defined as Qualys devices on the Local Controller. However, you can define unique schedules across different Local Controllers. For more information, see [Scheduling Topology Updates, page 2-39](#).

**Step 4** Enter the URL that identifies the device and report type in the URL field.

The URL provides the following information:

- **Server.** Identifies the server from which the report should be pulled. This value can be specified as a hostname or IP address that identifies the primary Qualys server.
- **Report type.** Real-time vs. Last Saved. The default value.

- *Real-time Report.* `qualysguard.qualys.com/msp/scan.php?ip=[addresses]`

The *addresses* attribute specifies the target IP addresses for the scan request.

IP addresses may be entered as multiple IP addresses, IP ranges, or a combination of the two. Multiple IP addresses must be comma separated, as shown below:

`123.123.123.1,123.123.123.4,123.123.123.5`

An IP address range specifies a start and end IP address separated by a dash (-), as shown below:

`123.123.123.1-123.123.123.8`

A combination of IP addresses and IP ranges may be specified. Multiple entries must be comma separated, as shown below:

`123.123.123.1-123.123.123.5,194.90.90.3,194.90.90.9`



**Note**

You must use a Scanner Appliance to scan private IP addresses on your internal network.

- *Last Saved Report.* `qualysapi.qualys.com/msp/scan_report_list.php?last=yes`

**Step 5** Enter the username of the account that MARS will use to access the Qualys device in the Login field.

**Step 6** Enter the password that corresponds to the account identified in [Step 5](#) in the Password field.

**Step 7** (Optional) To verify that the settings are correct and that the MARS Appliance can communicate with this Qualys device, click **Test Connectivity**.

If you receive error messages during this test, refer to [Troubleshooting QualysGuard Integration, page 9-8](#).

**Step 8** To add this device to the MARS database, click **Submit**.

Once you activate this device (click Activate in the web interface), you must define the schedule at which MARS should pull data from it. For more information, see [Schedule the Interval at Which Data is Pulled, page 9-8](#).

## Schedule the Interval at Which Data is Pulled

Once you activate one or more Qualys devices (where each device represents a report query run on the QualysGuard API Server), you must define the schedule at which MARS pulls data from them. The schedule, or update rule, that you define is the same for all Qualys devices. This update rule is based on the fixed IP address of 165.193.18.12, which is the Qualys Access IP. When you define an update rule using this address, all Qualys devices are updated based on that schedule. Even if you have more than one Qualys device on your network, you cannot stagger when MARS queries those Qualys devices. However, you can define unique schedules across different Local Controllers.

For more information on the broader use of update rules, see [Scheduling Topology Updates, page 2-39](#). To define the rule by which all Qualys devices will be discovered, follow these steps:

- 
- Step 1** Click **Admin > Topology/Monitored Device Update Scheduler**.  
The Topology/Monitored Device Update Scheduler page displays.
  - Step 2** Click **Add**.
  - Step 3** Enter *Qualys Devices* or another meaningful value in the Name field.  
This name identifies the rule in the list of rules that appears on the Topology/Monitored Device Update Scheduler page.
  - Step 4** Select the **Network IP** radio button, and enter 165.193.18.12. and 255.255.255.255 in the Network IP and Mask fields respectively.
  - Step 5** Click **Add** to move the device into the selected field.
  - Step 6** In the Schedule table, select **Daily**, and select a time value from **Time of Day** list.  
We recommend that you pull this data daily, during off-peak hours, however, you can define any interval required by your organization.
  - Step 7** Click **Submit**.  
The update rule appears in the list on the Topology/Monitored Device Update Scheduler page.
  - Step 8** Click **Activate**.



**Tip** To perform this discovery on demand, select the check box next to the rule you just defined and click **Run Now**.

---

## Troubleshooting QualysGuard Integration

[Table 9-1](#) identifies possible errors and likely causes and solutions.

**Table 9-1**      **Error Table for QualysGuard and MARS Integration**

| Error/Symptom                                                                                                                                                                                                                                                            | Workaround/Solution                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Test connectivity failed. Click the View Errors button for more information.<br>Server unavailable.                                                                                                                                                                      | <p>This error means that MARS was unable to connect to the Qualys device. Four possible issues can account for this message:</p> <ul style="list-style-type: none"> <li>You have entered an invalid hostname or IP address in the URL field. Verify the value was entered correctly.</li> <li>The traffic may be blocked by either a proxy server or firewalls and gateways on your network. Enable SSL traffic (TCP port 443) to traverse between the MARS Appliance and the Qualys device. Enter the correct settings for your proxy server on the Admin &gt; System Parameters &gt; Proxy Setting page.</li> </ul> |
| Fail to parse scan report.                                                                                                                                                                                                                                               | <p>This error means that MARS was unable to parse the scan report that it pulled from the Qualys device. Two possible issues can account for this message:</p> <ul style="list-style-type: none"> <li>Data corruption on the QualysGuard device.</li> <li>Format changes to the report due to an issue on the QualysGuard device or due to a software upgrade on the QualysGuard device.</li> </ul> <p>Verify that the QualysGuard device is running a supported version and that the device data is not corrupted.</p>                                                                                               |
| Invalid user credentials.                                                                                                                                                                                                                                                | <p>This error means that MARS was unable to authenticate to the Qualys device. Two possible issues can account for this message:</p> <ul style="list-style-type: none"> <li>The provided login credentials are incorrect. Verify these values were entered correctly, and verify that the provided account has sufficient privileges.</li> <li>Your account has expired. Renew your subscription services with Qualys.</li> </ul>                                                                                                                                                                                     |
| Test connectivity failed for qualys.<br>Unknown host: qualysapi.qualys.com<br>Please make sure that,                                                                                                                                                                     | <p>Make sure that the DNS server is configured correctly for the MARS Appliance. For more information on these DNS settings, see <a href="#">Specifying the DNS Settings, page 5-15</a>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <ul style="list-style-type: none"> <li>Proxy settings are configured correctly, If there is no direct connection exists from CS-MARS to Qualys server</li> <li>The hostname specified in the URL string is correct</li> <li>Login name and Password is valid.</li> </ul> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |





## Configuring Generic, Solaris, Linux, and Windows Application Hosts

---

Application hosts are simply hosts on your network that are running important applications. Many of the supported reporting devices and mitigation devices cannot be represented in MARS until the base host on which they are running is defined. Examples of such applications include CheckPoint Firewalls and all forms of web servers.

MARS provides for the definition of the following host types:

- **Generic.** Identifies no specific operating system, as well as any that are not directly supported.
- **Windows.** Identifies one of the Microsoft operating systems.
- **Solaris.** Identifies any of the Solaris family of operating systems.
- **Linux.** Identifies any of the Linux family of operating systems.

You should strive to define the application host as exactly as possible. This guideline applies to the vulnerability assessment information as well as the general settings. This detailed information helps MARS determine whether the host is susceptible to known attacks, such as those that specifically target on operating system or application/service running on the host.

This chapter contains the following sections:

- [Adding Generic Devices, page 10-1](#)
- [Sun Solaris and Linux Hosts, page 10-2](#)
- [Microsoft Windows Hosts, page 10-4](#)
- [Define Vulnerability Assessment Information, page 10-11](#)

### Adding Generic Devices

The MARS can support any syslog or SNMP devices, even if they do not appear on the list of devices supported by the MARS. You can enter any syslog or SNMP device into the network topology, configure it to report data to the MARS, and query it using a free-form query. For more information on free form queries, see [To Run a Free-form Query, page 20-2](#).

# Sun Solaris and Linux Hosts

To configure MARS to receive and process Solaris or Linux host log information, you must perform three tasks:

- [Configure the Solaris or Linux Host to Generate Events, page 10-2](#)
- [Configure Syslogd to Publish to the MARS Appliance, page 10-2](#)
- [Configure MARS to Receive the Solaris or Linux Host Logs, page 10-3](#)

## Configure the Solaris or Linux Host to Generate Events

MARS Appliance can receive syslog information from a Linux/Solaris host. To configure the Linux/Solaris applications, you must configure the following applications to write to syslog:

- xferlog
- inetd

To configure these applications to write to the system log, follow these steps:

- 
- Step 1** xferlog (which provides transfer logging information from the FTP server)
- For ftpd, add the following to `/etc/ftpd/ftpaccess`:
- ```
log transfers real,guest,anonymous inbound,outbound log syslog+xferlog
```
- Step 2** inetd trace messages (which provide the authentication information for services provided using inetd)
- For inetd, the line in `/etc/rc2.d/S72inetsvc` that reads:
- ```
/usr/sbin/inetd -s
```
- needs to be changed to:
- ```
/usr/sbin/inetd -t -s
```
- Other messages will automatically appear in the syslog and do not need to be specifically configured.
- Step 3** Once you have enabled the message generation, you must configure the syslogd daemon to publish messages to the MARS Appliance. For more information, see [Configure Syslogd to Publish to the MARS Appliance, page 10-2](#).
-

Configure Syslogd to Publish to the MARS Appliance

Once you have enabled the correct applications to write to the system log, you must configure the syslog daemon on the Solaris or Linux host to publish syslog messages to the MARS Appliance.

To configure the Solaris or Linux host to publish syslogs to the MARS Appliance, follow these steps:

-
- Step 1** Edit `/etc/syslog.conf` file and add the line below:
- ```
*.debug @MARS_hostname
```
- where `MARS_hostname` is the hostname or IP address of the MARS Appliance.
- Step 2** Run following commands to restart syslogd so that the changes are process:
- ```
/etc/init.d/syslog stop
```

`/etc/init.d/syslog start`


Once this line is added to the `syslog.conf` file and you have restarted `syslogd`, any messages sent to console are also sent to the MARS Appliance.

Configure MARS to Receive the Solaris or Linux Host Logs

To add a generic device to MARS, follow these steps:

Step 1 Click **Admin > Security and Monitor Devices > Add**.

Figure 10-1 Adding a Generic Device



Step 2 From the **Device Type** list, select **Add SW Security apps on a new host**.

Figure 10-2 Adding a Generic Device to receive logs

- Step 3** Enter the **Device Name**, and its **Reporting IP** address.
- Step 4** Select **Operating System** as **Generic**.
- Step 5** Select **Logging Info** and select **Receive**, then click **Submit**.
- Step 6** Click **Apply** to add the device.

Microsoft Windows Hosts

MARS processes data pulled from hosts running Microsoft Windows. This data includes the events found in the security event log as well application event and system event logs. You can use one of two methods to retrieve the logs from a host running Microsoft Windows, whether it is a server or workstation version:

- You can configure MARS to pull the logs from the host.
- You can configure the host to send the log data to the MARS Appliance.

These two methods are mutually exclusive; in other words, you cannot configure both methods. Your decision in which method to use depends on how much time you can spend preparing the host, the desired load on the MARS Appliance, and how near real-time you want MARS to process the event data.

The *pull method* not only requires system resources for correlating, but also for contacting and pulling the event data from each host. It also operates in a single process, completing the pull from one device before moving to the next. As a result, the pull method may take much longer to cycle through all of the reporting devices as the number of devices grows.

The *push method* is more efficient in terms of resource utilization on the MARS Appliance and in terms of how quickly the MARS Appliance can be made aware of event data, but it requires that you install and configure the Snare Agent for Windows on the Microsoft Windows host. The Snare Agent pushes event data from the servers to MARS in near real time, when an audit event occurs, the agent sends a syslog message to MARS that details the event. It is also more efficient and timely in that each Snare Agent is able to act independently rather than being bound by a single process as with the pull method.

The following sections describe these two methods:

- [Push Method: Configure Generic Microsoft Windows Hosts, page 10-5](#)
- [Pull Method: Configure the Microsoft Windows Host, page 10-6](#)

Push Method: Configure Generic Microsoft Windows Hosts

MARS can treat hosts running Microsoft Windows as reporting devices, monitoring the event log data generated by the host. The host needs to run InterSect Alliance SNARE Agent for Windows, which captures event log data and sends it to MARS. The push method requires four steps:

1. Install the SNARE agent on the Microsoft Windows host. For more information, see [Install the SNARE Agent on the Microsoft Windows Host, page 10-5](#).
2. Configure the SNARE agent to forward event data to the MARS Appliance. For more information, see [Enable SNARE on the Microsoft Windows Host, page 10-6](#)
3. Ensure that UDP 514 traffic can pass between the hosts and the MARS Appliance.
4. Identify that host in MARS so that it can correctly parse and correlate the event data. For more information, see [Configure the MARS to Pull or Receive Windows Host Logs, page 10-8](#).

Install the SNARE Agent on the Microsoft Windows Host

To install the SNARE agent, follow these steps:

-
Step 1	Log in to the target host using a username with proper administrative privileges. The username must have the permission to publish audit data as well as to install new programs.
Step 2	Download the SNARE Agent for Windows from the following URL that corresponds to the operating system type installed on the target host: http://www.intersectalliance.com/projects/SnareWindows/index.html#Download
Step 3	Double-click the SnareSetup<version>.exe file to start the install program.
Step 4	Click Next .
Step 5	Select the target install folder and click Next .
Step 6	Select Normal Installation in the Components list and click Next .
Step 7	Select the target Start menu location and click Next .
Step 8	Verify the selection options and click Install .
Step 9	SNARE is installed and started on the local host. A dialog box appears, prompting you to specify whether to allow SNARE to control the EventLog configuration for the Microsoft Windows host.
Step 10	Select Yes to enable SNARE to control the EventLog configuration for this Microsoft Windows host. The SNARE - Remote Event Logging for Windows user interface appears.

Step 11 To configure the Snare agent, continue with [Enable SNARE on the Microsoft Windows Host, page 10-6](#).

Enable SNARE on the Microsoft Windows Host

Once you have downloaded and installed the SNARE agent on the target Microsoft Windows host, you must configure the agent to forward the correct event data in the correct format to the MARS Appliance.



Note

The first time you install SNARE, a dialog box appears asking “Do you want Snare to take over control of your Event Log?” Select **Yes**.

To configure the SNARE agent, follow these steps:

-
- Step 1** Click **All Programs > InterSect Alliance > Snare for Windows** to run the SNARE - Remote Event Logging for Windows user interface.
- Step 2** The first time you install SNARE, a dialog box appears asking “Do you want Snare to take over control of your Event Log?” Select **Yes**.
- Step 3** Click **Setup > Audit Configuration...**
The Audit Configuration dialog box appears.
- Step 4** Specify values for the following fields:
- **Enter the local host name.** Specify the IP address or DNS name of the local host in the field.
 - **Enter the snare server ip or dns addr.** Specify the IP address or the DNS name of the MARS Appliance.
- Step 5** Verify that the following options are selected:
- **Enable SYSLOG header**
 - **Automatically set audit configuration**
 - **Automatically set file system audit configuration**
- Step 6** Click **OK** to close the Audit Configuration dialog box and save your changes.
- Step 7** Click **File > Exit** to close the SNARE - Remote Event Logging for Windows user interface.
The Snare agent is stopped and restarted to pick up the configuration changes.
-

Pull Method: Configure the Microsoft Windows Host

As an alternative to the push method, you can configure MARS to pull event log data (security, application, and system event logs) from Microsoft Windows hosts. The pull method requires four steps:

1. Ensure that the Windows host and MARS Appliance clocks are synchronized. It is recommend that you configure a NTP server for this purpose. For more information, see [Specify the Time Settings, page 5-10](#).
1. Select an existing or define a new user account on the Windows host that the MARS Appliance can use to pull event log records.

2. Ensure that the user account has the correct credentials. Verify that the user account belongs to the Administrator group and verify that it includes the privilege for managing and auditing security logs. For more information, see the procedure that corresponds to the operating system running on the host:
 - [Enable Windows Pulling Using a Domain User, page 10-7](#)
 - [Enable Windows Pulling from Windows NT, page 10-7](#)
 - [Enable Windows Pulling from a Windows 2000 Server, page 10-7](#)
 - [Enable Windows Pulling from a Windows Server 2003 or Windows XP Host, page 10-8](#)
3. Configure the Windows host to generate the correct event data.
4. Identify that host in MARS so that it can correctly parse and correlate the event data. For more information, see [Configure the MARS to Pull or Receive Windows Host Logs, page 10-8](#).
5. Specify the time interval at which the event log data should be pulled from all identified host running Microsoft. For more information, see [Windows Event Log Pulling Time Interval, page 10-10](#).

Enable Windows Pulling Using a Domain User

To enable Windows pulling using a domain user (`domain\username`), for example, `CORP\syslog`, do the following on the domain controller *before* you enable Windows pulling on your client:

-
Step 1	On the domain controller, click Administrative Tools > Default Domain Security Policy > Security Settings > Local Policies > User Rights Management .
Step 2	Grant the permission Manage auditing and security log to the domain user (<code>domain\username</code>).
-

Enable Windows Pulling from Windows NT

To enable MARS to pull event log data from a Windows NT host, follow these steps:

-
Step 1	From Start > Programs > Administrative Tools > User Manager , in the menu bar, choose Policies .
Step 2	In the submenu, choose User Rights , make sure the right of Manage auditing and security log is granted to the user account used for pulling event log records.
Step 3	In the submenu, choose Audit . Configure the audit policy according to your site's security auditing policy.
-

Enable Windows Pulling from a Windows 2000 Server

When there is no Active Directory Service (ADS) server sending domain information to your Windows 2000 server, you must set this property to *Disabled* on each host from which you want the MARS Appliance to pull syslogs.

To enable MARS to pull event log data from a Windows 2000 host, follow these steps:

-
Step 1	Go to Start > Settings > Control Panel > Administrative Tools > Local Security Policy .
-

The Local Security Settings applet appears.

Step 2 Configure the settings under the following Local Policy groups as specified:

- Security Settings > Local Security Policy > User Rights Management

Make sure the right of **Manage auditing and security log** is granted to the user account used for pulling event log records.

- Security Settings > Local Security Policy > Audit Policy

Configure the audit policy according to your site's security auditing policy and ensure that all entries under Effective Setting are set to **Success, Failure**.

Enable Windows Pulling from a Windows Server 2003 or Windows XP Host



Note

If you are selecting Microsoft Windows XP Home Edition, you must enable the Remote Procedure Call services under All Programs > Control Panel > Administrative Tools > Services. This service is enabled by default on Windows XP Professional.

To enable MARS to pull event log data from a Windows Server 2003 or Windows XP host, follow these steps:

Step 1 Go to **Start > Settings > Control Panel > Administrative Tools > Local Security Policy**.

The Local Security Settings applet appears.

Step 2 Configure the settings under the following Local Policy groups as specified:

- Security Settings > Local Security Policy > User Rights Management

Make sure the right of **Manage auditing and security log** is granted to the user account used for pulling event log records.

- Security Settings > Local Security Policy > Audit Policy

Configure the audit policy according to your site's security auditing policy.



Note

The pulling of an event log itself generates security event logs if certain events, such as **Log on/off**, are audited. We recommend you either set a default domain policy, or set the retention method for security event logs on your Windows system to be **Overwrite as needed**. Otherwise, when the log is full no new event log can be generated on the Windows system.

Configure the MARS to Pull or Receive Windows Host Logs

Once you've prepared the Microsoft Windows host, you must identify that host in MARS and identify whether the push or pull method is being used on that host.

To configure the MARS Appliance to either pull or receive logs, follow these steps:

Step 1 Select **Admin > Security and Monitor Devices > Add**

- Step 2** From the **Device Type** list, select **Add SW Security apps on a new host** or **Add SW security apps on existing host**.
- Step 3** Enter the **Device Name** and **IP addresses** if adding a new host.
- Step 4** Select the **Operating System > Windows** from the list.
- Step 5** (Optional) Enter **NetBIOS name**.

Figure 10-3 Window Log Configuration

The screenshot shows the 'Window Log Configuration' window with three tabs: 'General', 'Reporting Applications', and 'Vulnerability Assessment Info'. The 'General' tab is active. It contains the following fields and controls:

- *Device Name:** Text field with 'Softie III' entered.
- Access IP:** Four text fields with '192', '168', '2', and '5' entered.
- Reporting IP:** Four text fields with '192', '168', '2', and '5' entered.
- Operating System:** Dropdown menu with 'Windows' selected. A 'Logging Info' button is next to it.
- NetBIOS Name:** Text field with 'netBIOS_Name' entered.
- Monitor Resource Usage:** Dropdown menu with 'NO' selected.
- Enter interface information:** A section containing:
 - 'Add Interface' and 'Remove Interface/IP' buttons.
 - Fields for 'Name:', 'IP Address:', and 'Network Mask:'.
 - A checkbox next to 'ether0'.
 - Four text fields for IP address (192, 168, 2, 5) and four for network mask (255, 255, 255, 255).
 - An 'Add IP/Network Mask' button.

143263

- Step 6** Click on **Logging Info** to configure OS Logging Information. New pop-up window will appear.
- Step 7** From the Windows Operating System, select the correct option for either the server or workstation version:
- Microsoft Windows 2000
 - Microsoft Windows 2003 (Also used for Microsoft Windows XP platforms.)
 - Microsoft Windows Generic
 - Microsoft Windows NT



Note

If you are selecting Microsoft Windows XP Home Edition, you must enable the Remote Procedure Call services under All Programs > Control Panel > Administrative Tools > Services.

- Step 8** Select either the **Pull** or the **Receive** checkbox, based on the host configuration that you have performed.



Caution

Do not select both checkboxes. Doing so generates unpredictable results.

- Step 9** If you selected the Pull method, enter values for the following fields:

- **Domain name**—Identifies the domain name to which the host belongs.
- **Host login**—Identifies the username with security audit and log permissions.
- **Host password**—Identifies that password that authenticates the username provided in the Host login field.

Step 10 Click **Submit**.

Figure 10-4 Windows Logging

OS Logging Information

Windows Operating System: Microsoft Windows 2003

Logging mechanism: ☒ Pull ☒ Receive

Domain Name: my_domain

Host login: username

Host password:

Cancel Submit

Step 11 Click **Submit** to save your changes.

Step 12 Add Interface IP Address and Network Mask.

Step 13 Click **Apply**.

Step 14 Click the **Vulnerability Assessment Info** link to define the host information that MARS uses to determine false positive attacks against this host. Continue with [Define Vulnerability Assessment Information](#), page 10-11.

Step 15 Click **Done** to save the changes.

Step 16 To activate the device, click **Activate**.

If you selected the pull check box in [Step 8](#), verify that a value has been specified for the interval at which which MARS pulls an event log from the host. For more information, see [Windows Event Log Pulling Time Interval](#), page 10-10.

Windows Event Log Pulling Time Interval

You can now set the interval at which MARS pulls an event log from all Microsoft Windows host that are defined as reporting devices. This feature determines how often MARS requests logs from the Windows hosts that are configured a reporting devices.



Note

If you are using SNARE to push the log data to MARS, then you do not need to enable this setting.

To configure the Windows event log pulling time interval, follow these steps:

Step 1 Click **Admin > System Parameters > Windows Event Log Pulling Time Interval**.

Windows Event Log Pulling Time Interval

Step 2 Enter the new time interval in seconds. The default value is 300 seconds (five minutes).

Step 3 Click **Submit**.

Define Vulnerability Assessment Information

For each host that you define in MARS, you can specify information about that host that assists MARS in assessing whether that host is vulnerable to the attacks that MARS detects. For example, you can identify the operating system running on the host, even providing the latest or nearest patch level. When an attack is detected that is targeted toward a specific operating system, then MARS can quickly determine whether the host is running the operating system that is targeted.

For hosts that are defined as the base platform of a reporting device, you should define this information as part of that device definition.

However, as MARS, it begins to add discovered hosts to the list of hosts under **Management > IP Management**. You should periodically review these hosts to update their information if you do not have a vulnerability assessment software device or service, such as Qualys QualysGuard, running on your network.

To specify the vulnerability assessment information for a host, follow these steps:

Step 1 To select the desired host, do one of the following:

- Select **Management > IP Management**, select the check box next to the desired host, and click **Edit**.
- Select **Admin > Security and Monitor Devices**, select the check box next to the desired host, and click **Edit**.

Step 2 Click the **Vulnerability Assessment Info** tab.

Figure 10-5 Vulnerability Assessment Info for a Host

General **Vulnerability Assessment Info**

Specify OS and patch Information

☒ Select operating system from:

Microsoft Windows 2000(version: 5.0.4.2195,patch: SP 4) ☒ Allow Overwrite with VA

☐ Define new operating system:

Name: Version:

Patch: Vendor:

Current running services:

Step 3 Under Specify OS and patch Information, do one of the following:

- Select **Select operating system from**, and then select the operating system that matches the one running on this host from the list. Continue with [Step 4](#).
- Select **Define new operating system**, and continue with Step [a.](#)
 - a. Enter the name of the operating system in the Name field.
 - b. Enter the version number for this operating system in the Version field.
 - c. Enter the patch level associated with the version number the Patch field.
 - d. Enter the name of manufacturer of the operating system in the Vendor field.
 - e. Click **Apply** to save the operating system definition.

Result: The new operating system definition is added to the Select operating system from list, and it is the selected option.

If you define a custom operating system, you must select **Generic** in the Operating System list on the General page of the host and click **Apply**. Otherwise, you cannot select the new operating system in the Select operating system from list.

Step 4 To allow the information that you provided to be overridden by a vulnerability assessment service running on your network, select the **Allow Overwrite with VA** checkbox.

Step 5 To add more detailed information about the host, continue with [Identify Network Services Running on the Host, page 10-13](#).

Step 6 Click **Apply** to save the changes made to this host.

Step 7 Click **Done** to close the Host page

Identify Network Services Running on the Host

By identifying the network services that are running on a host, you are specifying the types of network activities that you expect for this host. This data is helpful in eliminating expected activities that might otherwise be flagged as suspicious by MARS; for example, if you have administrative servers that run network discovery applications or perform vulnerability assessment probes at scheduled times.

To identify the network services running on a host, follow these steps:

-
- Step 1** To select the desired host, do one of the following:
- Select **Management > IP Management**, select the check box next to the desired host, and click **Edit**.
 - Select **Admin > Security and Monitor Devices**, select the check box next to the desired host, and click **Edit**.
- Step 2** Click **Add New Service** under Current running services.
-  **Note** It may take five minutes or more for this dialog box to load. You can place the cursor over the title bar of the window that opens. This allows you to see if the window is still loading.
-
- Step 3** Enter as much detail on the service and its applications as you can.
- You can choose between selecting a service and defining a new service.
 - You can also choose between select an application or defining a new application.
- Step 4** Click **Submit**.
- Step 5** You can enter more services here by clicking **Add New Service**, or you can click **Submit** to continue.
- Step 6** Click **Submit** to complete the addition of the host.
-



Configuring Database Applications

Database applications are typically high-value assets, and as such, they are common targets for attacks. Database applications provide MARS with user activity, such as successful and failed login attempts, session durations, and activities indicative of privilege escalation.

This chapter explains how to bootstrap and add the following database applications to MARS:

- [Oracle Database Server Generic, page 11-1](#)

Oracle Database Server Generic

To configure CS-MARS to collect information from the Oracle database server, you must perform three tasks:

- configure the Oracle database server to generate a audit trail and record those events the database.
- represent the device in the web interface
- configure the interval at which CS-MARS should pull the logs from the Oracle database server.

Configuring the pull interval is a one-time operation that applies to all of the Oracle database servers monitored by the MARS Appliance.

This section contains the following topics:

- [Configure the Oracle Database Server to Generate Audit Logs, page 11-1](#)
- [Add the Oracle Database Server to MARS, page 11-2](#)
- [Configure Interval for Pulling Oracle Event Logs, page 11-3](#)

Configure the Oracle Database Server to Generate Audit Logs

You must configure the Oracle database server to write audit logs to the database. You may need your DBA support to perform most of these configurations. Once configured, MARS can retrieve the audit logs from your Oracle database server. The following examples are for an Oracle instance running on a UNIX/Linux application host.

To configure an Oracle database server to write audit logs, follow these steps:

Step 1 As sysdba execute cataudit.sql to create audit trail views:

```
[oracle@server]$ sqlplus /nolog
```

```
SQL> conn / as sysdba;
SQL> @$ORACLE_HOME/rdbms/admin/cataudit.sql
```

- Step 2** Enable auditing to the database by adding the following entry to the Oracle instance initialization file, usually named init<SID>.ora

```
AUDIT_TRAIL=DB
```

This file is usually located in \$ORACLE_BASE/admin/<SID>/pfile, where <SID> is the name of the Oracle instance.

If a binary initialization file is used for this instance, make sure you update it first. This file is usually located in \$ORACLE_HOME/dbs and named spfile<SID>.ora. Ask your DBA about the location of these files as well as the policies applied for this server.

- Step 3** Restart the database to activate the change made to the initialization file.

```
[oracle@server]$ sqlplus /nolog

SQL> conn / as sysdba;
SQL> shutdown immediate;
SQL> startup;
```

- Step 4** Turn on all the logs that you want to audit. The following example is turning on the “audit session”.

```
SQL> audit session;
Audit succeeded.
```

- Step 5** Repeat the previous step for all the logs that you want to audit.

- Step 6** Create a user account on this server and grant select privilege for the view dba_audit_trail. Our example assumes the user has login name “pnuser”.

```
SQL> grant select on dba_audit_trail to pnuser
```

You’ll use “pnuser” as the value for “User Name” in the MARS setup.

- Step 7** To test that everything was properly configured, audit logs are written to the database and “pnuser” has read access to them, execute the following commands:

```
[oracle@server]$ sqlplus pnuser/<password>@<oracle_server>

SQL> select count(*) from dba_audit_trail;

COUNT(*)
-----
          3
```

If the above count is anything but zero, congratulations, you have successfully configured the Oracle Server! You will have to repeat the above procedure for every Oracle server that you want to report audit logs to MARS.

Add the Oracle Database Server to MARS

To represent the Oracle database server in the web interface, follow these steps:

- Step 1** Click **Admin > Security and Monitor Devices > Add**.

- Step 2** From the **Device Type** list, select **Add SW Security apps on a new host** or **Add SW security apps on existing host**.
- Step 3** Enter the **Device Name** and **IP addresses** if adding a new host.
- Step 4** Click **Apply**.
- Step 5** From the **Select Application** list, select **Oracle Database Server Generic**.
- Step 6** Click **Add**.

- Step 7** Enter the **User Name**, **Password** and **Oracle Service Name**
- **User Name** – the Oracle Database User Name
 - **Password** – the Oracle Database User password
 - **Oracle Service Name** – the Oracle Service Name

The Oracle Service Name is the GLOBAL_DBNAME=username.server which can be found inside a file called listener.ora.

- Step 8** Click **Test Connectivity** to verify the configuration.
- Step 9** Click **Submit**.

Configure Interval for Pulling Oracle Event Logs

To specify the interval at which MARS should pull the event logs from all Oracle database servers on your network, follow these steps:

- Step 1** Click **Admin > System Parameters > Oracle Event Log Pulling Time Interval**.

Oracle Event Log Pulling Time Interval

Oracle Event
Log Pulling
Time
Interval:

(secs)

⏪ Back

Submit

143252

- Step 2** Enter the new time interval in seconds. The default value is 300 (five minutes).
- Step 3** Click **Submit**.
-



Configuring Web Server Devices

To use web logging with MARS, you need to configure the host, the webserver, and MARS. MARS can process up to 100 MB of web log data per receive from your host.



Note

Web logging is only supported on hosts running Microsoft IIS on Windows, Apache on Solaris or Linux, or iPlanet on Solaris.

This chapter explains how to bootstrap and add the following web sever devices to MARS:

- [Microsoft Internet Information Sever, page 12-1](#)
- [Apache Web Server on Solaris or RedHat Linux, page 12-7](#)
- [Sun Java System Web Server on Solaris, page 12-7](#)

Microsoft Internet Information Sever

You can add computers running Microsoft Windows to MARS as reporting devices. The Microsoft Windows computer needs to run [InterSect Alliance SNARE for IIS](#), from which MARS receives web log data.



Note

Synchronize clocks of the Microsoft Windows system and the MARS to ensure times match between them.

Install and Configure the Snare Agent for IIS


To configure IIS to publish logs to MARS, you must install and configure a log agent. This agent is free from the InterSect Alliance. You can download the Snare Agent for IIS Servers from the following URL:

<http://www.intersectalliance.com/projects/SnareIIS/index.html#Download>

After you have downloaded and install the SNARE on the the Windows webserver, you can continue with the procedures in this section that detail the correct configuration for MARS,

To configure SNARE for web logging, follow thees steps:


Step 1 Click **Start > Programs > InterSect Alliance > Audit Configuration**.

Figure 12-1 *Configure SNARE for Web Logging*

- Step 2** In **Target Host** enter the IP address of the MARS.
- Step 3** In **Log Directory**, enter the directory where the logs are to be placed.
- Step 4** In **Destination**, click the **Syslog** radio button.
- Step 5** Click **OK**.


To configure IIS for web logging

- Step 1** Click **Start > Programs > Administrative Tools > Internet Services Manager**.

Figure 12-2 *Configure IIS for Web Logging*


Step 2 In the **Tree** tab on the left, right-click **Default Web Site**.

Step 3 On the shortcut menu, select **Properties**.

Figure 12-3 *Enable Logging*

- Step 4** In the **Web Site** tab:
- Make sure **Enable Logging** is checked.
 - From the **Active log format** list, select **W3C Extended Log Format**.
 - Click **Properties**.

Figure 12-4 Select General Log Settings




- In the **General Properties** tab, set the **New Log Time Period** to **Daily**.



Note

The **Log file directory** *must* match the one previously set using the **Audit Configuration** program.

- In the **Extended Properties** tab, make sure all available properties are selected.

Figure 12-5 **Select Extended Log Events**

f. Click **OK**.

Step 5 Click **OK**.

MARS-side Configuration

To add configuration information for the host

- Step 1** Click **Admin > Security and Monitor Devices > Add**
- Step 2** From the **Device Type** list, select **Add SW Security apps on a new host** or **Add SW security apps on existing host**
- Step 3** Enter the **Device Name** and **IP Addresses** if adding a new host.
- Step 4** Select the **Windows** from **Operation System** list
- Step 5** Click **Logging Info**
- Step 6** For this configuration, you *must* check the **Receive host log** box

Figure 12-6 Windows Web Server Logging mechanisms

OS Logging Information

Windows Operating System: Microsoft Windows 2003

Logging mechanism: ☒ Pull ☒ Receive

Domain Name: my_domain

Host login: username

Host password:

Cancel Submit

143262

Step 7 Click **Submit**.**Step 8** Continue adding the interfaces.

- For the first interface, enter its name, IP address, and mask.
- For multiple interfaces, click **Add Interface**, and add each new interface's name, IP address, and mask.

Step 9 Add as many IP addresses and masks to the interface as you need by clicking **Add IP/Network Mask**.**Step 10** Click **Apply**.**Step 11** Click **Reporting Applications** tab.**Step 12** From the **Select Application** list, select **Generic Web Server Generic**.**Step 13** Click **Add**.**Figure 12-7 Selecting the Windows Web Log format**

Web log format: None

None

W3C_EXTENDED_LOG

Cancel Submit

143268

Step 14 Select **W3C_EXTENDED_LOG** format**Step 15** Click **Submit**.**Note**

Once you have configured and activated both sides, it takes two pulling intervals (default time of 10 minutes) before new events appear.

Apache Web Server on Solaris or RedHat Linux

Sun Java System Web Server on Solaris

**Note**

The Sun Java System Web Server was formerly known by the following product names: Netscape Enterprise Server, iPlanet Web Server, and Sun ONE Web Server,

Generic Web Server Generic

You can add computers running Solaris or Linux to MARS as reporting devices. The computer needs to run an opensource agent that sends web log data to MARS.

Solaris or Linux-side Configuration

Cisco provides an opensource logging agent and an associated configuration file for you to use. This agent can be downloaded from the software download center at the following URL:

<http://www.cisco.com/cgi-bin/tablebuild.pl/cs-mars-misc>

**Note**

Synchronize clocks of the UNIX or Linux system and the MARS to ensure times match between them.

Install and Configure the Web Agent on UNIX or Linux

For MARS to receive logs from a webserver, you must install the Web agent, (agent.pl version 1.1) on the target webserver and direct the agent to publish logs to the MARS Appliance.

**Note**

Before you install the agent, you must have **perl** and **curl** installed on your system.

To install the agent on a UNIX or Linux hosts, follow these steps:

-
- Step 1** Log into the host as the root user.
- Step 2** Create a directory called `/opt/webagent`.
- Step 3** Copy the files `agent.pl` and `webagent.conf` to the `/opt/webagent` directory.
- Step 4** Set the protection of the agent script (`agent.pl`) so it can be read and executed by all:
- ```
cd /opt/webagent
chmod 755 agent.pl
```
- Step 5** Edit the configuration file (`weblogagent.conf`):
- ```
logfile_location = access_log_path
MARS_ip_port = MARS_ip_address
username = a
```

```
password = b
```

Where the following values are provided:

- *access_log_path* identifies the absolute path name to the web server's access log
- *MARS_ip_address* is the IP address of the MARS Appliance

You do not need to edit the username or password in the file.


Note

You need a separate `weblogagent.conf` file for each access log you want to pull. We recommend naming them `weblogagent1.conf`, `weblogagent2.conf`, and so forth. Put these in the `/opt/webagent` directory also.

To run the agent using a configuration file other than `weblogagent.conf`, use the command:

```
agent.pl other_config_file
```

replacing *other_config_file* with the name of the web agent configuration file.

- Step 6** Edit the crontab file to push the logs to the MARS at regular intervals. The following example gets new entries from the access log and pushes them to MARS every five minutes:

```
crontab -e
5,10,15,20,25,30,35,40,45,50,55,0 * * * *
(cd /opt/webagent; ./agent.pl weblogagent1.conf)
5,10,15,20,25,30,35,40,45,50,55,0 * * * *
(cd /opt/webagent; ./agent.pl weblogagent2.conf)
```

Web Server Configuration

To configure the Apache web server for the agent

- Step 1** In the file `httpd.conf`, make sure the `LogFormat` is either common or combined *and* matches the format set on the MARS.
- Step 2** Stop and restart the Apache server for your changes to take effect.

To configure the iPlanet web server for the agent

- Step 1** In the iPlanet server administration tool, click the **Preferences** tab.
- Step 2** In the left menu, click the **Logging Options** link.
- Step 3** Make sure the **Log File** matches the log file name set on the MARS.
- Step 4** Make sure the **Format** radio button **Use Common Logfile Format** is checked.
- Step 5** If you have made any changes, click **OK**.
- Step 6** If necessary, shut down and restart the iPlanet web server.

MARS-side Configuration

To add configuration information for the host

- Step 1** Click **Admin > Security and Monitor Devices > Add**.
- Step 2** From the **Device Type** list, select **Add SW Security apps on a new host** or **Add SW security apps on existing host**.
- Step 3** Enter the **Device Name** and **IP Addresses** if adding a new host.
- Step 4** Select the either **Solaris** or **Linux** from **Operation System** list.
- Step 5** Click **Logging Info**.
- Step 6** For this configuration, you *must* check the **Receive host log** box.

Figure 12-8 Unix or Linux Web Server Logging mechanism

OS Logging Information

Logging mechanism: ☐ Pull ☒ Receive

Host login:

Host password:

143267

- Step 7** Click **Submit**.
- Step 8** Continue adding the interfaces.
- For the first interface, enter its name, IP address, and mask.
 - For multiple interfaces, click **Add Interface**, and add each new interface's name, IP address, and mask.
- Step 9** Add as many IP addresses and masks to the interface as you need by clicking **Add IP/Network Mask**.
- Step 10** Click **Apply**.
- Step 11** Click **Reporting Applications** tab.
- Step 12** From the **Select Application** list, select **Generic Web Server Generic**.
- Step 13** Click **Add**.

Figure 12-9 Linux Operating System Web Log Format

Web log format:

None
COMMON_ACCESS_LOG/COMBINED_LOG
SQUID_LOG
NETSCAPE_EXTENDED_LOG
NETCACHE_WEB_ACCESS_DEFAULT_LOG
W3C_EXTENDED_LOG

143264

Step 14 From the **Web Log Format** list, select appropriately.

Step 15 Click **Submit**.

**Note**

Once you have edited a device you must click **Activate** for the changes to take effect.



Configuring Web Proxy Devices

Web proxy devices provide MARS with additional data surrounding user requests of network services, such as HTTP, FTP, NNTP, and DNS. These device cache data and provide additional services around requests for that data. These additional services provide MARS with data about session requests, including authentication logs, denied session requests based on ACLs enforced by the web proxy device, and traffic logs.

This chapter contains the following section:

- [Network Appliance NetCache Generic, page 13-1](#)

Network Appliance NetCache Generic

This section contains the following topics:

- [Configure NetCache to Send Syslog to MARS, page 13-1](#)
- [Add and Configure NetCache in MARS, page 13-2](#)

Configure NetCache to Send Syslog to MARS

Synchronize clocks of the NetCache device and the MARS to make sure times match between them.



Note

MARS supports only HTTP proxy logs and MMS streaming media proxy logs.

To configure NetCache to send syslog to MARS, follow these steps:

-
Step 1	In Internet Explorer, enter the URL and log in to the NetCache device.
Step 2	Click the Setup tab.
Step 3	In the left side of the window, select HTTP , then Logging .
Step 4	In the right side of the window, under Web Access Log Enable , select the Enable the Web Access Log checkbox.
Step 5	Under Log Format , select one of the first four formats: • Web Access Log Default Format • Common Log Format

- Netscape Extended Format
- Squid Type Format

Step 6 Under Web Details Log Enable, verify the box is *not* selected.

Step 7 Click **Commit Changes** to save your changes.

Step 8 In the left side of the window, select **Streaming**, then **Logging**.

Step 9 In the right side of the window, under Streaming Access Log Enable, select the **Enables access logging for streaming protocol clients** check box.



Note

You can only enable access logging for streaming protocol clients if you have a streaming cache license.

Step 10 Under Streaming Access Log Format, select either of the options. If you select **Custom**, replace “x-client-port” with “x-username”.

Step 11 Under Streaming Details Log Enable, verify that the box is *not* selected.

Step 12 Click **Commit Changes** to save your changes.

Step 13 In the left side of the window, select **Streaming**, then **MMS**.

Step 14 Under **MMS Enable**, verify that the **Enables MMS protocol support** check box is selected.

Step 15 Click **Commit Changes** to save your changes.

Step 16 In the left side of the window, select **System**, then **Logging**.

Step 17 In the right side of the window, under Maximum Log File Size, enter a number less than or equal to 100 (megabytes).

Step 18 Under How to Switch Log Files, select **Push the log file to the following URL**.

Step 19 For the URL, enter:

`http://MARS_HOST/upload/UploadWebLogServlet`

Replace *MARS_HOST* with the hostname or IP address of the MARS Appliance.

Step 20 Verify that the User and Password fields are blank.

Step 21 Verify that the **Push the log files in compressed gzip format** check box is not selected.

Step 22 Under When to Switch, select the option that prevents the log files from becoming greater than 100 megabytes.

Step 23 Click **Commit Changes** to save your changes.

Add and Configure NetCache in MARS

To add the NetCache device in MARS, follow these steps:

Step 1 Select **Admin > Security and Monitor Devices > Add**.

Step 2 From the Device Type list, select **Network Appliance NetCache Generic**.

Device Type: Network Appliance NetCache Generic ▼

→ *Device Name:

→ *Reporting IP: ---

→ Web log format: ▼

→ Streaming media log format: ▼

COMMON_ACCESS_LOG
SQUID_LOG
NETSCAPE_EXTENDED_LOG
NETCACHE_WEB_ACCESS_DEFAULT_LOG

143266

- Step 3** Enter the device name and its reporting IP address.
- Step 4** From the Web log format list, select the web log format that matches the value you selected in [Step 5](#) of [Configure NetCache to Send Syslog to MARS, page 13-1](#).
- Step 5** From the Streaming media log format list, select a streaming media log format.
- Step 6** Click **Submit**.
-



Configuring AAA Devices

Authentication, authorization, and accounting (AAA) devices provide accountability throughout your network, ensuring that valid users are authorized to use the network services they request and providing detailed event logs regarding failures and successes in such requests.

The AAA server is a key component in the Network Access Control (NAC) initiative (see [Configuring Network Admission Control Features, page 2-42](#) and [Enable NAC-specific Messages, page 3-4](#)). Cisco Secure Access Control Server (ACS), which is the AAA server for NAC, returns access control decisions to the network access device on the basis of the antivirus credentials of the hosts that are requesting network services.

MARS supports the Cisco Secure ACS software and the Cisco Secure ACS Solution Engine, version 3.3 and later. In the case of Cisco Secure ACS software, support is provided by an agent that resides on the Cisco Secure ACS server. For the Cisco Secure ACS Solution Engine, this agent must reside on a remote logging host. This agent provides MARS with three event logs in syslog format. The logs are as follows:

- Passed authentication log (requires Cisco Secure ACS, 3.3 or later)
- Failed attempts log
- RADIUS accounting log

To support NAC and the 802.1x features, Cisco Secure ACS uses the RADIUS authentication protocol and the cisco-av-pair attributes. For more information on configuring Cisco Secure ACS as a posture validation server for NAC, see the following URL:

http://www.cisco.com/en/US/products/sw/secursw/ps2086/products_user_guide_chapter09186a00802335f1.html

For more information on the cisco-av-pair attributes, see the following URL:

http://www.cisco.com/en/US/products/sw/secursw/ps2086/products_user_guide_chapter09186a00802335ea.html

This chapter explains how to prepare the Cisco Secure ACS server or the Cisco Secure ACS Solution Engine to allow MARS to collect the event logs. It also describes how to configure MARS to receive and process these logs correctly. Using the web interface, you must define a host to represent the Cisco Secure ACS server (or the remote logging agent collecting logs for the Cisco Secure ACS Solution Engine) and then add the software application to that host.

Supporting Cisco Secure ACS Server

To configure a Cisco Secure ACS server to act as a reporting device, you must perform three tasks:

1. Configure Cisco Secure ACS server to generate the correct log files and details and define the AAA clients.
2. Install the PN Log Agent on the Cisco Secure ACS server and configure it to forward the correct log files.
3. Add the Cisco Secure ACS server to the MARS web interface

You can also configure Cisco Secure ACS to provide command authorization for the MARS Appliance. In this role, Cisco Secure ACS verifies that the MARS Appliance is authorized to execute specific commands on reporting devices and mitigation devices.

The following sections detail supporting a Cisco Secure ACS server:

- [Bootstrap Cisco Secure ACS, page 14-2](#)
- [Install and Configure the PN Log Agent, page 14-7](#)
- [Add and Configure the Cisco ACS Device in MARS, page 14-12](#)

Supporting Cisco Secure ACS Solution Engine

MARS supports the Cisco Secure ACS Solution Engine via a remote logging host. Cisco Secure ACS Remote Agent for Windows and Cisco Secure ACS Remote Agent for Solaris are applications that support Cisco Secure ACS Solution Engine for remote logging.

Even though the Cisco Secure ACS Solution Engine supports up to five appliance via a remote logging host, MARS currently supports only one Cisco Secure ACS Solution Engines per remote logging host. Otherwise, MARS cannot identify the IP address of the originating Cisco Secure ACS Solution Engine.

To enable this support, follow these steps:

1. Configure the Cisco Secure ACS Solution Engine to publish logs to the remote logging host. See [Bootstrap Cisco Secure ACS, page 14-2](#).
2. Install and configure either the Cisco Secure ACS Remote Agent for Windows and Cisco Secure ACS Remote Agent for Solaris on the target remote logging host.

For instructions on installing and configuring the remote agent, see [Installation and Configuration Guide for Cisco Secure ACS Remote Agents](#).

3. Install the pnLog Agent on the remote logging host.

For information on installing and configuring the pnLog Agent, see [Install and Configure the PN Log Agent, page 14-7](#).

4. Add the remote logging host to MARS as a Cisco ACS 3.x reporting device. To perform this task see [Add and Configure the Cisco ACS Device in MARS, page 14-12](#), and substitute the ACS server references with the remote logging host.

Bootstrap Cisco Secure ACS

Bootstrapping the Cisco Secure ACS includes the following task:


- [Configure Cisco Secure ACS to Generate Logs, page 14-3](#)

- [Define AAA Clients, page 14-5](#)
- (Optional) [Configure TACACS+ Command Authorization for Cisco Routers and Switches, page 14-6](#)


Configure Cisco Secure ACS to Generate Logs

To configure Cisco Secure ACS to generate the audit logs required by MARS, follow these steps:

- Step 1** Log in to the Cisco Secure ACS server or Solution Engine.
- Step 2** Select **System Configuration > Logging**.



- Step 3** Verify that CVS Failed Attempts, CVS Passed Authentications and CVS RADIUS Accounting Logging are enabled.
- Step 4** Click **CSV Failed Attempts**, and verify that the following attributes appear in the Logged Attributes list:
- User-Name
 - Caller-ID
 - NAS-Port
 - NAS-IP-Address



Step 5 Click **Submit**

Step 6 Click **CVS Passed Authentications**, and verify that the following attributes appear in the Logged Attributes list:

- User-Name
- Caller-ID
- NAS-Port
- NAS-IP-Address
- System-Posture-Token
- EAP Type Name

Step 7 Click **Submit**.

Step 8 Click **CVS RADIUS Accounting**, and verify that the following attributes appear in the Logged Attributes list:

- User-Name
- Calling-Station-Id
- Acct-Status-Type
- NAS-Port
- NAS-IP-Address

Step 9 To support the 802.1x features of NAC, select the following RADIUS accounting attributes:

- Framed-IP address
- cisco-av-pair

CSV RADIUS Accounting File Configuration

Enable Logging

☒ Log to CSV RADIUS Accounting report

Select Columns To Log

Attributes	Logged Attributes
Service-Type	User-Name
Framed-Protocol	Calling-Station-Id
Login-IP-Host	Acct-Status-Type
Login-Service	Framed-IP-Address
Class	NAS-Port
Termination-Action	NAS-IP-Address
Called-Station-Id	cisco-av-pair
NAS-Identifier	
Proxy-State	
Login-LAT-Service	
Login-LAT-Node	
Login-LAT-Group	
Acct-Delay-Time	
Acct-Input-Octets	
Acct-Output-Octets	
Acct-Session-Id	
Acct-Authentic	
Acct-Session-Time	
Acct-Input-Packets	
Acct-Output-Packets	

Up Down

Step 10 Click **Submit**.

For additional details on the RADIUS attributes supported by Cisco Secure ACS, see to the following URL:

http://www.cisco.com/en/US/products/sw/secursw/ps2086/products_user_guide_chapter09186a00802335ea.html

Define AAA Clients

To support the 802.1x features of NAC, you must also define the Cisco switches as AAA clients within Cisco Secure ACS. When defining a AAA client, verify the following settings:

- RADIUS (IETF) is selected in the Using Authentication box, as other RADIUS implementations may not support 802.1x correctly.
- The Log Update/Watchdog Packets from this AAA Client box is selected.

Figure 14-1 displays the correct settings for such a client.

Figure 14-1 Configure a AAA Client to Support 802.1x

Network Configuration

AAA Client Setup For 802.1xrouter

AAA Client IP Address: 20.1.1.1

Key: protego

Authenticate Using: RADIUS (IETF)

☐ Single Connect TACACS+ AAA Client (Record stop in accounting on failure).

☒ Log Update/Watchdog Packets from this AAA Client

☐ Log RADIUS Tunneling Packets from this AAA Client

☐ Replace RADIUS Port info with Username from this AAA Client

Buttons: Submit, Submit + Apply, Delete, Delete + Apply, Cancel, Back to Help

Help

- [AAA Client IP Address](#)
- [Key](#)
- [Network Device Group](#)
- [Authenticate Using](#)
- [Single Connect TACACS+ AAA Client](#)
- [Log Update/Watchdog Packets from this AAA Client](#)
- [Deleting a AAA Client](#)
- [Renaming a AAA Client](#)
- [Log RADIUS Tunneling Packets from this AAA Client](#)
- [Replace RADIUS Port info with Username from this AAA Client](#)

AAA Client IP Address

Type the IP address information for this AAA client.

If you want to designate more than one AAA client with a single AAA client entry in Cisco Secure ACS, you can specify the IP address for each AAA client to be represented by this AAA client entry. To separate each IP address, press **Enter**.

You can use the wildcard asterisk (*) for an octet in the IP address. For example, if you want every AAA client in your 192.168.13.1 Class C network to be represented by a single AAA client entry, enter 192.168.13.* in the AAA Client IP Address

For more information on defining AAA clients, see the following URL:

http://www.cisco.com/en/US/products/sw/secursw/ps2086/products_user_guide_chapter09186a00802335ef.html#wp342084

Configure TACACS+ Command Authorization for Cisco Routers and Switches

You can use the TACACS+ feature of Cisco Secure ACS to authorize the command sets that MARS is allowed to execute on a reporting device. The use of this feature is not required by MARS. However, if you are using this feature on your routers and switches, you must ensure that MARS is allowed to execute specific commands. Required commands are grouped under two operations: configuration retrieval and mitigation.

The following commands support configuration retrieval:

- all **show** commands
- **changeto system**
- **changeto context** <context_name>
- **enable**
- **page**
- **no page**
- **terminal length 0**

- **terminal pager lines 0**
- **write terminal**

The following commands support mitigation:

- **conf terminal**
- **interface** <interface_name>
- **shutdown**
- **set port disable** <port_name>

For more information on configuring command authorization sets in Cisco Secure ACS, see the following URL:

http://www.cisco.com/en/US/products/sw/secursw/ps2086/products_user_guide_chapter09186a00802335ec.html#wp697557

Install and Configure the PN Log Agent

MARS includes the PN Log Agent to monitor Cisco Secure ACS active log files (failed attempts, passed authentications, and RADIUS accounting). This agent pushes these log files via syslog to MARS. You can download the PN Log Agent from the software download center at the following URL:

<http://www.cisco.com/cgi-bin/tablebuild.pl/cs-mars-misc>



Note

If you are upgrading to a new version of the PN Log Agent, see [Upgrade PN Log Agent to a Newer Version, page 14-9](#).

As part of its operation, the PNLog Agent service writes error and informational message to the Application Log, which can be viewed using the Event Viewer. To learn more about these messages, see [Application Log Messages for the PN Log Agent, page 14-10](#).

To install and configure the PNLog Agent, follow these steps:

Step 1

Download the PN Log Agent and install it on the server running Cisco Secure ACS or on the remote logging host to which the Cisco Secure ACS Solution Engine is publishing its logs.



Note

If installing on a remote logging host, you must have configured the Cisco Secure ACS Remote Agent for Windows and Cisco Secure ACS Remote Agent for Solaris on the target remote logging host.

For instructions on installing and configuring the remote agent, see [Installation and Configuration Guide for Cisco Secure ACS Remote Agents](#).


Step 2

Select **Start > All Programs > Protego Networks > PNLogAgent > Pn Log Agent**

Step 3

Click **Edit > PN-MARS Config**.

Result: The PN Log Agent Configuration dialog box appears.




Step 4 In the MARS IP Address field, enter the address of the MARS Appliance, and click **OK**.

Step 5 Select **Edit > Log File Config > Add**.

Step 6 From the **Edit** pull down menu select **Add**.


Result: The Add/Edit File Details dialog box appears.



Step 7 From the Application Name list, select the **Cisco ACS-Failed Attempts**.

Step 8 Click on the ... button to select the appropriate log where all Cisco Secure ACS logs are stored. In this example after selecting **Failed Attempts** application, be sure to select the matching log file, **Failed Attempts active** log.

Result: The Open dialog box appears.




Step 9 Add all 3 applications and their active log files:

- Failed Attempts active
- Passed Authentications active

- RADIUS Accounting active

Result: The configured files appear in the List of Log Files to Monitor list.



Step 10 Select **File > Activate**.

Upgrade PN Log Agent to a Newer Version

You can determine which version of the PN Log Agent is running on your server by selecting **Help > About** in the PN Log Agent Configuration dialog box. This program is updated independently of the MARS Appliance software updates. Therefore, the version number does not correspond to any release of the MARS Appliance software.



Note

Beginning with the 4.1.3 release of the pnLog agent, the agent requires a minimum of Cisco Security Monitoring, Analysis, and Response System, release 4.1.3 running on the appliances to which it is reporting in order to operate correctly.

To upgrade to the new PN Log Agent from an existing installation, you must perform the following steps:

- Step 1** On the Cisco Secure ACS or syslog server where PN Log Agent is running, uninstall the old agent.
- To uninstall the old agent, click **Start > Control Panel > Add/Remove Programs**.
 - Select **PnLogAgent** in the list of currently installed programs, and click **Remove**.
 - Select **Yes** to confirm the removal.

Step 2 Reboot the server.

Step 3 Install the new agent. You can download this tool from the following URL:

<http://www.cisco.com/cgi-bin/tablebuild.pl/cs-mars-misc>

- Step 4** Re-configure the new agent, specifying the list of files and IP address of the MARS Appliance, etc. For information on configuring the pnLog Agent, see [Install and Configure the PN Log Agent, page 14-7](#).
-

Application Log Messages for the PN Log Agent

The PN Log Agent service writes events to the Application Log of Event Viewer on the Cisco Secure ACS server. The agent, identified in the log messages as PNLogAgentService, writes status messages, such as successful service start and stop. It also writes error messages for incomplete configuration and error conditions, such as when the service is out of memory.

[Table 14-1](#) categories the types of messages that can occur and explains their affects on the PNLog Agent service.

Table 14-1 Possible Application Log Messages for PN Log Agent

Type/Message	Effect on Service/Cause of Error
Fatal Errors	
Failed to start thread to monitor file.	Windows errors or configuration errors that will stop the service
The service failed to monitor the configured file. Please check service privileges.	
The service failed to obtain path from log file name and shall exit the thread now!	
The service failed to get the CS-MARS device IP Address. Please use the PnLogAgent to configure it.	
The service has detected an invalid IP address. Please use the PnLogAgent to configure the correct IP Address for CS-MARS.	
Error	
Network detected to be down while attempting to send syslog message	Network connectivity errors that will cause the service to not send syslog messages, but will keep the service running
Destination network unreachable while attempting to send syslog message	
Network dropped connection on reset condition while attempting to send syslog message	
Connection reset by peer while attempting to send syslog message	
Connection refused by target while attempting to send syslog message	
No route exists to host. Please check the network connectivity	
Attempt to send syslog returned error code: <error_code>	
The log file doesn't have all required attributes. Attribute missing: <missing_attribute>	Error in configuration
The number of attributes in the file header don't match the number of attributes in the value. Hence this log entry shall not be sent to CS-MARS.	
The service detected that the configured file is missing some mandatory header attributes. A list of mandatory attributes is available in the CS-MARS user documentation.	
The service failed to read the file pnWinEvent.dat and will now wait for an update to the configuration.	
Failure in reading from pnWiinEvent.dat. Service will wait for an update	
Warning	
The attribute <attribute_name> has a value that exceeds the CS-MARS limit for an individual attribute value and shall be split.	Warning in case some attribute data in the file exceeds MARS raw message length... MARS will store the data after splitting it into multiple events
Informational	

Table 14-1 Possible Application Log Messages for PN Log Agent (continued)

Type/Message	Effect on Service/Cause of Error
PnLogAgentService started	Informational messages describing expected operations for the service.
PnLogAgentService Exiting	
The service read the configuration and will attempt to process files.	
As the service has no logs configured, it shall wait for an update	
Exiting thread processing file as service stop received!	

Add and Configure the Cisco ACS Device in MARS

To add the host and Cisco Secure ACS software application to MARS, follow these steps:

-
- Step 1** Click **Admin > Security and Monitor Devices > Add**.
- Step 2** From the Device Type list, select **Add SW Security apps on a new host**.
You can also select Add SW Security apps on an existing host if you have already defined the host within MARS, perhaps as part of the Management >IP Management settings or if you are running another application on the host, such as Microsoft Internet Information Services.
- Step 3** In the Device Name field, enter the hostname of the server or the remote logging host.
- Step 4** In the Reporting IP field, enter the IP address of the interface in Cisco Secure ACS server or the remote logging host from which the syslog messages will originate.
- Step 5** Under Enter interface information, enter the interface name, IP address, and netmask value of the interface in Cisco Secure ACS server or remote logging host from which the syslog messages will originate.
This address is the same value as the Reporting IP address.
- Step 6** Click **Apply**.
- Step 7** Click **Next** to move the Reporting Applications tab.

↓

General	Reporting Applications
Enter reporting application:	
→ Device Name: Softie II	
→ Select application:	Select one Select one CheckPoint Opsec NG AI CheckPoint Opsec NG FP3 Cisco ACS 3.x Cisco CSA 4.x Cisco ICS 1.x Enterasys Dragon 6.x Entercept Entercept 2.5 Entercept Entercept 4.0 Foundstone FoundScan 3.0 Generic Web Server Generic ISS RealSecure 6.5 ISS RealSecure 7.0 IntruVert IntruShield 1.5 McAfee ePO 3.5 Add
Edit Remove	Device Type

143255

Step 8 In the Select Application box, select **Cisco ACS 3.x**, and then click **Add**.

Result: The Cisco ACS 3.x Windows Requirements page appears, explaining that you must have installed an agent on the server as described in [Install and Configure the PN Log Agent, page 14-7](#).



Note

The Cisco ACS 3.x option supports both Cisco Secure ACS 3.x and Cisco Secure ACS 4.0. No explicit 4.0 option exists for Cisco Secure ACS.

Step 9 Click **Submit** to add this application to the host.

Result: Cisco ACS 3.x appears in the Device Type list.

Step 10 Click the **Vulnerability Assessment Info** link to define the host information that MARS uses to determine false positive attacks against this host. Continue with [Define Vulnerability Assessment Information, page 10-11](#).

Step 11 Click **Done** to save the changes.

Result: The new host appears in the Security and Monitoring Information list.

Step 12 To activate the device, click **Activate**.



Configuring Custom Devices

Adding User Defined Log Parser Templates

MARS allows the user to enter any SYSLOG or SNMP device into the network topology, configure it to report data to the MARS and query the data using free-form query.

User needs to specify the incoming data format so that MARS can parse and retrieve session information from arbitrary logs.

In order to add a user defined log parser template, the following steps have to be taken:

-
Step 1	Add a custom Device or Application type
Step 2	Add a log parser template
Step 3	Add device with the above custom Device or Application type.
-

To add a custom Device/Application type:

-
Step 1	Go to Admin >Custom Setup tab
Step 2	Click the User Defined Log Parser Templates

Figure 15-1 User Defined Log Parser Template

System Setup System Maintenance User Management System Parameters Custom Setup

ADMIN | PN-MARS Standalone: pnmars34 v3.3 Login:

User Defined Log Parser Templates

Device/Application Type: None available **Add** Edit Delete

Log Templates

Log ID	Log Description	Mapped to Event Type

Step 3 On the next screen, click **Add** button which is located next to the Device/Application type list

Figure 15-2 Device Type Definition

System Setup System Maintenance User Management System Parameters Custom Setup

ADMIN | PN-MARS Standalone: pnmars34 v3.3

Device/Application Type Definition

→ *Type: ☒ Appliance ☐ Software

→ *Vendor:

→ *Model:

→ *Version:

Back **Submit**

Step 4 Choose the Type - Appliance or Software.

- Appliance - A hardware device that can send logs to the MARS Appliance
- Software - An application running on a host and the host can be configured to send logs to the MARS Appliance

Step 5 Enter the Vendor, Model and Version for the Device or Application. (For Example, Cisco PIX 7.0)

Step 6 Click **Submit**.

Figure 15-3 User Defined Device/Application Type

User Defined Log Parser Templates

Device/Application Type:

Log Templates for : Vendor1 Model1 1.2

Log ID	Log Description	Mapped to Event Type	Severity
0 to 0 of 0 25 per page			

143244

To add Parser Templates for a Device/Application

- Step 1** Go to the **Admin > Custom Setup** tab.
 - Step 2** Click the **User Defined Log Parser Templates**.
 - Step 3** Select the newly created/existing Device or Application from the dropdown.
 - Step 4** To add a log template, click **Add** which located in the Log Template area.
 A log template ties directly to the particular message that you want to parse. A log template is composed of one or more Event Types that describe the contents of the message. Using the Event Types, MARS parses the message when it is received.
 - Step 5** Enter **Log ID** - a tag that will identify the log message.
 The Log ID field provides an opportunity to map this message number or another moniker used by the device to the custom event type that you are developing. You can use this value to clarify the device messages for which you have developed custom event types.
 - Step 6** Enter **Description** - a description of the log message.
 - Step 7** Map the log to an Event Type.
- 

Note The MARS Appliance comes with a number of predefined Event Types. To display them, select **System** (for example) from the list above the Event Type select window and click **Get**
- Step 8** New **Event Types** can be added by clicking **Add** below the Event Type list.

Figure 15-4 Mapping Log to Event Type

Log Template for : Vendor1 Model1 1.2

Definition Patterns

→ *Log ID: Log1

→ Description: The first example log message to parse

Map to Event Type

→ *Event Type: Teardown Connection

User All Severity Get Search

Teardown Connection

Add Edit Delete

143246

Figure 15-5 Define Event Type

Event Type Definition

→ *Event Type: TEARCONN

→ Description: Teardown Connection

→ Severity: RED

→ CVE Name:

Cancel Submit

143247

- Step 9** Add new Event type and its information and click **Submit (optional)**
- Step 10** Click **Apply** - the **Patterns** link will become enabled.
- Step 11** Click the **Patterns** link.

Figure 15-6 Define Event Patterns

Patterns for Log Template : Log1

Definition	Patterns												
Add Edit Delete Test <table border="1"> <thead> <tr> <th>Position</th> <th>Key Pattern</th> <th>Parsed Field</th> <th>Value Type</th> <th>Value Format</th> <th>Value Pattern</th> </tr> </thead> <tbody> <tr> <td colspan="6">0 to 0 of 0 25 per page</td> </tr> </tbody> </table> Add Edit Delete Test		Position	Key Pattern	Parsed Field	Value Type	Value Format	Value Pattern	0 to 0 of 0 25 per page					
Position	Key Pattern	Parsed Field	Value Type	Value Format	Value Pattern								
0 to 0 of 0 25 per page													

Step 12 Click **Add** to input patterns.

The parsing patterns for the example above are specified to match the following example raw message reported in an event.

```
Teardown TCP connection 1000 faddr 67.126.151.132/80 gaddr 198.133.219.28/43246 laddr
10.1.1.30/890 (sudha) duration 01:00:02 bytes 1000000 (TCP FINs)
```

Step 13 The first step is to identify the values in the log that need to be parsed and stored in MARS events.

Step 14 Currently MARS supports the following parsed value fields in its events:

- Source address
- Destination address
- Source port
- Destination Port
- Protocol
- NAT Source address
- NAT Destination address
- NAT Source port
- NAT Destination Port
- NAT Protocol
- Device Time stamp
- Session Duration
- Received Time stamp
- Exchanged Bytes
- Reported User

Step 15 The parsing format can now be thought of as being made up of several KEY pattern followed by VALUE patterns. Both KEY and VALUE patterns are regular expressions based on the library PCRE which is perl-compatible regular expressions ([Appendix B, “Regular Expression Reference.”](#) for details on syntax). Note that a KEY can be an empty string. A log format consists of several KEY-VALUE sub-pattern pairs.

Figure 15-7 Define Pattern for a Log

Pattern definition for Log ID : Log1

→ Position:

→ Key Pattern:

→ Parsed Field:

→ Value Type:

→ Pattern Name:
Or enter new:

→ Description:

→ Value Pattern:

143249

- Step 16** In the above example, the **Position** refers to the position of each KEY-VALUE sub-pattern pair. These KEY-VALUE sub-pattern pairs are concatenated in the order of their positions and used for matching against the raw message in an event. It does allow arbitrary whitespace between KEY and VALUE patterns, as well as between KEY-VALUE sub-patterns.
- Step 17** In the above example, the **Key-Pattern** is “**Teardown**” is a simple regular expression that does not have any wildcards or repetitions.
- Step 18** The **Parsed Field** is one of fields of a MARS event that has been fully parsed. In the above case, it is the protocol field.
- Step 19** The **Value Type** gives indication to the parser on what kind of value to expect so that suitable parsing action can be applied on the matching sub-pattern string. By “**Choosing Protocol (String)**” as the value type above, we indicate that the protocol field is coming in the form of a string as defined in the file /etc/protocols in a UNIX system. For example, “**TCP**” is the string that will be captured by the value pattern. The **Value Type** will indicate that TCP is to be converted into its protocol number, 6.
- Step 20** **Pattern Name** is a mnemonic given to standard regular expression patterns available for the user who is specifying the log format. There are several common pre defined patterns with appropriate names. In the edit box right below the **Pattern Name** list, a user can add new value names to identify value patterns that may be commonly used in their logs. In the above figure, the value pattern captures all word-character strings that may also include the characters ‘-’, ‘/’ and ‘+’.

Figure 15-8 Sub Pattern Definition

Pattern definition for Log ID : Log1

→ Position:

→ Key Pattern:

→ Parsed Field:

→ Value Type:

→ Pattern Name:
Or enter new:

→ Description:

→ Value Pattern:

- Step 21** The above KEY-VALUE sub-pattern in the 2nd position of the log format. Notice that the key pattern isn't just a simple string, it is a regexp `\d+` which matches against all unsigned decimal numbers. The parsed field where the above value is stored is the Source Address which is specified as a dotted quad. Notice the somewhat complicated regexp that only capture IP addresses in the correct range (0.0.0.0 through 255.255.255.255).

Figure 15-9 Third Position of Pattern Definition

Pattern definition for Log ID : Log1

→ Position:

→ Key Pattern:

→ Parsed Field:

→ Value Type:

→ Pattern Name:
Or enter new:

→ Description:

→ Value Pattern:

- Step 22** The above is for a source port. PORT_NUMBER is the **Pattern Name**, provided for the above **Value Pattern** with the **Description** above.
- Step 23** Repeat for every position of **Pattern definition**.

Figure 15-10 The above example is a 12 KEY-VALUE sub-pattern pieces.

Patterns for Log Template : Log1

Definition		Patterns			
Add Edit Delete Test					
	Position	Key Pattern	Parsed Field	Value Type	Value Pattern
☐	1	Teardown	Protocol	Protocol String	[\w-/+]+
☐	2	connection [+~]? \d+ faddr	Source Address	IPv4 Dotted Quad	((([01]?[d\d]?[2[0-4]\d 25[0-5])\.\.){3}([01]?[d\d]?[2[0-4]\d 25[0-5])
☐	3	/	Source Port	Port Number	((0x[a-fA-F\d]{1,4}) (\d{1,5}))
☐	4	gaddr	Destination Address	IPv4 Dotted Quad	((([01]?[d\d]?[2[0-4]\d 25[0-5])\.\.){3}([01]?[d\d]?[2[0-4]\d 25[0-5])
☐	5	/	Destination Port	Port Number	((0x[a-fA-F\d]{1,4}) (\d{1,5}))
☐	6	laddr	NAT Dest Address	IPv4 Dotted Quad	((([01]?[d\d]?[2[0-4]\d 25[0-5])\.\.){3}([01]?[d\d]?[2[0-4]\d 25[0-5])
☐	7	/	NAT Dest Port	Port Number	((0x[a-fA-F\d]{1,4}) (\d{1,5}))
☐	8	\Q\E	Reported User	String	\w+
☐	9	\Qduration\E	Session Duration	Duration-Hours	(0x)?[a-fA-F\d]+
☐	10	:	Session Duration	Duration-Minutes	(0x)?[a-fA-F\d]+
☐	11	:	Session Duration	Duration-Seconds	(0x)?[a-fA-F\d]+
☐	12	bytes	Transmitted Bytes	Number	(0x)?[a-fA-F\d]+

1 to 12 of 12 25 per page

143234

Figure 15-11 Log template for the device type 'Vendor1 Model1 1.2'

User Defined Log Parser Templates

Device/Application Type: Vendor1 Model1 1.2

Log Templates for : Vendor1 Model1 1.2

Add Edit Delete				
	Log ID	Log Description	Mapped to Event Type	Severity
☐	Log1	The first example log message to parse	Teardown Connection	

1 to 1 of 1 25 per page

143236

Figure 15-12 New software based Apache Webserver application.

Device/Application Type Definition

→ *Type: ☐ Appliance ☒ Software

→ *Vendor:

→ *Model:

→ *Version:

143237

Figure 15-13 Sample Definition for Apache Webserver1.1

Log Template for : Apache Webserver 1.1

↓

Definition	Patterns	
→ *Log ID: HTTP Status OK → Description: HTTP Status 200 (OK) log Map to Event Type → *Event Type: HTTP Status - OK System All Severity Get HTTP Status Search HTTP Status - OK ☐ HTTP Status - Accepted ☐ HTTP Status - Bad Gateway ☐ HTTP Status - Bad Request ☐ HTTP Status - Conflict ☐ HTTP Status - Continue ☐ HTTP Status - Created ☐ HTTP Status - Expectation Failed ☐ HTTP Status - Forbidden ☐ HTTP Status - Found ☐ HTTP Status - Gateway Timeout ☐ HTTP Status - OK Add Edit Delete		143238

Step 24 Define the log template for a HTTP Status OK log message. And associate a system defined event type. In order to find the event type, specify the search string 'HTTP Status' and find it defined as above.

Step 25 The parsing patterns for 'HTTP Status OK' are specified to match the following example raw message reported in an event.

```
155.98.65.40 - - [21/Nov/2004:21:08:47 -0800] "GET /~shash/ HTTP/1.0" 200 1633 "-"
"Lynx/2.8.2rel.1 libwww-FM/2.14"
```

Figure 15-14 Key Pattern for HTTP Status OK

Pattern definition for Log ID : HTTP Status OK

→ Position:

→ Key Pattern:

→ Parsed Field:

→ Value Type:

→ Pattern Name:
Or enter new:

→ Description:

→ Value Pattern:

143239

Step 26 The key pattern is empty above since the log message starts with a value pattern.

Figure 15-15 Position 2 Key Pattern for HTTP Status OK

Pattern definition for Log ID : HTTP Status OK

→ Position:

→ Key Pattern:

→ Parsed Field:

→ Value Type:

→ Pattern Name:
Or enter new:

→ Description:

→ Value Format:

→ Value Pattern:

143240

Step 27 The **Parsed Field** above is a Date/Time field. In addition to **Value Pattern**, a value format is required since a date/time can be specified in arbitrarily different ways. Details on how to specify the value format are given in Appendix F. Several pattern names with a few of the commonly used date/time formats have been predefined.

Pattern definition for Log ID : HTTP Status OK

143241

Patterns for Log Template : HTTP Status OK

143242

Figure 15-18 User Defined Log Parser for Apache Webserver1.1

User Defined Log Parser Templates

Device/Application Type: Apache Webserver 1.1

Log Templates for : Apache Webserver 1.1

Log ID	Log Description	Mapped to Event Type	Severity
HTTP Status OK	HTTP Status 200 (OK) log	HTTP Status - OK	

1 to 1 of 1 25 per page

143243

Adding the new custom device or application is similar as adding predefined device or application. Below is the example of adding the newly user created Apache Webserver1.1:

- Step 1** Go to the **Admin** tab.
- Step 2** Click the **Security and Monitor Devices**.
- Step 3** Click **Add** to add a new device.
- Step 4** From the **Device Type** list, select **Add SW security apps on a new host**.
- Step 5** Fill in name and other host details and click **Apply**.
- Step 6** Click on **Reporting Applications**.
- Step 7** Select **Application** (e.g., Apache Webserver.1.1) from the list and click **Add**.

Figure 15-19 Adding the Customer Application to MARS

Device Type: Edit host with security applications

↓

General	Reporting Applications	Vulnerability Assessment Info
Enter reporting application:		
→ Device Name: Host1 → Select application: Apache Webserver 1.1 Add Edit Remove Device Type ☐ Apache Webserver 1.1		

143245

- Step 8** Select Reporting Method (SNMP TRAP or SYSLOG) on the resulting window and click **Submit**.
- Step 9** Click **Done**.
-



Policy Table Lookup on Cisco Security Manager

MARS and Cisco Security Manager (Security Manager) can be configured to provide round-trip policy audit features and improve traffic flow analysis and debugging. Using this feature, you can identify the ACL on a router or firewall that generated a syslog message received by MARS. It is important to understand that the integration between MARS and Security Manager is unique; MARS can provide users of Security Manager with better analytical tools.

When using MARS as your STM solution, you must understand that MARS suggests and makes changes to devices without notifying Security Manager of the suggested changes. Specifically, you can use the “Big Red” button to shutdown a port for support L2 devices. For a layer 3 device, MARS suggest ACL changes to block the traffic. In such cases, you can use Security Manager to manually mitigate using the ACL recommendations provided by MARS, thereby, ensuring that the configuration management solution stays abreast of the mitigation responses. Security Manager can also publish the same change to all like devices that it manages, providing a more robust containment.


For example, consider the following case where a user cannot connect to *destination X* from *source Y*. To troubleshoot this issue, an administrator can do the following:

1. Log into the MARS web interface, and using an on-demand query, determine whether an event has been received that shows that traffic from *source Y* to *destination X* has been blocked.
2. If such events are found, the administrator can continue by determining which ACL is actually blocking the traffic. To do so, the administrator would click the policy query icon in the row of one of the selected events. MARS then queries Security Manager to retrieve the list of ACLs that match that traffic flow, and assuming Security Manager was used to configure the routers and firewalls between *source Y* and *destination X*, then a list of matching ACLs are returned.
3. Next, the administrator can log into the Security Manager user interface and modify the identified policy, or ACL, to allow traffic between *source Y* and *destination X*.

This chapter describes how to configure Security Manager and MARS to ensure optimal functionality and seamless integration.

Overview of Cisco Security Manager Policy Table Lookup

When MARS receives a syslog from a Cisco PIX firewall, Cisco Adaptive Security Appliance (Cisco ASA), Cisco Firewall Services Module (Cisco FWSM), or Cisco IOS, and can derive the five tuple information required to establish an event (source IP, destination IP, source port, destination port, and protocol) the Security Manager Policy Table Lookup icon  appears in the Reporting Device column of the MARS session display. Clicking the icon invokes a query to the Security Manager, the result of which is to identify the access rule in the policy table of the device which created the traffic incident or event. [Figure 16-1](#) depicts the policy query process between MARS and Security Manager.

Figure 16-1 Cisco Security MARS Policy Table Query Process

The syslog that generated the MARS incident or event may not have sufficient information for Security Manager to uniquely identify the device or the access rule. In these ambiguous cases, Security Manager returns a list of all possible devices to MARS in a pop-up window. When the MARS user manually selects a reporting device, the policy table is then displayed for that device. Access rules that match the query criteria are highlighted.

**Note**

The policy table displayed by MARS is the Security Manager Committed View, not the Deployed View, meaning the displayed Security Manager policies are saved in the Security Manager database but not yet deployed on the device. If the deployed and committed views are not identical, the access rule generating the MARS event may not be visible in the policy table displayed by MARS. For further information on Cisco Security Manager operation, please access the documentation at the following URL:
http://www.cisco.com/en/US/products/ps6498/tsd_products_support_series_home.html

More About Cisco Security Manager Device Lookup

MARS requests the Policy Table of a Security Manager device by supplying the following criteria to Security Manager:

- Device Name—Derived from MARS Device Name
- IP Address—Derived from MARS Reporting IP
- Domain Name—If available, derived from the device name in MARS (for example, c3550-225-125.clab.cisco.com)
- Device Type—If available from MARS

The Device Lookup query includes the following actions between MARS and Security Manager:

1. Security Manager matches the MARS Device Name to the Security Manager host names. If only one matching host name is discovered, the process for Policy Table Lookup is invoked.
2. If the Security Manager host name is undefined, then Security Manager matches the MARS Device Name with the Security Manager Display Name (display names are unique, but the host name may be a substring of many display names.)
3. If there are multiple matches on host name and no unique display name matches, the domain name (if available) is used to narrow the choices.
4. If the domain name is not available, MARS Reporting IP is used to narrow the choices.
5. If a unique device cannot be identified, MARS displays a list of possible devices in a pop-up window that shows the IP address, host name, display name, and domain name for all possible device matches. The user manually selects the device and the process for the policy table lookup is invoked.

More About Cisco Security Manager Policy Table Lookup

The device lookup information is combined with the event information to perform the Security Manager policy table lookup.

The following MARS event information derived from the reporting device raw message is passed to Security Manager:

- **Event Five Tuple**—Source IP Address, Destination IP address, Source Port, Destination Port, and Protocol defining session. The event five tuple must match the five tuples of the target access rule. For ICMP logs, ICMP type and code, when available, are passed instead of the source and destination ports.
- **Action**—If available, permit or deny. If not available, access rules with both permit and deny are highlighted.
- **ACL name**—If available, the name of the ACL or Access Group that triggers the syslog. With the ACL name, Security Manager can reduce the number of matching access rules.
- **Interface**—If available, the interface names are parsed from the event's raw message.
- **Direction**—If available, keyword such as “inbound” and “outbound” identify the direction.

The device, five tuple, action, ACL name, interface, and direction information comprise the policy query criteria submitted to the Security Manager. MARS displays the policy table in a pop-up window. The matching access rule is displayed in highlight. If MARS was unable to provide the interface, direction, and action information, multiple matched access rules may be highlighted.

Sample Cisco PIX Firewall Syslog Messages with Direction and Protocol Information

```
10.33.10.2 <142>%PIX-6-302013: Built outbound TCP connection 2021 for
inside:10.1.1.10/4000 (10.1.1.10/4000) to dmz:192.168.1.10/80 (192.168.1.10/80)
```

```
10.33.10.2 <142>%PIX-6-302013: Built inbound TCP connection 2000 for
outside:1.234.58.149/12000 (1.234.58.149/12000) to inside:192.168.1.10/25 (100.1.4.10/25)
```

Sample Cisco PIX Firewall Syslog Messages with Access Group Name Information

```
10.33.10.2 <142>%PIX-4-106023: Deny tcp src inside:10.1.5.234/3010 dst outside:5.6.7.8/21
by access-group "Cisco Security Manager-acl-inside"
```

Sample Cisco IOS 12.2 Syslog Messages with ACL Name Information

```
100.1.20.2 Mon Jun 9 14:46:31 2003 <46>485232: Jun 9 14:46:29 PDT: %SEC-6-IPACCESSLOGP:
list Cisco Security Manager-acl-FastEthernet0/0 permitted tcp 1.234.51.255(12000) ->
100.1.4.10(25), 1540 packet
```

```
10.34.1.1 <46>146570: Dec 19 21:01:57 PST: %SEC-6-IPACCESSLOGP: list Cisco Security
Manager-acl-FastEthernet1/0 denied tcp 10.10.1.20(59399) -> 10.1.5.11(23), 1 packet
```

Prerequisites for Policy Table Lookup

- MARS Local Controller running software version 4.2.1 or more recent version.
- Cisco Security Manager version 3.0.1 or more recent
- MARS configured for operation with Cisco Security Manager as explained in the section, [Checklist for Security Manager-to-MARS Integration, page 16-6](#)

Restrictions for Policy Table Lookup

- A Local Controller can be configured to retrieve the policy tables from only one Cisco Security Manager server at a time.
- The Policy Table Lookup icon in MARS is displayed only for traffic logs which are reported by the following MARS device types:
 - Cisco Switch-IOS
 - Cisco PIX firewalls
 - Cisco Adaptive Security Appliance (Cisco ASA)
 - Cisco Firewall Services Module (Cisco FWSM)
- MARS displays the Cisco Security Manager security policy committed views, not the deployed views. The access rule causing the MARS event may not be visible in the policy table. To examine the deployed policies view of a device, you must login to the device or Cisco Security Manager directly.
- MARS examines only Layer 3 ACLs for traffic events on the supported reporting devices. The policy table lookup cannot directly indicate the cause of a traffic event caused by a deny not related to the displayed access rules, though the policy table can be displayed for the event (For instance, a no route deny, or a Network Address Translation [NAT] deny due to a NAT misconfiguration).
- The Security Manager Policy Table Lookup icon displays only for those events with 5-tuple event data (source and destination address, protocol, source and destination port). In the MARS web interface, the all matching events query displays the text “session five tuple” for events with no 5-tuple event data. These events will not have a policy query icon.
- The Security Manager Policy Table Lookup icon displays for NetFlow events even though they are not triggered by an ACL. This extra event data allows you to determine whether there is a policy permitting that traffic, which ensures you are able to tune accordingly.



Note

Because this is NetFlow data, it may not match the exact ACL or match multiple ACLs.

- The same events received by MARS can display the Security Manager Policy Table Lookup icon inconsistently between the low-latency, real-time event query and standard queries, such as sessions ranked by time. Specifically, the icon will not appear in the low-latency, real-time query, but it may appear in queries against sessionized events.

This behavior is expected. When MARS receives events, they are parsed, sessionized, written to an event shared buffer, and then written to the database. Because sessionization takes time, sometimes keeping an event in cache for 2 minutes, the low-latency event query displays events right after parsing, but before sessionization. Displaying the event at this point allows the low-latency query to achieve a close to real-time effect. For some events, parsing cannot determine some part of the 5-tuple data, such as a destination address. Later, sessionization later fills in such missing data using configuration data. As a result, the 5-tuple data displayed by the low-latency event query can be different from values stored in the database, which are used to populate the standard queries.

- An error can occur with the policy query if a device configuration is discovered using Security Manager but it is not submitted in Security Manager. The following error message is an example of this issue:

```
<190>2312080: *May 9 23:50:02.199: %SEC-6-IPACCESSLOGDP: list permit-all permitted
icmp 10.2.3.8 -> 10.4.21.2 (0/0), 1 packet
```

```
An error occurred while querying policies from Cisco Security Manager. Reason:  
Failed to retrieve policy information from CSM. Reason: Cisco Security Manager  
Internal error: Failed to get interfaces in the device!  
The device LC2DTM was discovered by CSM without any errors.
```

Before you perform policy queries, verify that all discovered devices have been submitted in Security Manager.

Checklist for Security Manager-to-MARS Integration

Security Manager-to-MARS integration deals with identifying the required and optional points of integration, configuring the applications and devices, and ensuring proper authorization among the two management platforms. This checklist assumes a greenfield install of both Security Manager and MARS.

The following checklist describes the tasks required to understand the decision-making process and the basic flow required to integrate MARS with a Security Manager server and the reporting and mitigation devices managed by that Security Manager server. Each step might contain several substeps; the steps and substeps should be performed in order. The checklist contains references to the specific procedures used to perform each task.

**Note**

For a comprehensive checklist on configuring MARS to operate most effectively as an STM, see [STM Task Flow Overview, page 1-1](#).

✓	Task
☐	1. Inventory overlapping reporting devices and mitigation devices. MARS supports round-trip policy audit analysis for reporting and mitigation devices that are both managed by Security Manager and monitored by a MARS Appliance. In other words, MARS can query the policy rules that generated an audit event log only when the policies are defined using Security Manager. As such, the first step in integrating MARS and Security Manager involves identifying those devices for which Security Manager is used to define policy rules. You must also ensure that devices are running a software versions supported by both MARS and Security Manager. This list focuses on those devices that Security Manager manages and should include all of the following devices: • Cisco ASA appliances, PIX appliances, and FWSM modules • Cisco Catalyst 6500 Series Switches • Cisco Routers running supported versions of Cisco IOS software Note MARS supports PIX 7.0 and ASA 7.0.1 releases; however, it does not support FWSM 3.1. FWSM support is restricted to FWSM 1.1, 2.2, and 2.3. For current device support information, see Supported Devices and Software Versions for Cisco Security MARS. Note FWSM support is supported only in Cisco Security Manager Enterprise Edition (Professional-50) and higher. The Professional version includes support for the management of Cisco Catalyst® 6500 Series switches and associated services modules; the Standard versions do not include this support. <i>Result:</i> The list of devices for which Security Manager manages the security and audit log policies is defined. The details of each device include device name, reporting IP address, management IP address, management protocol, administrative account information, and the logging features, levels, and protocols to enable. For more information, see: • Supported Devices and Software Versions for Cisco Security Manager 3.0 • Supported Devices and Software Versions for Cisco Security MARS • Selecting the Devices to Monitor, page 2-2 • Levels of Operation, page 2-1 • Deployment Planning Guidelines, page 2-1 in <i>Install and Setup Guide for Cisco Security Monitoring, Analysis, and Response System</i> • Device Inventory Worksheet, page 1-18

✓	Task
□	2. Identify and enable all required traffic flows. After you identify the devices managed by the Security Manager server, you must verify that the network services they use for management, reporting, and notification are permitted along the required traffic flows. Using the detailed Device Inventory Worksheet identified in Step 1., ensure that the management, logging, and notification traffic between the MARS Appliance and each supporting device, reporting device, and mitigation device is allowed by intermediate gateways. In addition, network services of supporting devices, such as DNS, e-mail, AAA, and NTP servers, must also be permitted to flow among the MARS Appliance, the supporting devices, and the reporting devices and mitigation devices on your network. Tip It is a recommended security practice to have all devices, including MARS Appliances, synchronized to the same time. <i>Result:</i> You have verified that all intermediate gateways permit the log, management, and notification traffic between the devices and the MARS Appliance. For more information, see: • Deployment Planning Guidelines, page 2-1, in <i>Install and Setup Guide for Cisco Security Monitoring, Analysis, and Response System</i> • Supporting Devices, page 2-1, in <i>Install and Setup Guide for Cisco Security Monitoring, Analysis, and Response System</i> • Required Traffic Flows, page 2-2, in <i>Install and Setup Guide for Cisco Security Monitoring, Analysis, and Response System</i> • Specify the Time Settings, page 5-10, in <i>Install and Setup Guide for Cisco Security Monitoring, Analysis, and Response System</i> • Event Timestamps and Processing in <i>Top Issues for the Cisco Security Monitoring, Analysis, and Response System</i> • Device Inventory Worksheet, page 1-18

✓	Task
☐	3. Bootstrap the reporting devices and mitigation devices managed by Security Manager. For each device identified in Task 1., you must prepare, or bootstrap, that device to ensure that the desired communications with MARS occur. Bootstrapping a device involves configuring the settings for that device, based on its role in the STM system. Perform the following subtasks as applicable to a device type and its role: • Enable management of the device by the MARS Appliance for mitigation and access. • Turn on the correct logging level and logging services. • Direct the logs to the MARS Appliance. • Enable discovery of the device settings. Note While many Cisco devices support the EMBLEM syslog format, this format is <i>not compatible</i> with MARS. As part of this task, you must verify that the devices are not reporting to the MARS Appliance using the EMBLEM format. You must configure the router and switch settings using the CLI, as Security Manager does not support those features. However, for ASA, FWSM, and PIX, you can use the Security Manager user interface to configure the management and log settings. Tip Any events published by a device to MARS prior to adding and activating the device in the web interface can be queried using the reporting IP address of the device as a match criterion. This technique can be useful for verifying that the device is properly bootstrapped. You may also need to enable alternate settings on the to provide richer data. For more information on these possible settings, see Task 5 in the Checklist for Provisioning Phase, page 1-2 found in the STM Task Flow Overview chapter. <i>Result:</i> The correct logging levels are enabled on the reporting devices and mitigation devices. The MARS Appliance can receive or pull any necessary logs from those devices, and it can retrieve configuration settings and push ACLS to the supported mitigation devices. While the MARS Appliance picks up and stores the events it receives, it does not inspect them until the reporting devices and mitigation devices are defined and activated in web interface. For more information, see: • Device Inventory Worksheet, page 1-18 • Bootstrap Summary Table, page 2-12 • Cisco Router Devices, page 3-1 • Cisco Switch Devices, page 3-9 User Guide for Cisco Security Manager 3.0 • Understanding Device Credentials See SNMP credentials. • Managing Firewall Devices (ASA, PIX, and FWSM) See device access, SNMP settings, logging policies, and static routes as needed. Note When defining SNMP settings for the FWSM and ASA, you should define these setting for the admin context. • Field definitions for the Logging Policies (ASA, PIX, and FWSM) • Managing Routers (Cisco IOS Routers) See device access, SNMP, 802.1x, NAC, and static routes as needed. • Using the Catalyst 6500/7600 Device Manager (Cisco Switches) See Spanning Tree Settings (STP).

✓	Task
□	4. Define the devices in MARS. After you identify and bootstrap the reporting devices and mitigation devices and enable the required traffic flows, you must represent those devices in MARS, which uses this information to communicate with the devices. You can do this by adding individual devices in the web interface or by importing a comma separated vector (CSV) file, which can define the required settings for basic device types and give you a headstart on defining the more complicated devices. After you add the devices, you must activate them by clicking Activate on any page in the web interface. To display all devices that are either added incorrectly or not activated in MARS, you can define one of two queries: • Select “Unknown Reporting Device” in the Devices field. This query returns the events only for those devices that are reporting events that do not matching the one of the reporting IPs defined in MARS. When MARS receives events, it first determines if the IP from which the events are received matches one of reporting IPs identified in the Reporting and Monitor Devices page. Only if MARS finds a match does it attempt to parse the events. Therefore, if the Reporting IP is defined incorrectly for a reporting device, the events from that device are not parsed. This query essentially identifies events that are not parsed. • Select the “Unknown Device Event Type” in the Events field. This query returns events from known devices that for some reason the event is not parsed by MARS (for example, if the MARS signature list is not current with the device event lists), and it returns events reported by unknown devices. For both queries, if you are looking for a specific reporting IP address, enter that address in the Keyword field to filter the results down to those that include that IP address. <i>Result:</i> All reporting devices and mitigation devices are defined and activated in MARS. When the devices are bootstrapped and defined in MARS, MARS begins to inspect the logs received from the devices. Until the devices are added in MARS, MARS picks up and stores the events it receives without inspecting them. For more information, see: • Device Inventory Worksheet, page 1-18 • Selecting the Access Type, page 2-10 • Add Reporting and Mitigation Devices Individually, page 2-17 • Add Multiple Reporting and Mitigation Devices Using a Seed File, page 2-20 • Adding Reporting and Mitigation Devices Using Automatic Topology Discovery, page 2-25 • Cisco Firewall Devices (PIX, ASA, and FWSM), page 4-1 • Cisco Router Devices, page 3-1 • Cisco Switch Devices, page 3-9 • Supported Reporting and Mitigation Devices, page 2 (CSV Keyword column) • Verify Connectivity with the Reporting and Mitigation Devices, page 2-26 • Activate the Reporting and Mitigation Devices, page 2-27

✓	Task
	5. Bootstrap each Security Manager server and add it to the correct MARS Local Controller. The required setup of the Security Manager server is quite simple: • Verify that HTTPS is enabled on the Security Manager server. • Define an account with requisite privileges for use by MARS when performing policy queries. • Ensure that the MARS Appliance, as a host, has permission to access the Security Manager server. • Ensure the HTTPS traffic flows are permitted between the two hosts. You can verify connectivity by clicking Test Connectivity when defining a Security Manager server. It does not perform a lookup, but it does authenticate to the Security Manager server. Note Each Local Controller can only manage one Security Manager server. All policy lookups for syslog messages received on a Local Controller are directed to the Security Manager server configured for that Local Controller. In addition, you should ensure that the user account used by MARS is exclusive on the Security Manager server. An exclusive account allows you to more easily detect any fraudulent use of the account. MARS does not retrieve any audit log data from the Security Manager server. It merely accesses the set of policy rules that are defined on the server. Once you prepare the Security Manager server, you must return to the MARS web interface and add the Security Manager server. <i>Result:</i> The correct settings are enabled on each Security Manager server. The MARS Appliance can request and receive queries from no more than one Security Manager server. After adding the Security Manager server to the MARS web interface, you can test the connectivity by performing a policy lookup query on any of the events that have been received from one of the managed reporting devices and mitigation devices. For more information, see: • Bootstrapping Cisco Security Manager Server to Communicate with MARS, page 16-12 • Add a Cisco Security Manager Server to MARS, page 16-13 • Supported Reporting and Mitigation Devices, page 2 • Procedure for Invoking Cisco Security Manager Policy Table Lookup from Cisco Security MARS, page 16-14
	6. Perform policy lookups as required. Once an event generated by one of the Security Manager-managed devices is received by MARS, you can perform a policy lookup operation. This operation allows you to, from a given incident or an event query, retrieve a list of the possible policies that could have affected the generation of that event from the Security Manager server. For more information, see: • More About Cisco Security Manager Device Lookup • More About Cisco Security Manager Policy Table Lookup • Prerequisites for Policy Table Lookup • Restrictions for Policy Table Lookup • Procedure for Invoking Cisco Security Manager Policy Table Lookup from Cisco Security MARS

✓	Task
	7. Using Security Manager for mitigation response. While MARS suggests ACL changes to mitigate attacks, and in the case of Layer 2 devices such as Cisco switches, it can push changes to layer 2 device via the “Big Red” button (which shuts down a port on a switch), you must ensure accuracy between the policy defined in Security Manager and the configuration running on the managed devices. This synchronization ensures an accurate understanding of your network configuration and improves your ability to troubleshoot issues using the policy analysis tools provided in Security Manager. Therefore, we recommend that you perform the device mitigation by applying the rules recommended by MARS with Security Manager. This approach also prevents you from having to manually synchronize your policy between Security Manager and the mitigation devices. As an added benefit, you can enable and remove containment rules on multiple devices via global rules, thereby further restricting the spread of possibly undetected infections. Using comments in the rules, you can document the attack responses, allowing for future analysis when considering global network stances and when developing attack response strategies.

Bootstrapping Cisco Security Manager Server to Communicate with MARS

To prepare the Security Manager server to be queried by MARS, you must configure the following settings:

- Define an admin account in Security Manager that MARS can use to perform queries. A separate account is recommended to provide a cleaner audit trail on the Security Manager server. The following security levels defined in Common Services 3.0 server satisfy the authorization requirements of MARS-to-Security Manager policy query:
 - Help Desk
 - Network Operator
 - Network Administrator
 - System Administrator



Note

Cisco does not recommend using System Administrator for this account. Instead, we recommend least privilege settings (only enabling those privileges required to perform the job). As such, we recommend defining an admin account with the Help Desk security level.

For more information on defining admin accounts on the Common Services 3.0 server, see:

http://www.cisco.com/en/US/products/sw/cscowork/ps3996/products_user_guide_chapter09186a008022f958.html#wp372210

- Enable HTTPS access to the Common Services 3.0 server by the MARS Appliance. If you are using AAA authentication, such as Cisco Secure ACS, on the Common Services 3.0 server, you must update the administrative access settings to ensure that the MARS Appliance has the necessary access to the Security Manager server.
- Before MARS can query the policies defined on the Security Manager server, you must enable HTTPS on the Security Manager server. For more information on enabling HTTPS, see:

http://www.cisco.com/en/US/products/sw/cscowork/ps3996/products_user_guide_chapter09186a08022f958.html#wp339451

Add a Cisco Security Manager Server to MARS

The Security Manager server is represented in MARS by defining a host with a software application residing on that host. Once you have identified the reporting devices to a Local Controller, you can add the Security Manager server that manages the policies for those reporting devices.

Each Local Controller can query one Security Manager server only; you cannot define more than one Security Manager server per Local Controller. You can define the same Security Manager server on multiple Local Controllers. When planning the zones for Global Controller/multi-Local Controller deployments, ensure that each Local Controller maps to the Security Manager server that manages the reporting devices monitored by that Local Controller.

To identify a Security Manager server to use for policy lookups from within the web interface of MARS, follow these steps;

-
- Step 1** Select **Admin > System Setup > Security and Monitor Devices > Add**.
- Step 2** Do one of the following:
- Select **Add SW Security apps on a new host** from the Device Type list, and continue with [Step 3](#)
 - Select **Add SW security apps on existing host** from the Device Type list. Select the device to which you want to add the software application and click Add. Continue with [Step 6](#).
- Step 3** Specify values for the following fields:
- **Device Name** — Enter the name of the device. This name must exactly match the hostname shown in the Cisco Security Manager user interface. MARS maps this name to the reporting IP address. This name is used in topology maps, queries, and as the primary management station in the Security and Monitoring Device list.
 - **Access IP** — This address is used to pull query data from a Security Manager server using HTTPS, enabling MARS to discover settings and perform policy queries from this device. This address represents the physical IP address of the Security Manager server. To learn more about the access IP address, its role, and dependencies, see [Understanding Access IP, Reporting IP, and Interface Settings, page 2-8](#).
 - **Reporting IP** — (Optional) If the Security Manager server is host to a reporting device other than Cisco Security Manager, enter the IP address of the interface in the Security Manager server from which MARS. This address represents the physical IP address of the Security Manager server. To learn more about the reporting IP address, its role, and dependencies, see [Understanding Access IP, Reporting IP, and Interface Settings, page 2-8](#).
 - **Operating System** — (Optional) If the Security Manager server is host to a reporting device other than Cisco Security Manager, you may need to specify the operating system type.
- Step 4** Under Enter interface information, enter the interface name, IP address, and netmask value of each interface in the Security Manager server from which configuration information will be queried.
- This address represents the physical IP address of the Security Manager server. This information is used to ensure that the topology generated by MARS represents all network connections for the Security Manager server. It is also used to calculate possible attack paths that might include the Security Manager server.
- Step 5** Click **Apply** to save these settings.

- Step 6** Click **Next** to access the Reporting Applications tab.
- Step 7** Select the **Cisco Security Manager ANY** from the Select Application list, and click **Add**.
- Step 8** If you entered an address in the Access IP field on the host that represents this Security Manager server, specify values for the following fields:
- **User Name**— Identifies the Cisco Security Manager administrative account to be used to discover configuration settings.
 - **Password** — Identifies the password associated with the User Name account.
 - **Access Type** — This value identifies the protocol used to discover configuration information. Select **HTTPS**.
 - **Access Port** — The default access port for HTTPS is port 443.
- Step 9** (Optional) To verify that the settings are correct and that the MARS Appliance can communicate with this Security Manager server, click **Test Connectivity**.
- Result:* If the username and password are correct and the MARS Appliance is configured as an administrative host for the device, the “Connectivity successful.” dialog box appears when the discovery operation completes. Otherwise, an error message appears, which you can click on the View Error link for more information.
- Step 10** To add this device to the MARS database, click **Submit**.
- Result:* The submit operation records the changes in the database tables. However, it does not load the changes into working memory of the MARS Appliance. The activate operation loads submitted changes into working memory.
- Step 11** Click **Activate**.
- Result:* Once the MARS Appliance is activated, it can query the Security Manager server to perform policy lookups.
-

Procedure for Invoking Cisco Security Manager Policy Table Lookup from Cisco Security MARS

Do the following steps to view a Cisco Security Manager policy table from the Cisco Security MARS:

-
- Step 1** Log on to MARS as an Administrator or Security Analyst.
- Step 2** Identify the incident or event to investigate.
- In this procedure, an incident to investigate appears on the **Recent Incidents** section of the Dashboard, as shown in [Figure 16-2](#).

Figure 16-2 Recent Incidents on MARS Summary Page

The screenshot shows the Cisco Security MARS Summary Page. The top navigation bar includes links for SUMMARY, INCIDENTS, QUERY / REPORTS, RULES, MANAGEMENT, ADMIN, and HELP. The main header displays 'CS-MARS Local Controller: LC44/LC44_zone v4.2' and the login status 'Login: Local: Administrator (pnadmin)'. Below the header, there's a 'Select Case' dropdown set to 'No Case Selected...'. The 'Page Refresh Rate' is set to '15 minutes'. On the left, a '24 Hour Events' summary shows: Netflow (0), Events (1,955,039), Sessions (612,812), and Data Reduction (68%). Below this, a '24 Hour Incidents' summary shows: High (248, 40%), Medium (145, 23%), Low (223, 36%), and Total (616, 100%). The 'All False Positives' section shows 0 confirmed. The main table, titled 'Recent Incidents', has columns: Incident ID, Event Type, Matched Rule, Action Time, Path, and Cases. It lists four incidents, including 'Inactive CS-MARS reporting device', 'WWW .bat File Request', 'Deny packet due to security policy', and 'Half-open SYN Attack'.

Step 3 Click Incident ID of the incident to examine. The Incident Page appears as shown in Figure 16-3.


Figure 16-3 MARS Incident Page Displaying the Cisco Security Manager Icon in the Reporting Device Field

The screenshot shows the Cisco Security MARS Incident Page for Incident ID 2289917483. The top navigation bar is the same as in Figure 16-2. The main header displays 'CS-MARS Local Controller: LC44/LC44_zone v4.2'. Below the header, there's a 'Select Case' dropdown set to 'No Case Selected...'. The 'Incident ID' is 2289917483 and the 'Session ID' is empty. The 'Rule Name' is '147.rule' and the 'Status' is 'Active'. The 'Action' is 'None' and the 'Description' is 'for showing multiple rule match'. Below this, a table shows the incident details with columns: Offset, Open, Source IP, Destination IP, Service Name, Event, Device, Reported User, Keyword, Severity, Count, Close, and Operation. The table has one row with the following data: 1, ANY, 10.1.1.147, ANY, ANY, ANY, ANY, ANY, ANY, ANY, 1, Close, Operation. Below the table, there's a section for 'Incident ID: 2289917483' with 'Expand All' and 'Collapse All' buttons. The main table, titled 'Incident Details', has columns: Offset, Session / Incident ID, Event Type, Source IP/Port, Destination IP/Port, Protocol, Time, Reporting Device, Reported User, Path / Mitigate, and False Positive. It lists four incidents, including 'Built/teardown/permitted IP connection', 'Built/teardown/permitted IP connection', 'Built/teardown/permitted IP connection', and 'Built/teardown/permitted IP connection'. The 'Reporting Device' field for the first three incidents shows a Cisco Security Manager icon.

Step 4 Click the Security Manager Policy Query icon in the Reporting Device field to invoke the Security Manager policy table lookup. One of the following three pop-up windows may appear:

- Multiple Events window—Lists all Security Manager device events in the session, this window appears in this step when there are two or more events in the session.
- Multiple Devices window—Lists all matching Security Manager devices that meet criteria available to MARS, this window appears in this step when there are two or more matching devices.
- Policy Table window—Lists the policy table of the reporting device. Access rules that match the MARS criteria are highlighted, this window appears in this step when there is one event and a unique Security Manager device identified.


In this procedure, the Multiple Events window appears, as shown in Figure 16-4.

Figure 16-4 MARS Multiple Events Pop-up Window

Step 5 Click the Security Manager icon in the **Policy** field of the appropriate event. One of the following two pop-up windows may appear:

- Multiple Devices window
- Policy Table window

In this procedure the Multiple Device pop-up window is displayed, as shown in [Figure 16-5](#).

Figure 16-5 MARS Multiple Devices Pop-up Window

Step 6 Click the radio button of the appropriate Security Manager Device. Click **Submit**. The Policy Table pop-up window appears for the selected device, as shown in [Figure 16-6](#).

Figure 16-6 Policy Table Pop-up Window

Cisco Security Manager



Feb 22, 2006 12:17:10 PM PST

Local Controller: LC44/LC44_zone v4.2 Login: Local: Administrator (pnadmin) :: [Close](#)

Event / Session / Incident ID	Event Type	Source IP/Port	Destination IP/Port	Protocol	Time	Reporting Device	Path / Mitigation	Tune
E:2305620876, S:2305620876, I:2289917478	Built/teardown/permitted IP connection	10.2.3.58 58062	10.1.1.147 443	TCP	Feb 21, 2006 5:30:33 PM PST	cherryWall		False Positive

Total policies returned: 14, Number of matched policies: 2, Jump to matched policy: [First](#) [Last](#)

No.	Permit	Source	Destination	Service	Interface	Dir.	Options	Category	Description	Prev/Next
Local (14 Rules)										
1		any	any	IP	inside	in	LOG	None		« »
2		10.10.0.0/16	any	Syslog	outside	in	LOG	None		
3		10.10.0.0/16	any	SNMP-Trap	outside	in	LOG	None		
4		10.4.0.0/16	any	Syslog	outside	in	LOG	None		
5		10.4.0.0/16	any	SNMP-Trap	outside	in	LOG	None		
6		10.10.0.0/16	any	ICMP-Echo-Reply	outside	in	LOG	None		
7		10.4.0.0/16	any	ICMP-Echo-Reply	outside	in	LOG	None		
8		10.2.0.0/16	any	ICMP-Echo-Reply	outside	in	LOG	None		
9		10.2.3.46	any	HTTPS	outside	in	LOG	None		
10		10.2.3.0/24	10.1.1.147	HTTPS	outside	in	LOG	None		« »
11		10.2.3.0/24	any	HTTPS	outside	in	LOG	None		
12		10.2.3.8	128.107.234.204	SMTP	outside	in	LOG	None		
13		10.2.3.8	171.70.168.183	DNS-UDP	outside	in	LOG	None		
14		10.2.3.0/24	any	SMTP	outside	in	LOG	None		

1 to 14 of 14 25 per page [v](#)

Copyright © 2003, 2005 Cisco Systems, Inc. All rights reserved. [Feedback](#)

The access rules that match the MARS query criteria appear in highlight. The access rules displayed in the Policy Table window are derived from the Security Manager committed view, not the deployed view. You must login to Security Manager or the specific device to examine or alter the access rule generating the MARS event or incident. If the committed and deployed views are identical, locating the policy is simplified. A MARS event can be generated from a deployed access rule not visible in the committed view.

Step 7 Login to Cisco Security Manager or the specific device to alter the security rule creating the MARS event.

This ends the Procedure for Accessing Cisco Security Manager Logs from Cisco Security MARS.

**Note**

Incidents, events, and sessions related to devices managed by Cisco Security Manager can also be viewed with an inline Query or a report.



Network Summary

This chapter describes the web interface and the components of the Summary tab of the web interface.


Navigation within the MARS Appliance

The MARS web interface runs within a single browser window. The MARS product functions are categorized with labeled tabs, each tab subdivided with subtabs.

Logging In

- Step 1** To login to the Local Controller, enter its IP or DNS address into the browser address field. The login box appears.

Figure 17-1 Local Controller Login Box



- Step 2** Enter your login name and password. If you do not have a login name, contact your network administrator.
- Step 3** From the **Type** drop-down list, select **Local** if you are logging in to a user account created on this MARS, or select **Global** if you are logging in to a user account created on the Global Controller to which this Local Controller reports.

Step 4 Click Login.

The first page to appear after a login is the Summary tab Dashboard page. The duration of the delay in displaying information results from a combination of the following causes:


- How long the Local Controller has been powered up and connected to the network.
- Amount of traffic on your networks
- Reporting syslog levels of the reporting devices
- Size of the network
- The number and type of reporting devices

For most networks, the Summary page populates shortly after configuration. Some values are only relevant after an interval of time. For example, the values in the **24 Hour Events** and **24 Hour Incidents** tables.

Basic Navigation

The Local Controller uses a tab-based, hyperlinked user interface. When you mouse over an alphanumeric string or an icon that is a clickable hyper-link, the mouse cursor changes to a pointing finger cursor . Figure 17-2 shows some of the clickable objects on the Dashboard page.

Figure 17-2 Links, Icons, and Filters



Incident ID	Event Type	Matched Rule	Action	Time	Path	Cases
I:10300958 	Inactive reporting device detected 	Try it-Dub05.03.08/09:01:23 	Alert	Aug 30, 2005 10:00:02 AM CDT	  C:119753 (Assigned) Follow-up 	143162

Callouts: 1 points to Incident ID, 2 points to Matched Rule, 3 points to All Rules filter, 4 points to Path icons, 5 points to Cases.

1	Link to the item's detail page or popup window.	2	Query icon links to query page. The corresponding query field is populated with the item.
3	Pulldown lists filter what is displayed.	4	Path icons launch Path or Incident Vector pop-up diagrams.

Click any of the seven tabs to navigate to the pages relevant to the tab's sub-tabs, as shown in Figure 17-3 though Figure 17-8.

Figure 17-3 Summary Tab






Figure 17-4 Incidents Tab

143160

Figure 17-5 Query/Reports Tab


143165

Figure 17-6 Rules Tab


143166

Figure 17-7 Management Tab

143167

Figure 17-8 Administration Tab

143151


Figure 17-9 Help Tab

143394

Help Page

The Help page, as shown in [Figure 17-10](#), provides URLs to online documentation and a feedback form to submit constructive comments to the MARS development engineering team.

Figure 17-10 Help Page



Click **About** to display the software version number running on the MARS.

Click **Documentation** to display URLs to MARS documentation on the Cisco Systems, Inc. website (<http://www.cisco.com>).

Your Suggestions Welcomed

The **Feedback** button appears at the bottom of most pages, as shown in [Figure 17-10](#).

When you click the feedback button, or navigate to the Feedback page, the feedback dialog box appears, as shown in [Figure 17-11](#).

Figure 17-11 Feedback Dialog Box

Standalone: LC20-Doc v4.1 Login: Administrator (pnadmin) :: [Close](#)

Contact Email:
The return email address can be set permanently via User Management.

Subject:

Include log files: ☐

Message:

[Submit](#)


143158

To send your comments to the MARS development engineering team, type in your email address and comments then click **Submit**. When you click the **Include log file** a MARS log file is sent with your message.

Summary Page

From the Summary pages, you can very quickly evaluate the state of the network. The Summary pages include the **Dashboard**, **Network Status**, and **My Reports**, as shown in [Figure 17-12](#).

Figure 17-12 Summary Tab



Dashboard



Note

When you first view the Summary page after upgrading the Local Controller, expect a small delay while the Java Server pages recompile.

Figure 17-13 The Working Areas on the Dashboard



1	Subtabs	5	Tabs
2	Case Bar (Local Controller only)	6	Recent incidents information
3	Links to Cases assigned to you.	7	HotSpot and Attack diagrams
4	Charts		

Recent Incidents

The first feature to notice about the Dashboard are the recent incidents that have fired. The Local Controller comes with pre-defined rules, and these incidents are the result of those rules firing. These rules are generic, globally applicable, and should serve you well as a starting point once you begin to tune the Local Controller.

Figure 17-14 Drilling-down into Incidents

Incident ID	Event Type	Matched Rule	Action	Time	Path	Cases
I:10300958 	Inactive reporting device detected 	Try it-Dub05.03.08/09:01:23 	Alert	Aug 30, 2005 10:00:02 AM CDT	 	C:119753 (Assigned) Follow up 

Arrows from numbered circles point to specific elements in the table:

- 1: Incident ID
- 2: Incident severity icon (Green)
- 3: Event Type
- 4: Query icon in Event Type
- 5: Matched Rule
- 6: Path icon
- 7: Vector icon
- 8: Case icon

1	Link to the Incident sessions detail page	5	Link to the rule details page
2	Incident severity icons  Red—Severe threat  Yellow—Possible threat  Green—Unlikely threat	6	Incident Path icon  launches the topology diagram popup window
3	Link to the Event Type Details page	7	Incident Vector icon  launches the incident attack vector diagram
4	Query icon links to Query page 	8	Link to the View Case page

Sessions and Events

Within a given time window, a session is a collection of events that all share a common end-to-end:

- Source and destination address
- Source and destination port
- Protocol

Event sessionization aggregates event data making it easier to sort and examine. Event sessionization lets the system treat events as single units of information and helps you understand if an attack truly has materialized. It gives you the context of the attack by giving you all the events on that session.

Sessionization works across NAT (network address translation) boundaries – if a session traverses a device that does NAT on that session, the Local Controller is able to sessionize events even if they are reported by two devices on either side of that firewall.

Networks start to show immediate action in the events and sessions categories. Note that the 24 Hour Events table and the Events and Sessions chart are different ways of presenting the same information.

Data Reduction

Data Reduction is a representation of how much event data the Local Controller collapsed into sessions. For example a data reduction of 66% measures three events per session on the average – this number is dependent on many variables particular to your network.

Figure 17-15 Data Reduction

24 Hour Events	
Netflow	442,302
Events	7,664,847
Sessions	5,896,067
Data Reduction	23%

143404

Page Refresh

The Page Refresh Rate polls the Local Controller according to the setting you assign. The default setting is fifteen minutes. The refresh setting remains the same until you log out. This setting only applies to the pages that have the Page Refresh pull-down.

Figure 17-16 Page Refresh

Page Refresh Rate	
15 minutes	▼
24 Hour Events	
Netflow	0
Events	2,132,436
Sessions	462,803
Data Reduction	78%

143401



Note

You can change the refresh rate with the dropdown list.

Diagrams

The Summary page has two diagrams: the Hot Spot Graph and the Attack Diagram. Local Controller uses the configuration and topology discovery information that you provide to generate these diagrams. The following table shows you the icons used in the diagrams.

You can start drilling-down into the diagrams by clicking any of the icons listed in [Table 17-1 on page 17-10](#). You can start drilling-down attack paths in the Attack Diagram by clicking the Path icon . Drilling-down into these diagrams is one of the fastest ways to uncover real-time information about your network.

Figure 17-17 Clickable Hot Spots: Brown = Attackers & Red = Compromised**Note**

Clouds can represent collections of gateways in the Hotspot graph. A gateway cloud is a device that is unknown to the Local Controller. You can discover gateway clouds by clicking them if you have the SNMP information.

Table 17-1 Icons and States in Topology


	Healthy	Attacker	Compromised	Compromised and Attacking
Clouds		—	—	—
Firewall				
Reporting Host				
Host				
IDS				
Network				
Router				
Switch				

To see the diagrams, you need the Adobe SVG viewer plug-in. The Adobe SVG viewer plug-in should automatically install.

**Note**

If you click **No** on the SVG auto-installer, the Local Controller does not prompt you to install it again. If you want to run the auto-installer, open the browser and click **Tools > Internet Options > General > Delete Cookies**.

Figure 17-18 The Hot Spot Graph and Attack Diagram



1	Displays SVG Help	2	Displays clouds for selected devices on a full page
3	Displays all devices on a full page	4	Selects zone to be displayed (Global Controller only)
5	Selects zone to be displayed (Global Controller only)		

Manipulating the Diagrams

- **Right-click** the diagram to zoom in and out, to reset the diagram to its original size, to set the diagram's viewing quality, to search, and to manipulate the SVG image.
- **Alt+click** to use the hand to move the image.
- **Ctrl+click** to use the magnifying glass to zoom in.
- **Ctrl+click and drag** to select an area.
- **Ctrl+shift+click** to use the magnifying glass to zoom out.

**Note**

If the Local Controller discovers an unknown device, it displays that device using a unique name in the form of the string “eth” followed by a hyphen (“-”), followed by the IP address in 32 bit notation, such as “eth-168034561”.

Display Devices in Topology


You can specify how to display a reporting device in the HotSpot Graph. By clicking the icon in the Device Display column, you can specify whether to display the device as an individual node on the graph or collapse it within a cloud. By having a device “hidden” in a cloud, you can cut down on the number of devices displayed in the graph, thus making it easier to read at a higher level.

A cloud identifies a collection of networks for which you do not want to define the complete physical topology. Much like when you draw a network diagram on a piece of paper, you can use a cloud to depict networks in which you have no direct interest, but which are needed to represent to complete the diagram. For example, you may want to display only gateway devices or mitigation devices, representing other reporting devices as part of a cloud.

To toggle the display status of a device, follow these steps:

-
- Step 1** Click **Admin > Security and Monitor Devices**.
- Step 2** Click the icon in the Device Display column of the device that you want to toggle.

Figure 17-19 *The Device Display icons*



The icon changes from a host icon to a host within a cloud or vice versa.

- Step 3** Click **Activate**.

Network Status

The Network Status page is where you come to get the big picture. On the Network Status page, you can see the charts for:

- *Incidents*

Rated by severity.

- *Attacks: All - Top Rules Fired*

Rated by the highest number of incidents fired.

- *Activity: All - Top Event Types*

Rated by the highest numbers of events of that type.

- *Activity: All - Top Reporting Devices*

Rated by the total number of events reported by each security device.

- *Activity: All - Top Sources*

The top IP addresses that appear as session sources, ranked by session count.

- *Activity: All - Top Destinations*


The top IP addresses that appear as session destinations, ranked by session count.

For all of the charts on this page, you can set different time frames, the size of the chart, view the latest report, and so on, by clicking on the buttons in the chart's window.

Reading Charts

These are stacked charts. You can tell which severity of incident your network has most experienced for the day by looking for the dominant shade. In the figure below, low priority green incidents cover less area than high priority red incidents because they have occurred less often.

Figure 17-20 A Day's Events and Netflow with the Legend Displayed




1	Displays values by hour, day, week, month, quarter (the last 3 months), or year.	2	Sets chart to represent the sum of all zones or each individual zone (Global Controller only).
3	Displays a larger version of the chart.	4	Displays the chart legend.
5	The chart legend		

To read the charts most efficiently, note that it is solely the thickness of a particular color that determines its value at that point – and that a spike (or drop) in any particular color could be caused by a spike (or drop) of a different color lower down in the stack.

A perfectly flat line indicates that Local Controller received no data during that time period.


Figure 17-21 A Flat Line in a Week's Top Rules Fired



1 The flat line in the Top Rules Fired chart

In the following Incidents chart, you can see the top incidents for the week, starting eight days in the past.

Figure 17-22 Eight Days of Incidents



1 A more drastic spike in red is not offset by the green incident

2 Incident spikes are built upon each other

My Reports

The My Reports page is where you can choose the reports that you want to view. As long as you are using the Local Controller with your log in name, the reports that you have selected appear here.

To set up reports for viewing

Step 1 Click the **Edit** button on the My Reports page.

Step 2 Select the radio button next to the report that you want to see as a chart.

Step 3 Click **Submit**.

Local Controller now displays the chart that you selected on the My Reports page.



Note

Reports must be scheduled to run periodically, that is, every hour or every day. If you activate a report, allow for some time for the data to accumulate.

You can display any number of charts on the My Reports page, however expect slower loading times for large numbers of charts.

The reports that you can select from are pre-defined. When you create your own reports, you can select those to display. See [Reports, page 20-22](#) for more information.



Case Management

This chapter contains the following sections:

- [Case Management Overview, page 18-1](#)
- [Hide and Display the Case Bar, page 18-3](#)
- [Create a New Case, page 18-4](#)
- [Edit and Change the Current Case, page 18-5](#)
- [Add Data to a Case, page 18-6](#)
- [Generate and Email a Case Report, page 18-7](#)

Case Management Overview

The Case Management feature can capture, combine, and preserve user-selected MARS data within a specialized report called a case. The following data can be added to a case:

- Text annotations
- Incident ID page
- Incident device information (source IP address, destination IP address, reporting device)
- Session Information page
- Query Results page
- Build Report page
- Report Results page
- View Case page (the current case can reference another case)

Any user can create or alter any case. You can assign a case to a MARS user on the same machine, and can change the status of a case to assigned, resolved, or closed. The contents of a case are displayed by category on a single GUI page (View Case), and can be automatically assembled into a single HTML case document. You can email the Case Document to any MARS user account or user group.



Note

When a case is closed, you can still email it, annotate it, add device information, and include a reference to another case.

Case information collected on incidents, sessions, queries, reports and mitigation logs are forensic evidence pertinent to the following:

- Audits (for example, regulatory compliance audits)
- Justifications for modifying ACLs or policy changes
- Notes for MARS false positive tuning
- Examples of allowed and prohibited behavior.

The case preserves and displays the selected data as it appeared when the data was added to the case, regardless of subsequent changes to the MARS state. For example, MARS data can be purged, topology can change from automatic discoveries or vulnerability scanning, and overall configuration can change when you edit rules or reports, but the data reported in the case remains the same as the time it was captured.

**Note**

As of MARS software version 4.1.1 the Case Management feature replaces the incident escalation feature.

The Case Management homepage is the Cases subtab of the Incidents tab as shown in [Figure 18-1](#).

Figure 18-1 Case Management Tab—Local Controller

1 Select Case: No Case Selected... View Cases New Case

2 Case Bar

3 Individual Cases

Case ID	Status	Owner	Summary	Created / Updated
C:121330	New	Martucci, Francesca (francy)	Confetti Attack	Created: Aug 26, 2005 2:01:43 PM CDT Updated: Aug 28, 2005 4:09:17 PM CDT
C:121284	Closed	Administrator (pnadmin)	New Case	Created: Aug 26, 2005 1:47:39 PM CDT Updated: Aug 26, 2005 1:51:18 PM CDT
C:119753	Assigned	Administrator (pnadmin)	Follow-up	Created: Aug 18, 2005 1:43:06 PM CDT Updated: Aug 30, 2005 12:41:55 PM CDT
C:119662	New	Lundell, Norm (nlundell)	Security Team	Created: Aug 16, 2005 8:52:39 AM CDT Updated: Aug 28, 2005 1:29:44 PM CDT
C:118855	New	Administrator (pnadmin)	New Case	Created: Aug 2, 2005 8:48:16 AM CDT Updated: Aug 30, 2005 9:32:44 AM CDT
C:118328	Assigned	McLutt, Blaine (rbmcnutt)	Sample Case for a View	Created: Jul 29, 2005 9:31:15 AM CDT Updated: Aug 2, 2005 8:47:40 AM CDT

1 to 6 of 6 25 per page

Copyright © 2003, 2005 Cisco Systems, Inc. All rights reserved. Summary :: Incidents :: Query / Reports :: Rules :: Management :: Admin :: Help :: Feedback

1	Case Bar	2	Dropdown Display Filters
3	Individual Cases		

All new, assigned, resolved and closed cases can be accessed from the Cases subtab.

To view the contents of a case, click the Case ID number of a case. The View Case page appears, as shown in [Figure 18-2](#).

To generate an HTML document of the **View Case** page content that can be emailed, click **View Case Document** at the bottom of the **View Case** page. Graphs and charts plotted from reports are also captured in the Case Document.

Figure 18-2 The View Case Page—Local Controller

1 → Current Case: C:109418 (Assigned) MARS Defends

2 → View Case: C:109418

Case ID	Status	Owner	Summary	Created / Updated
C:109418	Assigned	Local: Administrator (pnadmin)	MARS Defends	Created: Aug 16, 2005 11:07:22 AM PDT Updated: Aug 29, 2005 9:50:11 AM PDT

3 → Case History

User	Action	Comment	Time
Local: Administrator (pnadmin)	Case Opened		Aug 16, 2005 11:07:22 AM PDT
Local: Administrator (pnadmin)	State Changed: Assigned		Aug 16, 2005 11:07:22 AM PDT
Local: Administrator (pnadmin)	Summary Changed	Initial Summary: Case 4 norm	Aug 16, 2005 11:07:22 AM PDT
Local: Administrator (pnadmin)	Owner Changed: Local: Administrator (pnadmin)	Initial Owner: pnadmin	Aug 16, 2005 11:07:22 AM PDT
Local: Administrator (pnadmin)	Priority Changed: Yellow	Initial Priority: Yellow	Aug 16, 2005 11:07:22 AM PDT
Local: Administrator (pnadmin)	Summary Changed	Case 4 norm2	Aug 16, 2005 11:08:04 AM PDT
Local: Administrator (pnadmin)	Comment	c22	Aug 16, 2005 11:08:04 AM PDT
Local: Administrator (pnadmin)	Session Added: S:5458092		Aug 16, 2005 3:12:29 PM PDT
Local: Administrator (pnadmin)	Device Info Added:	3.1.5.5	Aug 16, 2005 3:12:29 PM PDT
Local: Administrator (pnadmin)	Case Referenced: C:108300 (New) New Case1		Aug 16, 2005 3:17:13 PM PDT
Local: Administrator (pnadmin)	Priority Changed: Red		Aug 16, 2005 3:19:22 PM PDT
Local: Administrator (pnadmin)	Report Added: Activity: All - Top Destination Ports (Peak View)		Aug 16, 2005 3:20:32 PM PDT
Local: Administrator (pnadmin)	Summary Changed	MARS Defends	Aug 29, 2005 9:50:11 AM PDT

4 → Sessions

Display	Session / Incident ID	Events	Source IP/Port	Destination IP/Port	Protocol	Time	Reporting Devices	Path / Mitigation	Tune		
☒	S:5458092, I:53366080, I:53366120, I:53366110, I:53366130, I:53366100	Inactive reporting device detected	0.0.0.0	0	3.1.5.5	0	N/A	Aug 16, 2005 3:00:02 AM PDT	pluto	N/A	False Positive

Devices

Display	Device

1	Case Bar—Identifies current case	2	View Case identifier—Shows the attributes of the case
3	Case History—Log of all changes made to the case	4	Summary of data added to the case

Case Management Considerations for the Global Controller

Case management on the Global Controller differs from the Local Controller implementation as follows:

- Cases are not created on a Global Controller. They can be viewed and modified.
- The Global Controller does not have a Case Bar. All Cases are selected from the Incident -> Cases page.
- The Cases page has an additional dropdown filter to display cases per Local Controller.

Hide and Display the Case Bar

The Case Bar displays by default. When displayed, the Case Bar appears at the top of each page. The Case Bar must be displayed to create or modify a case.

Hiding the Case Bar

To hide the Case Bar, perform the following steps:


- Step 1** Navigate to the **Cases** subtab (**Incidents > Cases**), as shown in [Figure 18-3](#).

Figure 18-3 Case Bar Displayed on the Incidents Page



- Step 2** Click **Hide Case Bar**.
The Case Bar no longer appears on all tabs, as shown in [Figure 18-4](#).

Figure 18-4 Case Bar Hidden on the Incidents Page



Displaying the Case Bar


To Display the Case Bar, follow these steps:

- Step 1** Navigate to the **Cases** subtab (**Incidents > Cases**) as shown in [Figure 18-4](#) .
- Step 2** Click **Show Case Bar**
The Case Bar, as shown in [Figure 18-3](#) now appears on all pages.


Create a New Case

To create a new case, perform the following procedure:

- Step 1** Display the Case Bar as described in the section, [Hide and Display the Case Bar](#).
- Step 2** Click **New Case**.
The Add a New Case Dialog box appears, as shown in [Figure 18-5](#).

Figure 18-5 Add a New Case Dialog Box

- Step 3** Select a severity color, change the state from new to assigned if appropriate, select the owner, replace the default summary name (default is New Case).
Figure 18-5 shows a case with case summary of Example_Case, assigned to the administrator with a yellow priority color (default is Green).
- Step 4** Type or paste any annotations into the text space.
- Step 5** Click **Create New Case**.
 The newly created case is numbered and becomes the current case displayed in the Case Bar as shown in **Figure 18-6**.

Figure 18-6 Case Bar Shows a Newly-Created Case as the Current Case

Proceed to the section [Add Data to a Case](#) for steps on how to combine various data into a single case.

Edit and Change the Current Case

Editing the Current Case

To edit the Current Case complete the following procedure:

- Step 1** Display the Case Bar and click **More**.
The Case Bar Expands to expose the editing options, as shown in [Figure 18-7](#).
See the section [Hide and Display the Case Bar](#) for procedures to display the case bar.

Figure 18-7 Expanded Case Bar



- Step 2** Change the severity, status, owner, or summary of the case as required.
- Step 3** Add an annotation in the text box as required.
- Step 4** Click **Submit**

Deselecting the Current Case

To replace the Current Case case with another, complete the following procedure:

- Step 1** Expand the Case Bar as explained in the previous procedure.
- Step 2** Click **Deselect**.
The Case Bar drop-down list displays **No Case Selected. . .** as shown in [Figure 18-4](#).
- Step 3** To select a different Current Case, select a case from the Case Bar drop down list.

Add Data to a Case

To add data to a case, complete the following steps:

- Step 1** Select the Current Case. See the section [Edit and Change the Current Case](#) for procedures on selecting the Current Case.
- Step 2** Navigate to the page to be captured in the case. In the example, the Query page is selected.
- Step 3** Click **Add this. . .** on the Case Bar.

Figure 18-8 Case Bar Add Button



- Step 4** To verify that the selected data was added to the case, click the case ID number in the Case Bar to display the View Case page.
- In the example shown in [Figure 18-8](#), the selected report should appear in the Reports section of the View Case page. A partial View Case page is shown in [Figure 18-2](#).
-

Generate and Email a Case Report

You can generate a case report of the case data and email the report to any MARS user group or individual user account. The email event is logged in the case history listings on the View Case page.

To add a new user account or user group, see “[Create a New User—Role, Identity, Password, and Notification Information](#)” section on page 22-10.

**Note**

Make sure that the MARS email server is configured. See “[Configure the E-mail Server Settings](#)” section on page 22-4 for further information.

To generate a case report and to email it, follow these steps:

- Step 1** Select a case from the Cases page or from the Case Bar dropdown list.
- Step 2** Click the Case ID number to navigate to the **View Case** page.
- Step 3** Click the check box in an item’s **Include** field to select or deselect that item for inclusion in the Case Document. By default, all items are selected.

**Tip**

Click **Show Include** to show only those items selected for the Case Document. **Show Include** does not function for cases created in Cisco Security MARS version 4.1.1.

- Step 4** Click **View Case Document** at the bottom of the **View Case** page.
- MARS generates and displays the case report.
- Step 5** Click **Email Case** at the bottom of the report page.
- The Case Email dialog box appears, as shown in [Figure 18-9](#).

Figure 18-9 Case Management Email Dialog Box

Choose Where to Email Case Information

Email Subject: Case Doc: C:126597 (New) New Case

Select All

User Groups: All Users Search

Admin
Notification
Operator
Security Analyst

Add Edit Delete

Cancel Submit

- Step 6** Click the check box of the user groups or individual users you want to receive the Case Document, then click **<< Add**.



Tip Select **All Users** from the dropdown menu to display all individual user accounts.

The selected recipients appear in the left-hand area of the dialog box.

- Step 7** Click **Submit** to send the Case Document to the recipients.

The email is sent and the case history is updated to show the email event as the latest item of the case history.



Incident Investigation and Mitigation

An incident is a chain of events that are correlated by a rule to signal an attack upon your network. MARS simplifies and expedites the detection, mitigation, reporting, and analysis of the incident. The Network Summary dashboard and the Incident pages help to detect recent incidents and show the rules and the events that compose them. Mitigation refers to the ability of the MARS to isolate the attacking and compromised network devices by identifying and configuring enforcing devices that act as choke points in the network. Queries and reports reveal the scope of a problem and gather data for analysis and regulatory compliance. All this information can be captured in a case report with Case Management and escalated to the relevant personnel.

Incidents Overview

An attack can consist of a reconnaissance activity (for instance, a port scan), followed by a penetration attempt (such as, a buffer overflow), and followed by malicious activity on the target host (for example, a local privilege escalation attack or the installation of backdoors).

An incident, which is generated by a Local Controller, collects the interesting events that constitute an attack scenario and uses rules to describe them. MARS provides you with pre-defined, system rules—which you can fine tune—and gives you the ability to create your own rules.

Incidents are sub-divided into instances to make it easier for you to investigate the attack scenario. Each instance alone is a full attack scenario.

For example, if your network is probed for a DoS attack and then attacked, a rule fires when it sees the follow up attack. The incident displays the instances of this attack.

Figure 19-1 A DoS probe followed by a DoS attack

Incident ID: 42998483   

Offset	Firing Event / Session / Incident ID	Event Type	Source IP / Port
Instance 1			
3		[1906920] Net Flood TCP 	+ Total: 5
Instance 2			
3	S:45754259, I:42998483  , I:42998484 	[1906910] Net Flood UDP  	10.4.17.4 
Instance 3			
1		[1905037] WWW SGI MachineInfo Info Leak 	10.1.1.21 
1	S:45775179, I:42998480  , I:42998481  , I:42998483  , I:42998487  , I:42998490  , I:42998492  , I:42998493  , I:42998495 	[1905110] WWW SuSE Installed Packages Info Leak  	10.1.1.21 

143431

The Incidents Page

Click the **Incidents** tab to navigate to the Incidents page. The Incidents page displays recent incidents.

Incidents are collections of events and sessions that meet the criteria for a rule, each having helped to cause the rule to fire. An incident's duration only includes the events that contributed to the incident firing.

Figure 19-2 Incidents Navigation

Incident ID	Event Type	Matched Rule	Action	Time	Path	Cases
I:347918013	Deny connection - no xlate	System Rule: Client Exploit - Sysbug Trojan		Sep 22, 2005 2:38:50 PM PDT		
I:342595342	Sudden increase of traffic to a port	System Rule: DoS: Network - Attempt		Sep 22, 2005 2:33:29 PM PDT - Sep 22, 2005 2:33:31 PM PDT		
I:347917458	PIX firewall login failed	System Rule: Password Attack: System - Attempt		Sep 22, 2005 2:33:25 PM PDT - Sep 22, 2005 2:33:26 PM PDT		C:214192 (Assigned) Acc_team

143428

1	The Incident ID— Link to the Incident Detail page.	2	Incident Severity Icon
3	The events that compose the Incident— Launches the Event Type Details popup window.	4	Query icon—Link to the Query page and populates the corresponding query field with the item.
5	The rule that fired to create the incident. Links to the rule page to display the details of the rule.	6	Time range of the incident.
7	Launches the Incident Path and Incident Vector diagrams Click to query on the matched rule	8	Link to the View Case page

The Incident page's table:

- *Incident ID*

An incident's unique ID.

- *Severity*

Low (green), medium (yellow), and high (red) icons.

- *Event Type*

The normalized signature sent from the reporting devices.

- *Matched Rule*

The rule whose criteria were met.

- *Action*

The description of the notification taken when this rule fires (epage, email, etc.)

- *Time*

A single time or a time range (see [Time ranges for Incidents](#), page 19-4 for more information)

- *Incident Path*

The icon that takes you to the incident's path diagram.

- *Incident Vector*

The icon that takes you to the source, event type, and destination diagram.

Time ranges for Incidents


The time column displays both single entries for time (Sep 6, 2003 12:09:54 PM PDT), and time ranges (Sep 6, 2003 12:06:43 PM PDT - Sep 6, 2003 12:06:47 PM PDT).

A single time tells you that all of the firing events were received in the same second. The duration of the incident includes only events that have fired that incident.

Incident Details Page

Clicking the Incident ID takes you to its Incident Details page. The Incident Details page is rich in information and information gathering tools. This page answers questions, such as who did it, what event types happened, when it happened, and to whom it happened.

Figure 19-3 The Incident Details Page



On the top of this page are the tools that let you search for Incident and Session ID and view the Matched Rule.

To Search for a Session ID or Incident ID

- Step 1** Enter the ID into the appropriate field.
 - Step 2** Click the **Show** button.
- To view a partially hidden rule
- Click the Show button next to the Rule Description.

Incident Details Table

Each row of the Incident Details table represents either a session or the information common to a group of sessions. You can see all of the collapsed session information by clicking the plus signs to expand the group. You can expand or collapse all of the incident's information by clicking the **Expand All** or **Collapse All** buttons.

Figure 19-4 Expanding a Row in a Table

Offset	Session / Incident ID	Event Type	Source IP/Port	Destination IP/Port	Protocol	Time	Reporting Device	Reported User	Path / Mitigate	False Positive
1		Built/teardown/permitted IP connection	Groups: 6, Total: 12							
1		Built/teardown/permitted IP connection	Groups: 6, Total: 12							
1		Built/teardown/permitted IP connection	0.0.0.0 0	0.0.0.0 0	TCP	Sep 5, 2005 11:20:05 AM PDT	cherryWall		Total: 4	
1		Built/teardown/permitted IP connection	10.2.3.42 51893	10.4.1.20 18184	TCP	Sep 5, 2005 11:20:09 AM PDT	cherryWall		Total: 2	
1	S:2008826703, I:2009826911, I:2009826889, I:2009826889, I:2009826890	Built/teardown/permitted IP connection	10.2.3.43 52499	10.4.1.251 443	TCP	Sep 5, 2005 11:20:05 AM PDT	cherryWall			False Positive
1		Built/teardown/permitted IP connection	10.4.1.200 1025	10.1.1.189 514	UDP	Sep 5, 2005 11:20:05 AM PDT	cherryWall		Total: 2	
1	S:2008826880, I:2009826911, I:2009826889, I:2009826889, I:2009826890	Built/teardown/permitted IP connection	10.4.2.11 22	10.2.3.33 40222	TCP	Sep 5, 2005 11:20:05 AM PDT	cherryWall			False Positive
1		Built/teardown/permitted IP connection	67.116.29.66 3684						Total: 2	
1	S:200882690, I:2009826911, I:2009826889, I:2009826889, I:2009826890	PIX firewall login failed, TCP access requested to the PIX firewall, TCP or UDP access permitted to the PIX firewall, SSH session disconnected for a reason	10.2.3.33 40224	10.4.5.1 22	TCP	Sep 5, 2005 11:20:05 AM PDT	cherryWall	pix		False Positive

Copyright © 2003, 2005 Cisco Systems, Inc. Summary :: Incidents :: Query / Reports :: Rules :: Management :: Admin :: Help :: Feedback

This high-density information table lets you drill deep into incidents. Click the Query  icon anywhere on this page to query on a particular criteria. Click the Raw Events  icon for raw events for a particular session. You can click the **Tune** link to tune incidents for False Positives, see [The False Positive Page](#), page 19-8 or click the **Mitigate** link to mitigate an attack.

Figure 19-5 Incident Table

1

2

3

Incident ID: 200982691

🔍

📄

Expand All

Collapse All

Offset	Session / Incident ID	Event Type	Source IP/Port	Destination IP/Port	Protocol	Time	Reporting Device	Reported User	Path / Mitigate	False Positive
1		Built/teardown/permitted IP connection	+ Groups: 6, Total: 12							
1	S:200882690, I:200982691 I:2009826889, I:2009826889, I:2009826890	PIX firewall login failed TCP access requested to the PIX firewall TCP or UDP access permitted to the PIX firewall SSH session disconnected for a reason	10.2.3.33 40224	10.4.5.1 22	TCP	Sep 5, 2005 11:20:05 AM PDT	cherryWall	pix	🔴	False Positive

4

5

6

7

8

9

10

11

12

143425

1	Incident ID	2	Severity icon
3	Path and Incident Vector icons. Launch popup windows to display Path and Incident Vector diagrams (L2 or L3 attack path information)	4	Offset number
5	Links to Session and Incident Detail pages of all incidents within the session	6	Links to the Event Type Details pages

7	Launches False Positive popup window	8	Link to the Device Information page
9	Query icon links to Query page	10	Click Device icon to launch popup window to display raw message information
11	Link to the Mitigation Information page	12	Link to the False Positive Tuning page

The following information describes some of the fine points of this table.

- *Instances*

Sometimes rows are split into instances. The *only* relationship among the different instances is that they fired the same rule in the same time frame.

- *Session/Incident ID*

This column shows the sessions that contributed to the incident, and the other incidents those sessions belong to.

- *Events column*

The Events column shows types of the firing events. Multiple firing events of the same types are shown once per session.

- *Time column*

An incident's duration only includes the events that contributed to the incident firing.

False Positive Confirmation

When investigating incidents, you will invariably come across false positive events. In some cases, firing events are classified automatically by MARS as system-confirmed false positives and unconfirmed false positives. Vulnerability scanning often identifies the false positive events, but at times you must investigate events to determine their validity.

To understand the false positive nomenclature and what tasks you are expected to perform within the user interface, we must study the possibilities among three variables surrounding possible attacks: legitimate attack, valid target, and attack detected. We examine these differences in [Table 19-1](#).

Table 19-1 **Attack Type Truth Table**

	Legitimate Attack	Valid Target	Attack Detected
invalid scenario	0	0	0
False Positive	0	0	1
invalid scenario	0	1	0
False Positive	0	1	1
False Negative	1	0	0
Attack/Alarm (noise)	1	0	1
True False Negative	1	1	0
Intrusion/True Alarm	1	1	1

Based on the valid cases in [Table 19-1](#), we can clearly distinguish the false positive terminology:

- A *legitimate attack* is an actual attempt by an attacker to gain access to or information about a specific host using a known exploit.
- A *valid target* is a host that is susceptible to the launched attack. A host can become an *invalid target* if it is properly patched or has some other preventative measure in place, such as a local firewall, virus scanner, or intrusion prevention software that guards against the attack.
- *Attack detected* refers to whether the monitoring device detected the attack and generated an alarm.
- A *false positive* is when the monitoring system generates an alarm for a condition that is benign. In this case, there is no legitimate attack, despite the alarm generation.
- An *unconfirmed false positive* is one where the monitoring system, based on data not available to the reporting device, has determined that an alarm is a false positive. Unconfirmed refers to the fact that the administrator must review and accept or reject the assessment of the false positive.
- A *false negative* is when the monitoring system fails to detect a legitimate attack.
- *Noise* refers to those alarms that are triggered due to attacks against invalid targets. While they can represent real attacks, the target cannot be compromised due to preventative measures. Attacks that fall within the noise category are of secondary importance in terms of investigation and mitigation.
- *Intrusion* identifies a successful attack against the host, where the host is compromised by the attacker.
- A *true false negative* identifies an intrusion that remains undetected by the monitoring system.
- A *true alarm* identifies an intrusion that is detected by the monitoring system.

When a Local Controller receives an event, it is evaluated against the conditions of the defined rules. If the event satisfies the conditions of a rule, then the incident triggers. When an event triggers an incident, we refer to that event as a *firing event*. False positive analysis is performed for such firing events to reduce the number of false alarms.

Using built-in event vulnerability data, learned topology paths, sessionized event data, ACL analysis of layer 2 and 3 reporting devices, supporting data from 3rd-party vulnerability analysis (VA) software (such as Foundstone and eEye), and information that you provide about hosts, MARS analyzes the firing events reported to it determine whether they hold up to a higher-level review.

In the case of MARS, a *system-confirmed false positive* is where, after further analysis, a firing event is determined to be invalid. Example system-confirmed false positives include:

- When an IDS device monitoring the network outside of a firewall reports an attack; however, the firewall drops that session as part of its standard access restrictions. Therefore, the attack never reaches the target.
- Cisco Security Agent detects an attack and blocks it.

An *unconfirmed false positive* is where, after further analysis, the firing event is believed to be invalid primarily due to the attack being against an invalid target. Example unconfirmed false positives include:

- A reporting device reports a valid attack against a host; however, the host is not susceptible to that attack because it targets a different operating system. You can reduce these types of false positives by employing OS fingerprinting technologies on the reporting devices.
- A reporting device reports a valid attack against a host's application; however, the host is not susceptible to that attack because it targets a different application.
- A reporting device reports a valid web attack against TCP port 80, however, dynamic probing determines that no services on the target host listen to TCP port 80.

For unconfirmed false positives, you must manually investigate the alarm and specify in Local Controller whether it is an actual false positive. For actual false positives, you should define a drop rule for the event. Defining a drop rule does not mean that the event is not stored in the database, you

have the option of dropping the event from incident evaluation and either shoring it in the database or not. Whether you store the event in the database or not, events matching the event type and target host can no longer act as firing events. By refining the event processing in this fashion, MARS frees up your time to focus on actual incidents by more accurately correlating events into incidents and reducing noise.

As part of your operational strategy, you should strive to refine event generation and processing to tune out the possibility for false positives. You can perform such tuning at the device level, by refining what traffic or action can generate an event, and at the Local Controller level by providing more information about your network, such as identifying the operating system of hosts attached to the network segments monitored by that Local Controller.

The False Positive Page

To navigate to the False Positives page, click **Incidents**, and click the **False Positives** sub-tab.

The False Positives page is where you can see groupings of False Positives.

You can filter categories by clicking on the **Select False Positive** drop-down list. Your choices are:

- *Unconfirmed false positive type*

For this type, the MARS needs user confirmation to determine if the target host is vulnerable to the event type in question.

- *User confirmed false positive type*

For this type, a user has provided confirmation that a firing event is a false positive.

- *User confirmed positive type*

For this type, a user has provided confirmation that a firing event is a true attack.

- *System determined false positive type*

For this type, the system has determined that a firing event is a false positive.

In the False Positives table, you can see how many sessions the false positive has appeared in, the event type, the false positive status confirmation icons, the event type information icon, the destination IP and its port, the destination IP information icon, its protocol, zone, and you can see the sessions that are related to the false positive.

Figure 19-6 False Positive Table

Session Count	Event	Destination IP/Port	Protocol	Zone	Related Sessions
192	[1905035] WWW HylaFAX Faxsurvey Command Exec  	10.4.17.2  80 	TCP 	CA	Show 

1 2 3 4 5 6

143423

1	Link to the Event Type Details page	2	Query icon links to the Query page and automatically populates the corresponding Query field
3	False Positive type and severity icon	4	Launches the Security Device Information popup window
5	Launches Port Information popup window	6	Launches False Positive Sessions Details popup window

The following table shows false positive status confirmation and severity icons: Tuning False Positives

Icon	Description
  	Low, medium, and high severity false positives that require confirmation.
  	Low, medium, and high severity user determined false positives.
  	Low, medium, and high severity system determined false positives.

From the Incidents page or the False Positives page, you can tune false positives – to verify if they are true or false.

To Tune a False Positive

-
- Step 1** Click one of the **Confirm False Positive** icons.   
 - Step 2** On the False Positive Confirmation page, review the information.
 - Step 3** If you decide that the event type is a false positive, click the **Yes** radio button, and follow the steps in: [To Tune an Unconfirmed False Positive to False Positive, page 19-9](#).
 - Step 4** If you decide that the event type is a true positive, click the **No** radio button, and follow the steps in: [To Tune an Unconfirmed False Positive to True Positive, page 19-9](#).
-

To Tune an Unconfirmed False Positive to False Positive

-
- Step 1** After you determine that a false positive is false, and you have clicked the **Yes** button, click **Next**.
 - Step 2** On the next page, decide whether or not you want MARS to keep this event type in the database by selecting the appropriate radio button:
 - **Dropping these events completely** (that stops logging those events)
 - **Log to DB only** (that logs the events to the DB)
 - Step 3** Once you have decided, click the **Next** button.
 - Step 4** On the next page, carefully review the information for the false positive and the new rule.
 - Step 5** When you are ready to commit this new information to the appliance, click the **Confirm** button.
-

To Tune an Unconfirmed False Positive to True Positive

-
- Step 1** After you determine that a false positive is true, and you have clicked the **No** button, click **Next**.
 - Step 2** Make a final confirmation that this is a true positive, and click the **Confirm** button.
-

To Activate False Positive Drop Rules

After you have completed tuning false positives, click **Activate** to immediately implement the changes.

Mitigation

Mitigation refers to the action of limiting an attacking network element's access to the network by modifying the configuration of an enforcement device, usually a switch, router, or firewall. CS-MARS can perform the following actions related to mitigation:

- Identify attacking and compromised hosts
- Plot Layer 2 and Layer 3 topology of the affected network segment to identify mitigation points and enforcement devices
- Recommend configuration commands for Layer 2 and Layer 3 enforcement devices
- Push (that is, download) recommended configuration commands to supported Layer 2 devices

With Telnet, SSH, or SNMP access to switches and routers, CS-MARS can recommend and push mitigation configurations to enforcement devices, as well as generate interactive topology and incident path diagrams. Without Telnet, SSH, or SNMP access, some mitigation information can still be obtained from Cisco switches running specific IEEE 802.1X Port Based Network Access Control protocol configurations, but recommended mitigation commands must be configured manually on the enforcement devices. See [Layer 2 Path and Mitigation Configuration Example, page 19-17](#) for further information and procedures for configuring Layer 2 devices to receive CS-MARS mitigation commands.

Static and Dynamic Network Information

Topology information obtained from access to relatively permanent Layer 2 and Layer 3 devices is called Static Information in the HTML interface. Dynamic Information refers to frequently changing information such as host names, or DHCP-leased IP addresses obtained through devices or agents that report dynamic events, such as 802.1X access control configurations, the Cisco Security Agent, or other security suite software. The CS-MARS can determine a mitigation point and an enforcement device if a Cisco 802.1X-enabled switch is running DHCP-snooping with RADIUS authentication through a Cisco Access Control Server (ACS). When a DHCP-snooping transaction is completed, the switch sends a log message to the ACS. The ACS logs are sent to the CS-MARS to report the Source IP address, user name, connection start and stop times, physical interface, and MAC address of each 802.1X client. Because 802.1X clients are often mobile, remember that 802.1X mitigation actions can occur only when the attacking host is currently connected to the network.



Note

For some 802.1X switch configurations, it is not possible for CS-MARS to determine the correct physical interface to which to push a mitigation command. This occurs for switches, such as the Cisco Catalyst 3550 Multilayer switch, where a FastEthernet and a Gigabit Ethernet port can have the same *module/port* designation (for example, 0/1). Because CS-MARS receives only the *module/port* information from the Cisco ACS logs, it cannot identify the specific port to mitigate. The following message appears in these circumstances:

No mitigation possible. Enforcement device exists but interface names conflict. Determine appropriate interface and mitigate manually.

802.1X Mitigation Example


In this procedure, an incident is observed on the Network Summary page, as shown in [Figure 19-7](#), and mitigated through 802.1X network mapping.

Prerequisites for Mitigation with 802.1X Network Mapping

To perform mitigation with 802.1X network mapping with CS-MARS, the following prerequisites are required:

- Cisco switch running Cisco CatOS or IOS and configured with IEEE 802.1X Port Based Network Access Control protocol
- The switch Reporting IP address must be configured on the CS-MARS Security and Monitoring Information page (**Admin > Security and Monitor Devices**).
- Cisco DHCP-Snooping enabled on the switch
- The switch performs Remote Access Dial-In User Service (RADIUS) authentication, authorization, and accounting through a Cisco Access Control Server (ACS).
- The Cisco ACS is running pnLogAgent to send logs to CS-MARS
- The Cisco ACS is configured to log Update (Watchdog) packets

Figure 19-7 Summary Page Displaying Incident to Mitigate




141315

Procedure for Mitigation with 802.1X Network Mapping

- Step 1** Click the Incident ID of the recent incident to Mitigate.
- Step 2** Click on the Incident ID to display the session summaries, shown in [Figure 19-8](#).

Figure 19-8 Incident Detail Page Displaying Red Mitigation Icon



Step 3 Click the red path information icon in the **Path/Mitigation** column.

The Mitigation pop-up window appears, with any possible Static topology and mitigation information, as shown in Figure 19-9.

CS-MARS recommends enforcement devices and mitigation commands. For static information, if the network is entirely discovered and CS-MARS has command level access to a Layer 2 enforcing device, the Push button appears red, otherwise it is gray. In Figure 19-9, CS-MARS does not have sufficient static information to identify a Layer 2 enforcement device, but can suggest mitigation commands for discovered Layer 3 devices (Cisco PIX firewall, and a Cisco router). Layer 3 mitigation commands must be configured manually on the Layer 3 devices.

Figure 19-9 Path Information Pop-up Window

Static information utilizes discovered Layer 3 and 2 network topology information to determine the optimal mitigation point.

Static Info **Dynamic Info**

Enforcement Devices: **S16229417 Path**

Suggested: cherryWall
Alternate: labCoreRouter1

Enforcement Device: **cherryWall**, Suggested
Default gateway: 67.116.29.125

L3 Enforcement Device Information

Device	Type	Manager	Children	Log To	Collects From	Info
cherryWall	Cisco PIX 6.3	PK-MARS on primars		PK-MARS on primars		

Interface Information

Direction	Interface Name	MAC Address	MAC Update Time
Inbound	outside	00:aa:00:00:00:0e	Oct 4, 2005 5:02:37 PM PDT
Outbound	outside	00:aa:00:00:00:0e	Oct 4, 2005 5:02:37 PM PDT

Recommended L3 Policies/Commands

access-list CSE-acl-outside
deny icmp host 67.116.29.117 host 10.4.2.1

Or

access-list CSE-acl-outside
deny icmp host 67.116.29.117 any

Or

plus 67.116.29.117 10.4.2.1 icmp

Push Cancel

Step 4 Click **Dynamic Info** to view Layer 2 mitigation recommendations derived from 802.1X configurations. The Dynamic Mitigation window appears with host name, IP address, MAC address, and connection status as shown in [Figure 19-10](#).

Figure 19-10 Dynamic Mitigation Information


Local Controller: pnmars/LC10.1.1.189 v4.1 Login: Local: Administrator (pnadmin) ::

Dynamic information utilizes current host location information reported by AAA servers and network devices as a result of enforcing mechanisms such as 802.1X, Cisco Network Admission Control etc.

↓

Static Info	Dynamic Info
Host IP Address: 20.1.1.210 Host MAC Address: 00-11-11-31-e3-48 Cisco ACS NAS IP: 10.1.1.243 Cisco ACS NAS Port: 313 AAA User Name: cisco	

Enforcement Device:

Enforcement Device (Device Name:Module/Port)	Start Time	End Time	Update Time
dot1x_switch:3/13	Oct 3, 2005 2:18:53 PM PDT	Present	Oct 3, 2005 2:18:53 PM PDT

Recommended Policy/Command

```

configure t
interface 3/13
shutdown
          
```


Copyright © 2003, 2005 Cisco Systems, Inc.
All rights reserved. 141313


Step 5 Review the enforcement device.

Step 6 Review the Recommended Policies/Commands.

Step 7 Click **Push** to download the recommended mitigation command to the enforcement device. The mitigation confirmation dialog appears, as shown in [Figure 19-11](#).
If the Push button is gray, the mitigation command must be manually configured on the enforcement device.

**Note**

The **Push** button is red and functional when the 802.1X target host is present on the network, and CS-MARS has command access to the enforcement device otherwise, it appears gray and is not functional.

Figure 19-11 Mitigation Confirmation Dialog


The screenshot shows a web interface for incident management. At the top, there's a header bar with 'INCIDENTS' on the left and a login/status area on the right showing 'Login: Administrator, Administrator (pnadmin) :: Jun 14, 2004 3:50:57 PM PDT ::' and a 'Close' button. Below the header, the main content area is titled 'Download Mitigation Command'. It contains a form with the following fields: 'Device Name:' (CatSw), 'Port/Interface Name:' (5/9), 'Access Type:' (SNMP), 'SNMP RW Community String:' (a masked field with seven dots), and 'Policy/Command:' (set port disable 5/9). Below this form is a confirmation dialog box with the text 'Are you sure you want to download the mitigation command to this device?' and two buttons: 'Yes' and 'No'. On the far right edge of the dialog box, the number '143376' is visible.

Step 8 Click **Yes** to confirm.

Display Dynamic Device Information

To display current, session, and all historical information for an IP address on an 802.1X connection, follow these steps:

- Step 1** Click on the Incident ID to display the session summaries as shown in [Figure 19-8](#).
- Step 2** Click on the **Source IP/Port** or **Destination IP** link of a session.
When examining an attacking host, the Source IP address is more relevant.
- Step 3** The current connection information pop-up window appears to display any static connection information.
- Step 4** Click **Dynamic Info** to display current connection information, as shown in [Figure 19-11](#).
Dynamic information can be derived from 802.1X configurations, Cisco Security Agents, or from other security software suites. The current connection information is the most recent network information available for the selected IP address.
- Step 5** Click **Session** to display the connections related to the specific session, as shown in [Figure 19-13](#).

Figure 19-12 Dynamic Information—Current Connection Status

CISCO SYSTEMS

Local Controller: pnmars/LC10.1.1.189 v4.1

Login: Local: Administrator (pnadmin) :: Close

View CasesNew Case

Dynamic information includes host location information reported by AAA servers and network devices as a result of enforcing mechanisms such as 802.1x, Cisco Network Admission Control etc.

Static Info

Dynamic Info

IP Address: 20.1.1.210

current

session

all

Host Name	MAC Address	AAA User	Enforcement Device (IP:Module/Port)	Reporting Device	Start Time	End Time	Update Time
N/A	00-11-11-31-E3-48	cisco	20.1.1.1:0/2	dot1x ACS (Cisco,ACS,3.x)	Oct 4, 2005 4:12:37 PM PDT	Present	Oct 4, 2005 4:12:37 PM PDT

Copyright © 2003, 2005 Cisco Systems, Inc. All rights reserved.

Feedback

Step 6 Click **All** to display the entire dynamic information for the specified IP address, as shown in Figure 19-13.

Figure 19-13 Dynamic Information History of a Specified IP Address

CISCO SYSTEMS

Local Controller: pnmars/LC10.1.1.189 v4.1

Login: Local: Administrator (pnadmin) :: Close

View CasesNew Case

Dynamic information includes host location information reported by AAA servers and network devices as a result of enforcing mechanisms such as 802.1x, Cisco Network Admission Control etc.

Static Info

Dynamic Info

IP Address: 20.1.1.210

current

session

all

Host Name	MAC Address	AAA User	Enforcement Device (IP:Module/Port)	Reporting Device	Start Time	End Time	Update Time
N/A	00-11-11-31-E3-48	N/A	20.1.1.1:0/18	dot1x ACS (Cisco,ACS,3.x)	Sep 30, 2005 3:55:00 PM PDT	Sep 30, 2005 3:59:00 PM PDT	Sep 30, 2005 3:55:00 PM PDT
N/A	N/A	cisco	20.1.1.1:N/A	dot1x ACS (Cisco,ACS,3.x)	Sep 30, 2005 3:55:00 PM PDT	Sep 30, 2005 4:44:14 PM PDT	Sep 30, 2005 3:55:00 PM PDT
N/A	00-11-11-31-e3-48	N/A	10.1.1.243:3/14	dot1x ACS (Cisco,ACS,3.x)	Sep 30, 2005 3:59:01 PM PDT	Sep 30, 2005 4:44:14 PM PDT	Sep 30, 2005 3:59:01 PM PDT
N/A	N/A	N/A	10.1.1.243:3/14	dot1x ACS (Cisco,ACS,3.x)	Oct 3, 2005 11:16:55 AM PDT	Oct 3, 2005 2:18:52 PM PDT	Oct 3, 2005 11:16:55 AM PDT
N/A	00-11-11-31-e3-48	N/A	10.1.1.243:N/A	dot1x ACS (Cisco,ACS,3.x)	Oct 3, 2005 11:16:55 AM PDT	Oct 4, 2005 4:42:27 PM PDT	Oct 4, 2005 4:09:47 PM PDT
N/A	N/A	cisco	10.1.1.243:N/A	dot1x ACS (Cisco,ACS,3.x)	Oct 3, 2005 11:16:55 AM PDT	Oct 4, 2005 4:50:17 PM PDT	Oct 4, 2005 4:09:47 PM PDT
N/A	N/A	N/A	10.1.1.243:3/13	dot1x ACS (Cisco,ACS,3.x)	Oct 3, 2005 2:18:53 PM PDT	Oct 4, 2005 4:09:46 PM PDT	Oct 3, 2005 2:18:53 PM PDT
N/A	N/A	N/A	10.1.1.243:3/14	dot1x ACS (Cisco,ACS,3.x)	Oct 4, 2005 4:09:47 PM PDT	Oct 4, 2005 4:42:27 PM PDT	Oct 4, 2005 4:09:47 PM PDT
N/A	00-11-11-31-E3-48	N/A	20.1.1.1:0/1	dot1x ACS (Cisco,ACS,3.x)	Oct 4, 2005 4:42:28 PM PDT	Oct 4, 2005 4:50:17 PM PDT	Oct 4, 2005 4:42:28 PM PDT

Step 7 Click the **Push** button if available or mitigate from the device. If you select the push button, a confirmation screen appears.



Note

To mitigate a device of Access Type SNMP you must have the SNMP Read/Write Community String.

Click the **Yes** button to confirm the mitigation command and have it take effect.

Virtual Private Network Considerations

Currently, MARS cannot display accurate Path/Mitigation information or compute the complete route of an attack originated by a host with a source IP address on a virtual private network (VPN). MARS can identify the attacking host if the VPN IP address of the host was supplied by a Cisco 3000 Series VPN Concentrator configured as a MARS reporting device.

**Note**

You must be able to recognize from your knowledge of your network that the IP address of the attacking host is an IP address allocated to a VPN.

To identify a host attacking from a VPN, perform a query of “Cisco VPN User connected/disconnected” events for the Cisco VPN Concentrator device. The attacking host name or next network element is disclosed in the raw messages of the events.

Layer 2 Path and Mitigation Configuration Example

This section provides a starting point for configuring MARS to perform Layer 2 (L2) path analysis and mitigation using a Cisco switch. It contains the following sections:

- [Prerequisites for Layer 2 Path and Mitigation, page 19-17](#)
- [Components Used, page 19-17](#)
- [Network Diagram, page 19-18](#)
- [Procedures for Layer 2 Path and Mitigation, page 19-19](#)
- [Add the Cisco Catalyst 6500 with SNMP as Access Type \(Layer 2 only\), page 19-20](#)
- [Add the Cisco 7500 Router with TELNET as the Access Type, page 19-21](#)
- [Verify the Connectivity Paths for Layer 3 and Layer 2, page 19-22](#)
- [Perform Mitigation, page 19-26](#)

Prerequisites for Layer 2 Path and Mitigation

- You need to have the SNMP community strings and IP addresses for the Layer 2 switches and routers.
- You must have STP (Spanning Tree Protocol) configured correctly on the switches.


Components Used

- a Cisco Catalyst 5000 with SNMP access enabled
- a Cisco Catalyst 6500 for Layer 2 with SNMP access enabled
- a Cisco 7500 Router with SNMP or TELNET access enabled
- a MARS running software Version 2.5.1

Network Diagram

This section uses the network setup shown in the [Figure 19-14](#).

Figure 19-14 Network Setup



Mitigation uses the Layer 2 path data obtained via SNMP or Telnet protocol to download a mitigation command from the MARS to the device. The Layer 2 path is based on MAC addresses, the Layer 2 forwarding table, and the Layer 3 path. MAC addresses and the Layer 2 forwarding table are obtained using SNMP.

To make the Layer 2 path and mitigation work correctly:

- The associated routers must be discovered via SNMP or a combination of SNMP and Telnet, including the MSFC module in the Catalyst switch.
- The SNMP community string is necessary for L2 switches to be discovered



Note

L2 devices must be added manually; there is no automatic discovery for these devices. Make sure all the L2 devices (switches) have the SNMP RO community strings specified in the web interface, even if the access type is not SNMP. The SNMP RO community string is always required on Layer 2 devices for L2 mitigation.

- If the switches are interconnected, make sure STP (Spanning Tree Protocol) is enabled and configured on them.

For example, given a topology such as the one in the preceding figure, follow these instructions to discover these devices.

Procedures for Layer 2 Path and Mitigation

Add the Cisco Catalyst 5000 with SNMP as the Access Type.

Step 1 Click **Admin > Security and Monitor Devices > Add**.

Figure 19-15 Configure Cisco Switch CatOS

Device Discovery-Cisco Switch-CatOS ANY

Note:

1. Enter the reporting IP (the IP address where events originated from) to ensure that the system processes the events.
2. * denotes a required field.

Device Type: Cisco Switch-CatOS ANY

Supervisor Module

→ *	Device Name:	CatSw
→ *	Access IP:	10 1 1 11
→ *	Reporting IP:	
→ *	Access Type:	SNMP
	Login:	
	Password:	
	Enable Password:	
	Config Path:	
	File Name:	
	SNMP RO Community:	MySNMPCommStr

Test Connectivity

Cancel

Submit

143361

Step 2 From the **Device Type** drop-down list, select **Cisco Switch-CatOS ANY**.

Step 3 Enter the **Device Name** of the switch.

Step 4 Enter the **Access IP** address and **Reporting IP** address (the IP address of the device as it appears to the MARS) of the switch. The **Reporting IP** address is usually the same as the **Access IP** address, but if you are using FTP as Access Type, it must be a different IP address. The **Reporting IP** address is required if the device is sending syslog data to the MARS

Step 5 From the **Access Type** drop-down list, select **SNMP** or **TELNET**. Note that which fields need to be completed, along with which you can complete, depend on which Access Type you select.

SNMP:

- For the **Login ID**, enter the user name and **Password** needed to access the switch.
- For **Enable Password**, enter the password to get into Cisco enable mode.

- Enter its **SNMP RO Community**.

TELNET:

- For the **Login ID**, enter the user name and **Password** needed to access the switch.
- For **Enable Password**, enter the password to get into Cisco enable mode.
- Enter its **SNMP RO Community**.

Step 6 Click the **Test Connectivity** button to have the MARS discover the device.

Step 7 Click the **Submit** button.

Add the Cisco Catalyst 6500 with SNMP as Access Type (Layer 2 only).

Step 1 Click **Admin > Security and Monitor Devices > Add**.

Figure 19-16 Configure Cisco Switch CatOS

Device Discovery-Cisco Switch-CatOS ANY

Note:

1. Enter the reporting IP (the IP address where events originated from) to ensure that the system processes the events.
2. * denotes a required field.

Device Type: Cisco Switch-CatOS ANY

Supervisor Module

→ *Device Name:	KittenSw
→ *Access IP:	10 1 1 12
→ *Reporting IP:	
→ *Access Type:	SNMP
Login:	
Password:	
Enable Password:	
Config Path:	
File Name:	
SNMP RO Community:	MySNMPCommStr

Test Connectivity

Cancel

Submit

143386

Step 2 From the **Device Type** drop-down list, select **Cisco Switch-CatOS ANY**.

Step 3 Enter the **Device Name** of the switch.

Step 4 Enter the **Access IP** address and **Reporting IP** address of the switch. The **Reporting IP** address is usually the same as the **Access IP** address.

SNMP:

- For the **Login ID**, enter the user name and **Password** needed to access the switch.

- For **Enable Password**, enter the password to get into Cisco enable mode.
- Enter its **SNMP RO Community**.

TELNET:

- For the **Login ID**, enter the user name and **Password** needed to access the switch.
- For **Enable Password**, enter the password to get into Cisco enable mode.
- Enter its **SNMP RO Community**.

Step 5 Click the **Test Connectivity** button to have the MARS discover the device.

Step 6 Click the **Submit** button.

Add the Cisco 7500 Router with TELNET as the Access Type

Step 1 Click **Admin > Security and Monitor Devices > Add**.

Figure 19-17 *Configure Cisco IOS 12.2*

Device Discovery–Cisco IOS 12.2

Note:

1. Enter the reporting IP (the IP address where events originated from) to ensure that the system processes the events.
2. * is denotes a required field.

Device Type: Cisco IOS 12.2

→ *Device Name:	MainRouter
→ *Access IP:	10 1 1 1
→ *Reporting IP:	10 1 1 1
→ *Access Type:	TELNET
Login:	myuserid
Password:	*****
Enable Password:	*****
Config Path:	
File Name:	
SNMP RO Community:	MySNMPCommStr

Test Connectivity

Cancel

Submit

143369

Step 2 From the **Device Type** drop-down list, select **Cisco Switch-IOS 12.2**.

Step 3 Enter the **Device Name** of the switch.

Step 4 Enter the **Access IP** address (optional) and **Reporting IP** address of the switch. The **Reporting IP** address is usually the same as the **Access IP** address, but if you are creating an FTP device it must be a different IP address.

If you have entered an Access IP address, from the **Access Type** pull-down menu, select **FTP**:

FTP:

- For the **Login ID**, enter the user name and **Password** needed to access the switch.
- For **Config Path**, enter the path of the configuration file on the FTP server.
- For **File Name**, enter the switch configuration file name on the FTP server.
- Enter its **SNMP RO Community**.

SNMP:

- For the **Login ID**, enter the user name and **Password** needed to access the switch.
- For **Enable Password**, enter the password to get into Cisco enable mode.
- Enter its **SNMP RO Community**.

SSH:

- For the **Login ID**, enter the user name and **Password** needed to access the switch.
- For **Enable Password**, enter the password to get into Cisco enable mode.
- Enter its **SNMP RO Community**.

TELNET:

- For the **Login ID**, enter the user name and **Password** needed to access the switch.
- For **Enable Password**, enter the password to get into Cisco enable mode.
- Enter its **SNMP RO Community** (mandatory).

Step 5 Click the **Test Connectivity** button to have the MARS discover the device.

Step 6 Click the **Submit** button.

Verify the Connectivity Paths for Layer 3 and Layer 2

Once you have a session, you can view the Layer 3 and Layer 2 topology paths. There are several ways to obtain a session.

- **To view sessions that are part of an Incident:**

Step 1 Click the **Incidents** tab to navigate to the Incidents page. Click an **Incident ID** of an incident you want to view (in this example we use Incident number 356120290). The [Incident Details](#) screen appears.

Figure 19-18 Incident Details screen

Matched Rule: System Rule: Server Attack: RPC - Success Likely
Description: This correlation rule detects specific attacks on RPC services on a host followed by suspicious acti... [Show](#)

Offset	Open	Source IP	Destination IP	Service Name	Event	Device	Severity	Counts	Zone	Close	Action/Operation	Time-range
		ANY	ANY	ANY	System Rule: Server Attack: RPC - Success Likely	ANY	ANY	1	ProtegoHQ			0hh:30mm:0ss

Incident ID: 356120290

Offset	Session / Incident ID	Event Type	Source IP/Port	Destination IP/Port	Protocol	Time	Zone	Reporting Device	Graph	False Positive	Mitigation	
1	S:372321252, I:356120249, I:356120250, I:356120253, I:356120254, I:356120255, I:356120257, I:356120258, I:356120261, I:356120262, I:356120266, I:356120267, I:356120269, I:356120274, I:356120277, I:356120278, I:356120279, I:356120282, I:356120283, I:356120287, I:356120290	Windows LSARPC Access	67.125.41.172 3077	10.4.14.2 445	TCP	Jun 21, 2004 12:56:55 PM PDT	ProtegoHQ	ids3			Tune	Mitigate
2	S:372468056, I:356120290, I:356120293	Windows RPC DCOM Overflow, Windows SMB/RPC NoOp Sled	10.1.252.250 3967	65.54.143.118 135	TCP	Jun 21, 2004 1:31:40 PM PDT	ProtegoHQ	firewall			Tune	Mitigate
3		Windows LSASS RPC Overflow	+ Total: 2									
3	S:372468056, I:356120290, I:356120293	Windows RPC DCOM Overflow, Windows SMB/RPC NoOp Sled	10.1.252.250 3967	65.54.143.118 135	TCP	Jun 21, 2004 1:31:40 PM PDT	ProtegoHQ	firewall			Tune	Mitigate
5		Net Flood UDP	10.4.14.2	10.1.1.132								
5		Net Flood TCP	10.4.14.2	+ Total: 3								

143374

Step 2 In the [Incident Details screen](#), in the same row as the Event Type you want to examine (in this example we use Windows RPC DCOM Overflow), click the graph icon under the Graph column to view the topology paths.

- To view sessions by performing a Query:

Step 1 Click **QUERY / REPORTS** and submit a query using the appropriate query criteria. Note that in our example, we limit the scope of the query so it runs faster. In the following [Query Event Data screen](#) we use the result format **All Matching Sessions** and query events from **Source IP 10.1.252.250** and **Destination IP 65.54.153.118** over the last **10** minutes.

Figure 19-19 Query Event Data screen**Query Event Data**

Click the cells below to change query criteria:

Query type: *Sessions ranked by Time, 0hh:10mm:0ss* [Edit](#) [Clear](#)

Source IP	Destination IP	Service	Events	Device	Severity	Zone	Operation	Rule	Action	Reported User
H-10.1.252.250	H-65.54.153.118	ANY	ANY	ANY	ANY	ANY	None	ANY	ANY	ANY

Keywords: [None] [Edit](#)

[Apply](#)

Result Format: All Matching Sessions [v](#)

Order/Rank By: Time [v](#)

Filter by Time:

☒ Last: 0 Days 0 Hrs 10 Mins

☐ Start: 2004 June 22 18 Hrs 38 Mins

End: 2004 June 22 18 Hrs 48 Mins

☐ Real Time

Use Only Firing Events: ☐

Maximum rank returned: 100

[Apply](#)

Step 2 After you **Apply** changes to and **Submit** your query, the [Query Results screen](#) appears.

Figure 19-20 Query Results screen

Query Event Data

Click the cells below to change query criteria:

Query type: *Sessions ranked by Time, 0hh:10mm:0ss* [Edit](#) [Clear](#)

Source IP	Destination IP	Service	Events	Device	Severity	Zone	Operation	Rule	Action	Reported User
H-10.1.252.250	H-65.54.153.118	ANY	ANY	ANY	ANY	ANY	None	ANY	ANY	ANY

Keywords: [None] [Edit](#)

[Save As Report](#) [Save As Rule](#) [Clear](#) [Apply](#) [Submit](#)

Query Results [v](#) [v](#)


[Expand All](#) [Collapse All](#)

Session / Incident ID	Events	Source IP/Port	Destination IP/Port	Protocol	Time	Zone	Reporting Devices	Graph	False Positive	Mitigation
S:381559066	Built/teardown/permitted IP connection v	10.1.252.250 v	65.54.143.118 v		Total: 5					
	Windows RPC DCOM Overflow v	10.1.252.250 v 1421 v	65.54.143.118 v 80 v	TCP v	Jun 22, 2004 5:31:15	ProtegoHQ	firewall v	v	Tune	Mitigate


Step 3 In the [Query Results screen](#), in the same row as the Event Type you want to examine (in this example we use Windows RPC DCOM Overflow), click the icon under the Graph column to view the topology paths. The first topology path to appear is the [Layer 3 topology graph](#):

Figure 19-21 Layer 3 topology graph

Topology Path Graph



Under **Topology Path Graph**, click the **Layer 2 Path** button to view the [Layer 2 topology graph](#):

Figure 19-22 Layer 2 topology graph

Perform Mitigation

Once you identify the compromised host (in this example, **10.1.252.250** connected to **CatSw**), it is critical to prevent it from attacking other hosts in the same subnet or other parts of the network. The MARS provides one-click mitigation that lets you isolate the compromised host from the rest of the network.

To perform mitigation, perform these steps:

- Step 1** On the [Incident Details screen](#), click the **Mitigate** link that corresponds with the **Session** or **Event Type** you want to mitigate (in this case, **Windows RPC DCOM Overflow**). The [Mitigation Information screen](#) appears.

Figure 19-23 Mitigation Information screen

INCIDENTS | Login: Administrator, Administrator (pnadmin) :: Jun 14, 2004 3:50:57 PM PDT :: [Close](#)

Mitigation Information

Enforcement Devices	Enforcement Device - Suggested
CatSw (L2) (suggested)	Name: CatSw
KittenSw (L2) (alternate)	Device type: Cisco Switch-CatOS ANY
MainRouter (alternate)	Outbound Interface: sc0
firewall (alternate)	IP Address: 10.1.1.11
	Mac Address: 00:60:47:f8:7a:ff Jun 22, 2004 5:31:15 PM PDT
	Zone: CA
	Managed by: pnmars
	Status: Active
	Default gateway: 10.1.1.1

Recommended Policy/Command

☒ set port disable 5/10

[Push](#) [Cancel](#)


143377

This screen contains information about the device, along with recommended policies or commands for mitigating the compromised host (in the example, 10.1.252.250).

- Step 2** If the device where the mitigation command to be downloaded is a Layer 2 device (such as in the example [Mitigation Confirmation Dialog](#)), a red **Push** button appears that you can click to mitigate the compromised host. If you select the push button, the [Mitigation Confirmation Dialog](#) appears.

**Note**

If the device where the mitigation command to be downloaded is a Layer 3 device, the **Push** button shown in red on the [Mitigation Information screen](#) is greyed out and you must use the suggested commands directly on the device to mitigate the compromised host.

Figure 19-24 Mitigation Confirmation screen


INCIDENTS | Login: Administrator, Administrator (pnadmin) :: Jun 14, 2004 3:50:57 PM PDT :: [Close](#)

Download Mitigation Command

Device Name:	CatSw
Port/Interface Name:	5/9
Access Type:	SNMP
SNMP RW Community String:	*****
Policy/Command:	set port disable 5/9

Are you sure you want to download the mitigation command to this device?

[Yes](#) [No](#)

143376

**Note**

The SNMP RW community string must be enabled for the MARS to download a mitigation command to a device using the Access Type SNMP.

Step 3

Click **Yes** to confirm the mitigation of the device.



CHAPTER 20

Queries and Reports

This chapter discusses the following topics:

- [Queries](#)
- [Viewing Events in Real-time](#)
- [Perform a Long-Duration Query Using a Report](#)
- [Perform a Batch Query](#)
- [Reports](#)

Queries

On the Query page, you can run reports as on-demand queries, or create your own query. Many links from other pages bring you to the query page, which then partially populate the query's criteria. Once you have submitted a query, you can save it as a report or a rule.

Figure 20-1 *The Local Controller Query Table*

1	Click to set the query type and time range criteria.	2	Click Clear to return query values to default values.
3	Quick query fields permit entry of values without opening dialog box for the field.	4	Click on a field value to open the dialog box for that field.
5	Save the query as a report or as a rule.	6	Click Submit Inline to run the query.

To Run a Quick Query

- Step 1** Enter a source IP, destination IP, or a service into the quick query field.
- Step 2** Click the **Submit Inline** button to run the query.

Figure 20-2 Running a Quick Query

Source IP	Destir
ANY	ANY
10.2.2.2	143441

To Run a Free-form Query

- Step 1** Enter a source IP, destination IP, or a service into the quick query field.

Figure 20-3 Running a free-form query

Specify raw message keywords:

Open (	Search String	) Close	Operation	Highlight
	pop3		OR	
	imap		None	
			AND	
			OR	
			NOT	
			None	
			None	
			None	
			None	
			None	
			None	
			None	
			None	

- Step 2** Click the name of the query ([None] appears as the name if you have none saved) or Edit to enter the rest of the query. You can also click the parentheses icon () to add parentheses for nested queries or click the trash can icon () to remove parentheses.
- Step 3** Under Search String enter strings to query; under Operation, select the operation (AND, OR, NOT). For the final item in the list, select **None**.
- Step 4** Click the **Apply** button.
- Step 5** Click the **Submit** button to run the query.

**Note**

The free-form query cannot be saved as a rule.

To Run a Batch Query

- Step 1** Enter your data for either a simple or free-form query. If your query is expected to take a long time to run, instead of **Submit Inline**, you may given the option of having it run as a batch query.

Figure 20-4 Construct a Query to Run in Background (Batch Query)

Query type: *Event Types ranked by Sessions, 0h:10m*

Source IP	Destination IP	Service	Events	Device	Reported User	Keyword	Operation	Rule	Action
ANY	ANY	ANY	ANY	ANY	ANY	ANY	None	ANY	ANY

- Step 2** Click **Submit...** to make your selection.

Figure 20-5 Choosing the Query Submission Method

Choose Query Submission Method

This query will likely take a significant amount of time to complete.

To have the query run in the background, select "Submit Batch." The results will be sent to you via email (assuming a correct entry in your user profile), and will be saved for viewing later. If you desire, the query can be run again at a future time and the previously computed results will be reused.

To run the query immediately, select "Submit Inline." The results will be displayed in your browser as soon as the query completes; no results will be saved and no email will be sent.

To submit as a standard inline query, click **Submit Inline**. To submit your query as a batch query, click **Submit Batch**. Your query is submitted, and you are automatically taken to the **Batch Query** tab.

If your query is very large, you may only be give the options of **Save as Rule**, **Save as Report**, or **Submit Batch**.

Figure 20-6 Change Query Criteria

Query Event Data

Click the cells below to change query criteria:

Query type: *Event Types ranked by Sessions, 2dd:0hh:0mm:0ss*

Source IP	Destination IP	Service	Events	Device	Severity	Zone	Operation	Rule	Action	Reported User
[10.1.1.6] 10.1.1.6	my group [10.0.0.0 / 255.0.0.0] n-10.0.0.0/8	BackOrifice (src port: ANY, dst port: 31337, proto: TCP), BackOrifice (src port: ANY, dst port: 31338, proto: TCP)	ANY	ANY	ANY	ANY	None	ANY	ANY	ANY

Keywords: [None]

To submit your query as a batch query, click **Submit Batch**. Your query is submitted, and you are automatically taken to the **Batch Query** tab.

Figure 20-7 **Select Batch Query**

Page Refresh Rate

Never ▼

Batch Query Selection

Owner	Query	Status	Submitted	Time Range
Administrator, Administrator (pnadmin)	Query Type: Event Types ranked by Sessions Time: 4ww:2dd:0hh:10mm:0ss	Not Run	Never	May 5, 2004 11:52:25 AM PDT - Jun 4, 2004 12:02:25 PM PDT

View HTML ▼ View Results Resubmit Stop Delete

1 to 1 of 1 25 per page ▼ 143434

Step 3 To watch the status of the query in real-time, you can use the drop-down list to change the **Page Refresh Rate** from **Never** (the default) to 1 minute, 3 minutes, 5 minutes, 10 minutes, 15 minutes, or 30 minutes.

Step 4 To view the results of the batch query as it is running, click **View Results**. This can be done while the query is in progress.

If the email address in your user profile on the MARS is valid, the results of your batch query are emailed to you when the query has completed, and can also be viewed by clicking **QUERY / REPORTS > Batch Query > View Results**.



Note

When you click **View Results** while the query is in progress, the results compiled up to that moment are recomputed. This can make the display take longer to appear than after the results are compiled.

To Stop a Batch Query

Step 1 Click **QUERY/REPORTS**, then click the **Batch Query** tab.

Step 2 Click **Stop**. The **Status** of the query changes to **Finished**.

To Resubmit a Batch Query

You can resubmit a batch query if you want to restart it. A resubmitted batch query will use previously computed results, thus resulting in a faster query than one submitted for the first time.

Step 1 Click **QUERY/REPORTS**, then click the **Batch Query** tab.

Step 2 Click **Resubmit**. The **Status** of the query changes to **In Progress**.

To Delete a Batch Query

Step 1 Click **QUERY/REPORTS**, then click the **Batch Query** tab.

Step 2 Click **Delete**.

Step 3 In the confirmation window, click **Delete** to confirm.



Note

You can only see your own batch queries and their results. The batch queries of others and their results are not viewable by you, and your batch queries and their results are not viewable by others.

Selecting the Query Type

Figure 20-8 Clicking the Query Type or Edit link



You can select different query criteria by clicking the **Query Type** link or **Edit** button. This lets you determine a query's result format, rank, time, whether it only uses firing events, and the number of rows returned.

Figure 20-9 The Query Criteria: Result Page

The screenshot shows a configuration page for query criteria. It includes the following fields and options:

- Result Format:** A dropdown menu currently set to "Event Type Ranking".
- Order/Rank By:** A dropdown menu currently set to "Session Count".
- Filter by Time:**
 - ☐ Last: 0 Days 0 Hrs 10 Mins
 - ☒ Start: 2005 February 1 18 Hrs 00 Mins
 - End: 2005 February 29 18 Hrs 00 Mins
 - ☐ Real Time
- Use Only Firing Events:** A checkbox (unchecked) followed by a dropdown menu set to "Any Status".
- Maximum rank returned:** A text input field containing the value "100".

Result Format

- *Event Type Ranking*

Returns the most reported event types. Ranked by either: number of sessions containing at least one of the event type or by bytes transmitted in sessions that contain events that meet the query criteria.

- *Event Type Group Ranking*

Returns either pre-defined or user defined grouped event types. Ranked by either: number of sessions containing at least one event type contained in the group or by bytes transmitted in sessions that contain events that meet the query criteria.

- *Source IP Address Ranking*

Returns source IP addresses. Ranked by number of sessions with that source IP address or by bytes transmitted in sessions that contain events that meet the query criteria.

- *Network Ranking*

Returns top networks that exists in MARS. Ranked by either: number of sessions that contain events that meet the query criteria or by bytes transmitted in sessions that contain events that meet the query criteria. If a network is excluded, it is excluded from all results.

- *Network Group Ranking*

Returns top network groups that exists in MARS. Ranked by either: number of sessions that contain events that meet the query criteria or by bytes transmitted in sessions that contain events that meet the query criteria. If a network is excluded, it is excluded from all results.

- *Source Network Ranking*

Returns top source networks that exists in MARS. Ranked by either: number of sessions that contain events that meet the query criteria or by bytes transmitted in sessions that contain events that meet the query criteria. If a network is excluded, it is excluded from all results.

- *Source Network Group Ranking*

Returns top source network groups that exists in MARS. Ranked by either: number of sessions that contain events that meet the query criteria or by bytes transmitted in sessions that contain events that meet the query criteria. If a network is excluded, it is excluded from all results.

- *Destination Network Ranking*

Returns top destination networks that exists in MARS. Ranked by either: number of sessions that contain events that meet the query criteria or by bytes transmitted in sessions that contain events that meet the query criteria. If a network is excluded, it is excluded from all results.

- *Destination Network Group Ranking*

Returns top destination network groups that exists in MARS. Ranked by either: number of sessions that contain events that meet the query criteria or by bytes transmitted in sessions that contain events that meet the query criteria. If a network is excluded, it is excluded from all results.

- *Destination IP Address Ranking*

Returns destination IP addresses. Ranked by either: number of sessions with that destination IP address or by bytes transmitted in sessions that contain events that meet the query criteria.

- *Source Port Ranking*

Returns source ports. Ranked by either: number of sessions with that source port or by bytes transmitted in sessions that contain events that meet the query criteria.

- *Destination Port Ranking*

Returns destination ports. Ranked by either: number of sessions with that destination port or by bytes transmitted in sessions that contain events that meet the query criteria.

- *Protocol Ranking*

Returns most used protocols. Ranked by either: number of sessions with that protocol or by bytes transmitted in sessions that contain events that meet the query criteria.

- *Reporting Device Ranking*

Returns most active reporting devices. Ranked by either: number of sessions that contain events from the device or by bytes transmitted in sessions that contain events that meet the query criteria.

- *Reporting Device Type Ranking*

Returns most active reporting device types. Ranked by either: number of sessions that contain events from a device of that type or by bytes transmitted in sessions that contain events that meet the query criteria.

- *Reported User Ranking*

Returns information about users from reporting devices such as: Windows clients, Solaris clients, etc. Ranked by either: number of sessions that contain events from a reported user or by bytes transmitted in sessions that contain events that meet the query criteria.

- *Matched Rule Ranking*

Returns top firing rules. Ranked by number of incidents.

- *Matched Incident Ranking*

Returns incidents. Ranked by either: number of sessions that contain events that meet the criteria that contributed to the incident or by bytes transmitted real time in sessions that contain events that meet the query criteria.

- *All Matching Sessions*

Returns all sessions that contain events that meet the criteria. Sessions that contain a common set of event types are grouped together. They are also sub-grouped by session source IP address and session destination IP address. Sessions in the same sub-group are ordered by time. Real Time results are available for this Result Type.

- *All Matching Events*

Returns events. Ranked by time with the most current first. Real Time results are available for this Result Type.

- *All Matching Event Raw Messages*

Returns the raw messages associated with events. Ranked by time with the most current first. Real Time results are available for this Result Type.

- *NAT Connection Report*

Returns NAT connections. Ranked by time with the most current first.

- *MAC Address Report*

Returns MAC addresses. Ranked by time with the most current first.

- *Unknown Event Report*

Returns events that are not fully processed by the MARS. In some cases, event information such as the five tuple (source IP, source port, destination IP, destination port, and protocol) might not be present, hence can not be queried in real time.

Order/Rank By

This selection determines the ranking or order of the query's results. These selections are determined by the kind of Result Format that you use when you run the query.

- *Session Count*

The number of sessions that contain events that meet the criteria that contributed to the incident.

- *Bytes Transmitted*

The number of bytes transmitted in sessions that contain events that meet the query criteria.

- *Time*

Most current results appear first.

- *Incident Count*

Largest number of incidents appear first.

Filter By Time

- *Last*

The present time minus the number of days, hours, and minutes entered.

- *Start/End*


Absolute literal time ranges defined by the date to the minute.

- *Real Time*

Streams rolling real-time results from recent past to current time. Result Formats that work in real time are: [•All Matching Sessions, page 20-7](#), [•All Matching Events, page 20-7](#), and [•All Matching Event Raw Messages, page 20-7](#).

Real Time results appear in a normal browser window. Moving the scroll bar stops the “rolling” behavior. Clicking the Resume button on the bottom of the page allows the scrolling to resume.

Figure 20-10 Click the Resume Button to Start the Page Rolling



1	Top row visible	2	Bottom row visible
3	Total rows queried since start	4	Number of new queries pulled when this page last refreshed

Use Only Firing Events

Select this if you want only events that fired incidents to return information.

Maximum Number of Rows Returned

Select the number of rows that you want displayed.

Selecting Query Criteria

To Select a Criterion


Step 1 Select the criteria that you want to edit by clicking it.

Figure 20-11 *Clicking any to narrow your criteria*



Step 2 Move the items that you want to query from the right to the left of the filter by selecting the check box next to them, and clicking the Equal and Not Equal buttons.

Figure 20-12 Selecting Variables



Step 3 You can select a variety of different variables, events, devices, addresses from the filter page. The following number correspond with the numbers in the preceding graphic:

1. Check the boxes next to the items in the **Sources Selected** field to select them, and click the **Toggle Equal** button to change them between equal and not equal.
2. Click the **Select All** button to select all items in the **Sources Selected** field. (Note: if you have items highlighted in the Sources Selected field, clicking **Select All** will de-select them.)
3. Use the **Equal** and **Not Equal** buttons to bring highlighted items from the **Sources Available** field into the **Sources Selected** field.
4. Filter sources from this drop-down list.

5. Enter search text, and click **Search** to move items that match the search criteria from the **Sources Available** field to the **Sources Selected** field.
6. To add a new item to the sources, click the **Add** button. To edit or delete an existing source, click the **Edit** or **Delete** button. See [IP Management, page 23-3](#) for more information.
7. Click an item or items in the Sources Selected field, and use the **Remove** button.
8. To move IP values up into the Sources Selected field, click the **Equal**  (Up) icon, or the **Not Equal**  (Up) icon.
9. Check the radio button next to **IP** or **Range**, and enter an IP address or a range of IP addresses into their respective fields.
10. Select items in the Sources Selected field by clicking them. Enter a group name, and click the **Grouped As** button to group them.
11. Once you have chosen the query criteria that interests you, click **Apply** to return to the Query page. Repeat this selection process for other query data.

Step 4 Click the **Submit** button to run the query.

Query Criteria

The following list describes the selections in the Query Event Data table.

Source IP

- *Pre NAT source addresses*

Specifies that the constraints entered are the session endpoints.

- *Post NAT source addresses*

Specifies that the constraints entered are the source as appearing at the destination.

- *ANY*

No constraint is placed on the source IP addresses.

- *Variables*

Signify any one IP address, only useful for queries in tandem with the same variable.

- *IP addresses*

IP addresses present on devices in the system or user entered dotted quads.

- *IP ranges*

The range of addresses between two dotted quads.

- *Networks*

Topologically valid networks.

- *Devices*

The hosts and reporting devices present in the system.

Destination IP

- *Post NAT destination addresses*

Specifies that the constraints entered are the session endpoints.

- *Pre NAT destination addresses*

Specifies that the constraints entered are the destination as appearing at the source.

- *ANY*

No constraint is placed on the source IP addresses.

- *Variables*

Any one IP address, only useful for queries in tandem with the same variable.

- *IP addresses*

IP addresses present on devices in the system or user entered dotted quads.

- *IP ranges*

The range of addresses between two dotted quads.

- *Networks*

Topologically valid networks.

- *Devices*

The hosts and reporting devices present in the system.

Service

- *ANY*

No constraint is placed on the source or destination ports or protocol.

- *Service variables*

Any one set of destination port and protocol, only useful for queries in tandem with the same variable.

- *Defined services*

Services on the database.

Event Types

- *ANY*

No constraint on the event type.

- *Event types*

Events that have been merged into types.

- *Event type groups*

Groups of event types.

Device

- *Devices*

The reporting devices present in the system. This restricts the query to a subset of the devices that report to the MARS.

Severity/Zone

- *ANY*

No constraint on the event type severity.

- *Green*

Low-severity events

- *Yellow*

Medium-severity events

- *Red*

High-severity events

- *Zone*

Events reported by devices in the indicated zone.

Operation

- *None*

Defines a single-line query.

- *AND*

Boolean “and” that defines a two or more line query.

- *OR*

Boolean “or” that defines a two or more line query.

- *FOLLOWED-BY*

Time conditional query (e.g.: Y must happen after X) that defines a two or more line query.

Rule

- *Empty field – Rules Chosen field*

When this field is empty, it acts like an ANY selection. No constraint is placed on the sub-set of events.

- *Rule*

Restricts the query to the sub-set of events that contributed to the incidents of the specified rules firing.

Action

- *Empty field – Empty Actions Chosen field*

When this field is empty, it acts like an ANY selection. No constraint is placed on the sub-set of events.

- *Actions*

Restricts the query to the sub-set of events that contributed to the incidents of rules that have the specified notifications as part of their actions. (See page ??? rules for more information?)

Saving the Query

You can save query criteria to re-use as reports or rules.

To save a query as a report

This takes the query that you are using and creates a report. For more information on creating reports, see [Reports, page 20-22](#).

To save a query as a rule

This takes the query to the rules page, populating the rules with the selected query criteria. Likely, you must identify additional criteria to complete the rule. For more information on creating rules, see [Rules, page 21-1](#).

Viewing Events in Real-time

The Real-time Event viewer is a query option that permits you to view real-time events as follows:

- View raw events as they stream to MARS before they are sessionized, with a maximum 5-second delay
- View a sessionized event stream—more delay is possible when there are many events in a session

The real-time events display as a continuously scrolling screen. You can configure query criteria to filter what is displayed. When viewing raw events, sessionization is not impeded, all the parsed raw events are sessionized per normal MARS operation. MARS

Restrictions for Real-time Event Viewer

Real-time event queries should be made *only* from a browser instance that was used to login to MARS. The real-time query will not have reliable results if it is executed from a browser instance spawned from the original login instance (for example, a new browser window launched with **Ctrl+N**, **File>New>New Window**, or **right-click** {link on MARS interface}>**Open in New Window**).

Multiple real-time queries can operate in multiple browser instances at the same time, but you *must* login to MARS with each browser instance.

Procedure for Invoking the Real-Time Event Viewer

To invoke the real-time event viewer, complete the following steps:

-
- Step 1** Navigate to the **Query** home page as shown in [Figure 20-13](#).

Figure 20-13 Query Home Page

Step 2 Click **Edit**. The Query edit dialog appears, as shown in Figure 20-14.

Figure 20-14 Configuring Real-Time Event Viewer Query

Step 3 Do the following substeps:

- From the **Result Format** dropdown list, select **All Matching Events** or **All Matching Event Raw Messages**.
- Click the **Real Time** radio button, and select **Raw events** or **Sessionized Events** from the dropdown list.

All Matching Events with **Raw events** displays Event ID, Event Type, Source IP/Port, Destination IP/Port, Protocol Time, and Reporting Device fields.

All Matching Events Raw Messages with **Raw events** displays Event ID, Event Type, Time, Reporting Device, and Raw Message fields.

Either Result Format type with **Sessionized Events** displays Event/Session/Incident ID, Event Type, Source IP/Port, Destination IP/Port, Protocol, Time, Reporting Device, Path/Mitigation, and Tune fields.

c. Click **Apply**.

The Query Event Data screen appears with the **Save as Report** and **Save as Rule** buttons gray and inactive, as shown in [Figure 20-15](#).

Figure 20-15 Real-Time Event Query to Submit

Load Report as On-Demand Query with Filter

Select Group...

Select Report...

Incident ID:

Session ID:

Query Event Data

Click the cells below to change query criteria:

Query type: Events ranked by Time, Real Time(raw events)

Source IP	Destination IP	Service	Events	Device	Reported User	Keyword	Operation	Rule	Action
ANY	ANY	ANY	ANY	ANY	ANY	ANY	None	ANY	ANY
			Apply						

Step 4 Modify the parameters of the Query Event Data filter as you require and click **Submit**.



Note The Operation, Rule, and Action parameters of the Query Event Data filter do not function for the real-time event viewer.

Real-time results begin to scroll up from the bottom of the page within 5 seconds, as shown in [Figure 20-16](#). Real-time raw events are shown in this example.

Figure 20-16 Viewing Events in Real-Time

Event ID	Event Type	Time	Reporting Device	Raw Message
87589898	Unknown Device Event Type	Mar 15, 2006 5:57:10 PM PST	earth2	10.4.1.1 Sat Nov 2 16:02:12 2002 <134>%PIX-6-302001: Built inbound TCP connection 10056383 for faddr 219.51.92.21/64776 gaddr 67.118.229.242/80 laddr 10.1.1.30/80
87589899	Unknown Device Event Type	Mar 15, 2006 5:57:10 PM PST	earth2	10.4.1.1 Sat Nov 2 16:02:12 2002 <134>%PIX-6-302001: Built inbound TCP connection 7756979 for faddr 249.205.234.83/59027 gaddr 67.118.229.242/80 laddr 10.1.1.30/80
87589900	Unknown Device Event Type	Mar 15, 2006 5:57:10 PM PST	earth2	10.4.1.1 Sat Nov 2 16:02:12 2002 <134>%PIX-6-302001: Built inbound TCP connection 15424022 for faddr 46.144.232.118/31134 gaddr 67.118.229.242/80 laddr 10.1.1.30/80
87589901	Unknown Device Event Type	Mar 15, 2006 5:57:10 PM PST	earth2	10.4.1.1 Sat Nov 2 16:02:12 2002 <134>%PIX-6-302001: Built inbound TCP connection 14518954 for faddr 27.64.245.11/35092 gaddr 67.118.229.242/80 laddr 10.1.1.30/80
87589902	Unknown Device Event Type	Mar 15, 2006 5:57:10 PM PST	earth2	10.4.1.1 Sat Nov 2 16:02:12 2002 <134>%PIX-6-302001: Built inbound TCP connection 15166103 for faddr 195.167.19.52/31447 gaddr 67.118.229.242/80 laddr 10.1.1.30/80
87589903	Unknown Device Event Type	Mar 15, 2006 5:57:10 PM PST	earth2	10.4.1.1 Sat Nov 2 16:02:12 2002 <134>%PIX-6-302001: Built inbound TCP connection 4438904 for faddr 95.55.162.89/34335 gaddr 67.118.229.242/80 laddr 10.1.1.30/80
87589904	Unknown Device Event Type	Mar 15, 2006 5:57:10 PM PST	earth2	10.4.1.1 Sat Nov 2 16:02:12 2002 <134>%PIX-6-302001: Built inbound TCP connection 3063834 for faddr 108.48.250.124/13434 gaddr 67.118.229.242/80 laddr 10.1.1.30/80
87589905	Unknown Device Event Type	Mar 15, 2006 5:57:10 PM PST	earth2	10.4.1.1 Sat Nov 2 16:02:12 2002 <134>%PIX-6-302001: Built inbound TCP connection 15782919 for faddr 128.81.130.55/17423 gaddr 67.118.229.242/80 laddr 10.1.1.30/80
87589906	Unknown Device Event Type	Mar 15, 2006 5:57:10 PM PST	earth2	10.4.1.1 Sat Nov 2 16:02:12 2002 <134>%PIX-6-302001: Built inbound TCP connection 15580904 for faddr 58.74.186.118/9997 gaddr 67.118.229.242/80 laddr 10.1.1.30/80
87589907	Unknown Device Event Type	Mar 15, 2006 5:57:10 PM PST	earth2	10.4.1.1 Sat Nov 2 16:02:12 2002 <134>%PIX-6-302001: Built inbound TCP connection 7688160 for faddr 44.209.154.112/31382 gaddr 67.118.229.242/80 laddr 10.1.1.30/80
87589908	Unknown Device Event Type	Mar 15, 2006 5:57:10 PM PST	earth2	10.4.1.1 Sat Nov 2 16:02:12 2002 <134>%PIX-6-302001: Built inbound TCP connection 744799 for faddr 201.87.206.66/24688 gaddr 67.118.229.242/80 laddr 10.1.1.30/80
87589909	Unknown Device Event Type	Mar 15, 2006 5:57:10 PM PST	earth2	10.4.1.1 Sat Nov 2 16:02:12 2002 <134>%PIX-6-302001: Built inbound TCP connection 2107207 for faddr 208.8.105.2/26237 gaddr 67.118.229.242/80 laddr 10.1.1.30/80

Scroll speed:


The Real-time event viewer display is governed by the following controls:

- **Scroll Speed**—Select one of four scrolling rates.
- **Pause button**—Suspends the scrolling display, there is no timeout for pause.

- **Restart** button—Restarts the display from the current time. This button appears when you pause the scrolling display.
- **Resume** button—Restarts the display from the time when paused. This button appears when you pause the scrolling display.
- **Clear**—Terminates the real-time query.

Step 5 Click the active links within a real-time event record to view the related pop-up windows. For example, the Reporting Device Information pop-up window is shown in [Figure 20-17](#).

Figure 20-17 Reporting Device Information Pop-up Window



Should errors occur during the display of events, a message box appears, as shown in [Figure 20-18](#).

Figure 20-18 Real-time Event Viewer Error Message



Click **OK** to clear the message box, and restart the Real-time event viewer by clicking **Submit**.



Tip

To view the most recent real-time events, you can click **Submit** at any time, or **Pause** and **Restart** to reinitialize the Real-Time Event Viewer. The most recent events are always at the bottom of the output queue, and their freshness when you view them is limited by the number of events in the queue and the scroll speed of the display.

This ends the [Procedure for Invoking the Real-Time Event Viewer](#).

Perform a Long-Duration Query Using a Report

This section explains how to create and view a long-duration query on the MARS. There are two ways to perform a long-duration query on the MARS:

1. Modifying an existing report.

Advantages:

- The report is compiled relatively quickly.
- You can compile data gathered over a longer time period

Disadvantage.

This type of query can only be used without any changes to query criteria other than time range, and can only be used with the following reports:

- Activity: All - Top Destination Ports
- Activity: All - Top Destinations
- Activity: All - Top Event Types
- Activity: All - Top Reporting Devices
- Activity: All - Top Sources
- Activity: Attacks Seen - Top Reporting Devices
- Activity: Denies - Top Destination Ports
- Activity: P2P Filesharing/Chat - Top Event Types
- Activity: Scans - Top Destination Ports
- Activity: Scans - Top Destinations
- Activity: Unknown Events - All Events
- Activity: Web Usage - Top Destinations by Sessions
- Activity: Web Usage - Top Sources
- Attacks: All - Top Rules Fired
- Attacks: All - Top Sources

2. Performing a batch query.

Advantages:

- You can modify any of the query criteria.
- Best suited for data that spans a short time period.

Disadvantages

- This type of query can be slow and may take a substantial amount of time to complete.
- Only Admin users can perform a batch query.

If you want to observe activity on your MARS over a long period, you can change the duration of time over an existing report that runs on a regular basis, such as hourly or daily, whether they are shipped with the MARS or created by you.

**Note**

Trying to run a long-duration query using a report that only runs “on demand” has the same effect as running a query; it can take just as long because it has to compile data, whereas data from the regularly-run reports has been precompiled on an ongoing basis.

To query using a report, follow these steps:

Step 1 In the **QUERY / REPORTS** tab, click the **Reports** tab to obtain the Main Report window.

Figure 20-19 Main Report Window

Report Selection

	Name	Schedule	Format	Recipients	Query	Description	Status	Submitted	Time Range
☐	Activity: All - NAT Connections	Run on demand only	Normal	None	Query Type: NAT connections ranked by Time Time: May 1, 2004 8:21:50 PM PDT - Jun 6, 2004 8:31:50 PM PDT	This report lists Network Address Translations performed on non-denied sessions as reported to MARS.	Finished: Jun 16, 2004 4:40:36 AM PDT	Jun 15, 2004 8:32:09 PM PDT	May 1, 2004 8:21:50 PM PDT - Jun 6, 2004 8:31:50 PM PDT
☐	Activity: All - Top Destination Ports	Run on demand only	Trend	None	Query Type: Destination Ports ranked by Sessions Time: 1hh:0mm:0ss	This report ranks the UDP and TCP destination ports of all events seen by MARS over the past hour. This report is used by pages in the Summary tab.	Finished: Jun 10, 2004 4:16:58 PM PDT 4:17:02 PM PDT	Jun 10, 2004 4:16:58 PM PDT	Jun 10, 2004 3:16:58 PM PDT - Jun 10, 2004 4:16:58 PM PDT
☒	Activity: All - Top Destinations	Every hour	Normal	None	Query Type: Destination IPs ranked by Sessions Time: 28ww:4dd:0hh:10mm:0ss	This report ranks the session destinations of all events seen by MARS over the past hour. This report is used by pages in the Summary tab.	Finished: Jun 17, 2004 2:15:52 PM PDT 2:15:52 PM PDT	Jun 17, 2004 2:15:52 PM PDT	Nov 30, 2003 1:05:52 PM PST - Jun 17, 2004 2:15:52 PM PDT

Step 2 Navigate to and then click the radio button next to the regularly-scheduled report you want to modify (in this example, we use **Activity: All - Top Destinations**). Click the **Query** column to edit the report. The Build Report window appears.

Figure 20-20 Build Report window

Build Report

Click the cells below to define the report:

Name	Schedule	Format	Recipients	Query	Description	Status	Submitted	Time Range
Activity: All - Top Destinations	Every hour	Normal	None	Query Type: Destination IPs ranked by Sessions Time: 28ww:4dd:0hh:10mm:0ss	This report ranks the session destinations of all events seen by MARS over the past hour. This report is used by pages in the Summary tab.	Finished: Jun 16, 2004 7:15:42 PM PDT	Jun 16, 2004 7:15:42 PM PDT	Nov 29, 2003 6:05:42 PM PST - Jun 16, 2004 7:15:42 PM PDT

Submit Previous Next

Time Range:

☒ Last: 200 Days 0 Hrs 10 Mins

☐ Start: 2004 June 16 19 Hrs 42 Mins
End: 2004 June 16 19 Hrs 52 Mins

Submit Previous Next

Step 3 In the lower portion of the Build Report window, change the **Time Range** the report (**Activity: All - Top Destinations**) covers to the duration you want it to cover.

Step 4 Click the **Submit** button to run the report and return to the Main Report window.

View a Query Result in the Report Tab

To view a query in the Report tab, follow these steps:

Figure 20-21 Main Report window (bottom)

☒	Activity: All - Top Destinations	Every hour	Normal	None	Query Type: Destination IPs ranked by Sessions Time: 28ww:4dd:0hh:10mm:0ss	This report ranks the session destinations of all events seen by MARS over the past hour. This report is used by pages in the Summary tab.	Finished: Jun 17, 2004 2:15:52 PM PDT	Jun 17, 2004 2:15:52 PM PDT	Nov 30, 2003 1:05:52 PM PST - Jun 17, 2004 2:15:52 PM PDT
☐	Activity: All Events and Netflow - Top Destination Ports	Run on demand only	Trend	None	Query Type: Destination Ports ranked by Sessions Time: 1hh:0mm:0ss	This report ranks the UDP and TCP destination ports of all events (including Netflow events) seen by MARS over the past hour. This report is used by pages in the Summary tab.	Finished: Jun 8, 2004 9:29:03 PM PDT	Jun 8, 2004 8:28:51 PM PDT - Jun 8, 2004 9:28:51 PM PDT	
☐	Activity: All Sessions - Top Destination Ports by Bytes	Run on demand only	Normal	None	Event type: Info/AllSession, Query Type: Destination Ports ranked by Bytes Transmitted Time: 0hh:10mm:0ss	This report ranks all destination ports by bytes transferred.	Not Run	Jun 8, 2004 9:29:20 PM PDT	Jun 8, 2004 9:19:20 PM PDT - Jun 8, 2004 9:29:20 PM PDT
☐	Activity: All Sessions - Top Destinations by only Bytes	Run on demand only	Normal	None	Event type: Info/AllSession, Query Type: Destination IPs ranked by Bytes Transmitted Time: 0hh:10mm:0ss	This report ranks all destinations by bytes transferred.	Not Run	Jun 8, 2004 9:29:57 PM PDT	Jun 8, 2004 9:19:57 PM PDT - Jun 8, 2004 9:29:57 PM PDT

View HTML
View Report
Resubmit
Add
Edit
Delete

143799

Step 1 At the bottom of the Main Report window, click the radio button next to the report (**Activity: All - Top Destinations**).

Step 2 From the drop-down list on the bottom of the Reports page, select either:

- **View HTML:** to view the report as an HTML file.
- **View CSV:** to view the report as a CSV (comma-separated values) file.

Step 3 Click the **View Report** button.



Note

The **Status** column of the report lets you know whether the report has finished before viewing. You can view a partially-completed report, but it might not contain all the data you want to examine. You can also refresh the screen to update the **Status** column.

Perform a Batch Query

This type of long-duration query can take a long time to perform and is more suitable for a shorter duration of time.



To perform a batch query, follow these steps:

Figure 20-22 *Query window*

Click the cells below to change query criteria:

143796

20-20

Figure 20-23 **Query Event Data window**

Query Event Data

Click the cells below to change query criteria:

Query type: **Sessions ranked by Time, 0hh:10mm:0ss** [Edit](#) [Clear](#)

Source IP	Destination IP	Service	Events	Device	Severity	Zone	Operation	Rule	Action	Reported User
H-10.1.252.250	H-65.54.153.118	ANY	ANY	ANY	ANY	ANY	None	ANY	ANY	ANY

Keywords: [**None**] [Edit](#)

[Apply](#)

Result Format: All Matching Sessions

Order/Rank By: Time

Filter by Time:

☒ Last: Days Hrs Mins

☐ Start: 2004 June 22 18 Hrs 38 Mins

End: 2004 June 22 18 Hrs 48 Mins

☐ Real Time

Use Only Firing Events: ☐

Maximum rank returned:

[Apply](#)

Step 3 In the Query Event Data window, you can change the query criteria. (For more information on query criteria, see [Query Criteria, page 20-10](#)). By clicking on various parameters you can change the nature of the query. In this case we are specifying a Source IP address of **10.1.1.6**, a Destination IP address range previously saved as **mygroup**, and setting the duration of the query to the past **2** days. Click either **Apply** button to apply your changes to the query. The Query Save/Submit window appears.

Figure 20-24 **Query Save/Submit window**

Query Event Data

Click the cells below to change query criteria:

Query type: *Event Types ranked by Sessions, 2dd:0hh:0mm:0ss*

Source IP	Destination IP	Service	Events	Device	Severity	Zone	Operation	Rule	Action	Reported User
[10.1.1.6] 10.1.1.6	my group [10.0.0.0 / 255.0.0.0] n-10.0.0.0/8	BackOrifice (src port: ANY, dst port: 31337, proto: TCP), BackOrifice (src port: ANY, dst port: 31338, proto: TCP)	ANY	ANY	ANY	ANY	None	ANY	ANY	ANY

ANY

Keywords: [None]

- Step 4** The Query Save/Submit window asks you to choose from the options of **Save as Rule**, **Save as Report**, or **Submit Batch**. To submit your query as a batch query, click **Submit Batch**. Your query is submitted, and you are automatically taken to the Batch Query tab.

Figure 20-25 Batch Query tab

Page Refresh Rate

1 minute ▼

Batch Query Selection

Owner	Query	Status	Submitted	Time Range
☒ Administrator, Administrator (pnadmin)	Query Type: Event Types ranked by Sessions Time: 0hh:10mm:0ss	Finished: Jun 21, 2004 8:07:08 PM PDT	Jun 21, 2004 8:07:02 PM PDT	Jun 21, 2004 7:57:02 PM PDT - Jun 21, 2004 8:07:02 PM PDT
☐ Administrator, Administrator (pnadmin)	Query Type: Event Types ranked by Sessions Time: 4ww:2dd:0hh:10mm:0ss	Not Run	Never	May 5, 2004 11:52:25 AM PDT - Jun 4, 2004 12:02:25 PM PDT
☐ Administrator, Administrator (pnadmin)	Query Type: Event Types ranked by Sessions Time: 2ww:0dd:0hh:0mm:0ss	Finished: Jun 13, 2004 2:17:43 PM PDT	Jun 13, 2004 12:58:32 PM PDT	May 30, 2004 12:58:32 PM PDT - Jun 13, 2004 12:58:32 PM PDT
☐ Administrator, Administrator (pnadmin)	Event type: != Built/teardown/permitted IP connection, Query Type: Event Types ranked by Sessions Time: 4ww:2dd:0hh:10mm:0ss	Stopped: 16%	Jun 13, 2004 12:42:35 PM PDT	May 14, 2004 12:32:35 PM PDT - Jun 13, 2004 12:42:35 PM PDT
☐ Administrator, Administrator (pnadmin)	Query Type: Event Types ranked by Sessions Time: 1ww:6dd:0hh:10mm:0ss	Finished: Jun 13, 2004 1:37:15 PM PDT	Jun 13, 2004 12:40:35 PM PDT	May 31, 2004 12:30:35 PM PDT - Jun 13, 2004 12:40:35 PM PDT

View HTML View HTML View CSV

View Results Resubmit Stop Delete

143785

- Step 5** To watch the status of the query in real-time, you can use the Batch Query tab drop-down list to change the **Page Refresh Rate** from **Never** (the default) to 1 minute, 3 minutes, 5 minutes, 10 minutes, 15 minutes, or 30 minutes.

- Step 6** To view the results of the batch query as it is running, click the radio button next to your query (here it's highlighted in green) and click **View Results**. This can be done while the query is in progress.

If the email address in your user profile on the MARS is valid, the results of your batch query are emailed to you when the query has completed. You can also view the results of your batch query by clicking **QUERY / REPORTS > Batch Query > View Results**.



Note

When you click **View Results** while the query is in progress, the results compiled up to that moment are recomputed. This can make the display take longer to appear than after the results are compiled.

Reports

Using the Reports page, you can build repeatable queries, edit and delete current reports, run reports, and view reports in either HTML or CSV (comma separated value) formats.

Predefined System Reports are treated as global reports. Global Controller receives report data once its connected to the Local Controller. Previous report results (prior to managing the Local Controller) will not be pushed up to Global Controller. Thus viewing of reports will not include the information before the Local Controller becomes active.

When you view a report, you are viewing the last instance that ran. If you want to view an up-to-the-minute report, resubmit the report before viewing it.

Report results are purged from the database after a purge interval, as tabulated in [Table 20-1](#).

Table 20-1 Maximum Database Retention Limits for Report Results

Cisco Security MARS Model	Maximum Number of Stored Reports ¹	Database Purge Interval ²
CS-MARS-20-K9	1,000 ranking reports 5,000 event/session reports	3 months
CS-MARS-50-K9	1,000 ranking reports 5,000 event/session reports	3 months
CS-MARS-100-K9	1,000 ranking reports 5,000 event/session reports	6 months
CS-MARS-100E-K9	1,000 ranking reports 5,000 event/session reports	6 months
CS-MARS-200-K9	1,000 ranking reports 5,000 event/session reports	6 months
CS-MARS-GC-K9	1,000 ranking reports 5,000 event/session reports	12 months
CS-MARS-GCM-K9	1,000 ranking reports 5,000 event/session reports	12 months

1. Table values are for Cisco Security MARS Release 4.1.5. In Release 4.1.4 and prior, the maximum number of ranking reports is 100, maximum number of event/session reports is 1,000.
2. As of Cisco Security MARS Release 4.1.5. In Release 4.1.3, and 4.1.4, report results are retained for one year in the MARS database before they are automatically purged. In Releases prior to Release 4.1.3, report results are retained indefinitely. The purge interval cannot be changed.

Report Type Views: Total vs. Peak vs. Recent

Where alerts provide up-to-the-minute views of high-priority incidents, reports aggregate sessions into different views. Reports correlate based on the three data points:

- Period of time
- Query criteria
- View type

The *period of time* defines boundaries around the analyzed session data based on when it was recorded. *Query criteria* restrict the set of sessions that will be aggregated to that which matches your criteria. Criteria can include source address, destination address, network service, event, reported user, and reporting device. The *view type* defines how to aggregate the matched data into a meaningful report view—one that matches the type of study in which you are interested.



Note

In each view type, you can refine the report criteria to filter out expected activity—the data you know about. You can filter this activity by refining the query criteria. These criteria should be tuned to a specific network. Reports can be valuable in detecting behaviors beyond the normal traffic flows of your network. You can determine the expected activities using reports that are not filtered and vetting those results against normal network use.

MARS provides three view types, each of which restricts the matched sessions to a user-defined limit of *N*. The following view types exist:

- **Total View.** For each result type matching the query criteria, this view counts the occurrences of that result type that transpire during the specified time period. It presents the total count of the top *N* matched result types, ranked by number of sessions, as determined by which ones occurred most frequently over the period of time. You can use these reports to determine your network's condition relative to the studied sessions. For example, you can use this view to identify attacks that launched at frequent intervals. This view does not present spikes in network activity; it simply presents the top occurring result types.
- **Peak View.** Within MARS, all report result data is stored in 10-minute time slices. The Peak View studies each of the 10-minute time slices within the specified time period to which one contained the highest number of matched sessions for a specific result type. It also determines an additional nine peaks within the time period, where each peak identifies a unique result type relative to the other peaks.

Each peak value is charted relative to the other nine peaks. For each time slice containing a peak value, the Peak View lists the top *N* matched result types that occurred. It is possible to have multiple peaks within the same time slice, as it is the result type, not the time slice, that must be unique across peaks.



Note

To be detected within this view, the result type must peak above normal traffic. Therefore, you must tune the query data to filter out expected traffic.

Unlike the Total View, the Peak View does not focus on the overall top occurring results, instead it identifies a high volume of traffic over a short time period. Its purpose is to detect temporary bursts of traffic on your network that overshadow normal traffic usage. These bursts identify possible issues, such as worm outbreaks.

- **Recent View.** This view is similar to Total View; however, it identifies the top *N* result types that occurred within the past hour. It then plots all occurrences of those result types over the selected time period.
- **CSV.** Generates the Total View but presents the report in the CSV format for processing by another tool or script. This option is intended for use with e-mail notifications where post-processing is required.

Creating a Report

You can create a report through the **Query** page, or you can create a report from scratch on the **Reports** page. These instructions detail creating a report from the **Reports** page, but are applicable to editing reports and to creating reports from the **Query** page.

To Create a New Report

- Step 1** On the Reports page, click the **Add** button.
- Step 2** In the **Report Name** and **Report Description** fields, enter a report name and description. Click the **Next** button.
- Step 3** Select the schedule parameters for the report.
- Step 4** Select a View Type for the report. You can receive these reports in your email or view them in the UI. Your choices are: **Total View**, **Peak View**, **Recent View**, and **CSV** (see [Report Type Views: Total vs. Peak vs. Recent](#), page 20-23). Click the **Next** button.
- Step 5** Select users in the Recipients Available field by expanding the user groups, clicking users or user groups, and clicking the **Add** button. See [User Management](#), page 23-8 for more information.

- Step 6** Repeat [Step 5](#) for other users. Click the **Next** button.
- Step 7** Build or modify the query. To edit the query time range, either click the Report type link or click the **Edit** button. See [Result Format, page 20-5](#) for information on query parameters; see [Query Criteria, page 20-10](#) for more information on building queries. Click **Apply** to save your changes; click **Next** when the query is complete.
- Step 8** Click **Submit** to save your report.
-

Working With Existing Reports

To View a Report

- Step 1** Click the radio button next to the report.
- Step 2** From the drop-down list on the bottom of the page, select either:
- **View HTML**: to view the report as an HTML file.
 - **View CSV**: to view the report as a CSV file.
- Step 3** Click the **View Report** button.

**Note**

If you chose to view the report as a CSV file, you need to save the file to your computer and open the CSV file in a third-party application.

To Run a Report

- Step 1** Click the radio button next to the report.
- Step 2** Click the **Run Now** button.

**Note**

Due to caching issues, reports with a time range of less than one hour are not recommended.

See [Table 20-1, “Maximum Database Retention Limits for Report Results”](#) for information on how long report results are retained in the database per MARS model number.

To Delete a Report


- Step 1** Click the radio button next to the report.
- Step 2** Click the **Delete** button to delete the report.
- Step 3** On the Delete Confirmation page, click **Delete**.
-

To Edit a Report

You can not edit system generated reports. Editing report criteria is meant for minor tweaking to previously generated report.

- Step 1** Click the radio button next to the report.
- Step 2** Click the **Edit** button to edit the report.
- Step 3** Navigate using the **Previous** and **Next** buttons, or clicking on the report criteria.

Figure 20-26 Navigating to the Recipients column by clicking its criteria



Name	Schedule	Format	Recipients
Test	Run on demand only	Normal	None

- Step 4** Edit the report, and click the **Apply** button to apply changes to the report.
- Step 5** Click the **Submit** button to finalize the report.



Note

Changing the report's query criteria will not re-generate a new result. New edited criteria is based on the previously generated report. In some situation such as filtering out specific IP source, user should create a new report.



Note

Email notification of a global generated report will be sent from the Global Controller and not the Local Controller.



Rules

This chapter discusses MARS Inspection and Drop rules in the following sections:

- [Rules Overview, page 21-1](#)
- [Constructing a Rule, page 21-5](#)
- [Working with System and User Inspection Rules, page 21-17](#)
- [Working with Drop Rules, page 21-21](#)
- [Setting Alerts, page 21-23](#)
- [Rule and Report Groups, page 21-24](#)

Rules Overview

An inspection rule is a real-time filter that detects interesting patterns of network activity. These patterns can signify attacks or false positives, and they inform you of network configuration errors and other anomalous network behavior. An attack might be straightforward, or it could be a probe, an attack, and then a follow-up to the attack. Whatever the method of attack, attacks share common traits, and you can use rules to define these traits to identify and mitigate attacks.

Rules create incidents. Rules connect the information you receive from your networks' reporting devices, linking them together to form a chain of events that describes an unfolding intrusion. They classify incoming events as firing events by matching them against the rule criteria. They also determine when a false positive is either dropped completely or kept as information in the database.

A rule is either active or inactive. Active means the rule is operating and is being applied to incoming events. Inactive indicates that the rule is inoperative and not consuming CS-MARS resources.



Note

A rule cannot be deleted, it can be made active or inactive.

[Figure 21-1](#) shows a portion of the Inspection Rules page of the Rules tab.

Figure 21-1 Top Portion of Inspections Rules Page

Inspection Rules:

Group: All View: Active

Edit Change Status Duplicate Add

Offset	Open	Source IP	Destination IP	Service Name	Event	Device	Reported User	Keyword	Severity	Count	Close	Operation
1	(	ANY	SAME, \$TARGET01, ANY	ANY	Penetrate/Backdoor/Rookit/Connect, Penetrate/Backdoor/Trojan/Connect, Penetrate/Backdoor/Trojan/SYN, Penetrate/Backdoor/CommandShell, Penetrate/Backdoor/RemoteControlApp/Connect	ANY	None	ANY	ANY	1	)	OR
2		SAME, \$TARGET01, ANY	ANY	ANY	Penetrate/Backdoor/RemoteControlApp/Response, Penetrate/Backdoor/Trojan/Response, Penetrate/Backdoor/Trojan/SYN-ACK	ANY	None	ANY	ANY	1	)	FOLLOWED-BY
3	((	SAME, \$TARGET01, ANY	DISTINCT, ANY	SAME_ANY_DEST_PORT	AttacksProtected, FirewallPolicyViolation/ACL, FirewallPolicyViolation/NAT	ANY	None	ANY	ANY	25	)	OR
4		SAME, \$TARGET01, ANY	ANY	ANY	DoS/Network/TCP, DoS/Network/UDP, DoS/Network/ICMP, DoS/Network/Misc, DoS/Distributed, Probe/HostInfo/All, Propagate/CopyFiles, Propagate/Worm, Penetrate/Backdoor/CovertChannel	ANY	None	ANY	ANY	1	)	OR
5		ANY	SAME, \$TARGET01, ANY	ANY	Persist/All, Penetrate/Backdoor/CovertChannel	ANY	None	ANY	ANY	1	)	

Prioritizing and Identifying

Your first order of business is to prioritize your network's assets; in other words, figure out what is going to cost you the most money if it goes down. Next, identify your networks' most exploitable weaknesses. Choose which ones you are willing and able to close, and rank the remaining weaknesses by risk and exploitability.

Use this ranked list to guide your time and energy expenditures when customizing the CS-MARS rule set.

Think Like a Black Hat

Ignore for a moment the benign users who do legitimate business on your networks.

Get inside the mind of the black hat that wants to take your network down. The person who should concern you is the one with a plan.

Good plans have a sequence of steps, contingencies, and metrics to determine success or failure. The more fully you can anticipate these plans, the fewer attacks will be able to execute unhindered and unobserved. The black hat is looking for wide-open doors and easy access. Failing that, the black hat is going to look for specific and obvious exploitable weaknesses.

Planning an Attack

Start to detail your plan. You want to penetrate a network. You'd like to avoid detection and identification if possible. You want root access on a host.

How do you get root access? You do not have a preexisting account, and physical access isn't feasible. The first few options that come to mind are password guessing, password brute force, or exploiting a known weakness on the host.

You decide to exploit services running on the host, so you need to find out what it is running. To do this, you have a number of techniques: port scans, OS fingerprinting, banner probing, etc.

Once you've identified a vulnerable service or software, you can attack it with a catalogue of exploit software. Depending on what you find and your available exploits, there are a number of different effects, usually allowing you to execute arbitrary code.

You now own the host. What happens next is up to you. You have many options: you can install a root kit, you can crash the machine, etc. You have full access—you can do just about anything on to/from that host.

Back to Being the Admin

You must now express the plan in terms of information that is reported to you. This attack plan contains an attack with a follow up of some kind. You might write your plan like:

- probe
- attacker to target, buffer overflow
- attacker to target, root login (compromised host)

At this point, the black hat has compromised the host. What happens next is up to the attacker. This makes the next few steps especially hard to predict. They want to be able to manipulate the world, they want to make change. Your newly compromised host is the instrument for change. You can specify additional potential steps in the plan that make it even more urgent to take care of the situation immediately. Such as:

- target to FTP server, code download
- target to secondary target, buffer overflow

The attacker is now using your compromised host as a launching point for further attacks.

One you've mapped out the anticipated attack to watch for, you can define a monitoring plan. The following task flow outlines the tasks involved in implementing a monitoring plan:

-
Step 1	Ensure your reporting devices are providing all the data you need. This step involves ensuring that each device is generating logs about the events that you expect to occur as the result of the probes and attacks. Depending on the device type, this can involve several substeps, such as specify a logging level, enable logging for the specific event, and ensuring that the reporting device publishes events to the Local Controller appliance. It can also involve enabling administrative access to the reporting device from the Local Controller appliance.
Step 2	Configure CS-MARS to pull events from the reporting devices on your network. This step involves adding each reporting device to Local Controller. If the reporting device type is not directly supported, you must define a custom device type for the reporting device. To add a supported reporting device, see Adding Reporting and Mitigation Devices, page 2-16 . To define a custom device type, see Adding User Defined Log Parser Templates, page 15-1 and To add a custom Device/Application type, page 15-1 .
Step 3	Ensure that the event types that you need to study are accepted and processed by Local Controller. If they are not, you must define a custom log parser template for each event and a custom device template to which the custom log parser templates are associated. For device types supported by CS-MARS, this should not be necessary. To define a new parser template, see Adding User Defined Log Parser Templates, page 15-1 and To add Parser Templates for a Device/Application, page 15-3 .

**Note**

You cannot define a custom log parser template for a reporting device that is supported out of the box. In this case, to define log parser for an unsupported event type, you must still define a custom device type before you can define the log parser.

Step 4

Check to see if a system rule will capture the information that you want, otherwise write your own user inspection rule. Define user inspection rules that monitor for the event types and correlate those events into a structure that will help you identify the incident. You can also specify who should be notified and how if the rule fires.

Types of Rules

**Note**

A rule cannot be deleted, it can be made active or inactive.

Inspection Rules

An inspection rule states the logic by which the CS-MARS tests whether or not a single network event or series of events is a noteworthy incident. An event or series of events with attributes that match the attributes specified in an inspection rule causes the rule to trigger (or “fire”) to create an incident. Incidents may be attacks, network configuration errors, false positives, or just anomalous network activity. The over 100 inspection rules that ship with MARS are called System Inspection Rules. The number and structure of system rules are updated in signature upgrades and with more recent software releases. Both types of upgrades are performed from the Admin > System Maintenance > Upgrade page.

You can create custom inspection rules by editing or duplicating system inspection rules, by adding your own from the Inspection Rules page, or by using the Query interface. Customized inspection rules are called User Inspection Rules and are displayed on the Inspection Rules page.

Inspection rules can be created on both the Global Controller and the Local Controllers.

Global User Inspection Rules

Global Inspection Rules are inspection rules you create on a Global Controller then push to the Local Controller. From the Local Controller, you can edit only the Source IP Address, Destination IP Address, and Action fields of a Global Inspection Rule. To change the arguments of the other fields, you must edit the rule on the Global Controller. When you edit a global inspection rule on the Local Controller then edit it again on the Global Controller, the Global Controller version overwrites the Local Controller version. Global Inspection rule names are displayed with the prefix “Global Rule.”

Drop Rules

Drop rules allow false positive tuning on a MARS, and are defined only on the Local Controller Drop Rules page. They allow you to refine the inspected event stream by specifying events and streams to be ignored and whether those data should be stored in the database or discarded entirely. Drop rules are applied to events as they come in from a reporting device, after they have been parsed and before they have been sessionized. Events that match active drop rules are not used to construct incidents. Because the Global Controller does not receive events from reporting devices, rather it receives them from Local Controllers, you cannot define drop rules for the Global Controller.

Constructing a Rule

Each step of your plan corresponds to a line of a rule. Each line identifies a set of conditions. A rule can have a single line, two lines, or multiple lines. You link these lines together using the logical operators, “AND, OR, FOLLOWED-BY (in time).”

For more information on the conditions and operators found in a rule, see [Table 21-1 on page 21-6](#).

The first step of the example plan, identified in [Back to Being the Admin, page 21-3](#), involved probing the target host. You can express a probe by selecting the appropriate event type groups as the line’s event type criteria. Also, you want to use dollar variables (\$TARGET)¹ to constrain your host to ensure that

For more information on the conditions and operators found in a rule, see [Table 21-1](#).

The first step of the example plan, identified in the section [Back to Being the Admin, page 21-3](#), involved probing the target host. You can express a probe by selecting the appropriate event type groups as the line’s event type criteria. Also, you want to use dollar variables (\$TARGET)² to constrain your host to ensure that the probe and attacks that are reported have happened to the same host. Then you need to figure out the logical step for the next line. In this case, the probe could be optional depending on the time frame that the probe was sent and its subtlety.

Rule logic is simple. You have a row. Every row has cells. The logical expressions connecting different cells are “and,” while the expressions connecting items inside a cell are either “or” or “and not”, depending which clause is chosen—the equal to or not equal to.

By studying the system inspection rules, you can identify three commonly used rules: attempts, success likely, and failures. The most common rule structure is the basic three-line rule that identifies an attempted attack. It is expressed as:

```
(Probe AND
Attack) OR
Attack)
```



Note

To clarify this pseudocode, keep in mind that uppercase AND, OR and FOLLOWED-BY identify a logical operator between two rule lines. Lowercase “and” identifies a logical operator between two cells. Lowercase “or” and “and not” identify a logical operator between two items within a cell.

Success likely rules extend the attempt rules by identifying suspicious activities originating from the attacked host. The general structure of these rules is:

```
((Probe AND
Attack) OR
Attack)) FOLLOWED BY
(Suspicious Activity[1]..Suspicious Activity[n])
```

Failures identify an event from a reporting device that the device classifies as a failure. Often, these rules simply match to known syslog or SNMP messages indicating some failure on the device. You can define alerts to keep you abreast of device failures. These rules follow one of two general structures: a one line failure—

1. A variable, such as (\$TARGET), serves two purposes in the rule: 1.) It captures the number of times the same cell value is matched upon—the count for that cell, e.g., ten login failures from the same source address. 2.) It correlates the same value of a cell across rule lines, e.g., a probe from a source address AND an attack from that same source address.
2. A variable, such as (\$TARGET), serves two purposes in the rule: 1.) It captures the number of times the same cell value is matched upon—the count for that cell, e.g., ten login failures from the same source address. 2.) It correlates the same value of a cell across rule lines, e.g., a probe from a source address AND an attack from that same source address.

Failure

—or multi-line failures separated by the *OR* operator—

1..N Failure OR

Failure

In the HTML interface, system rules are displayed in rows and columns. The row number is called the Offset. A rule can have more than one row (or offset), as shown in [Figure 21-2](#).

Figure 21-2 Rule with Multiple Offsets

☐		Rule Name:		System Rule: Backdoor: Active										Active	
		Action:		None										Time Range: 0h:30m	
		Description:		This correlation rule detects a connection to a backdoor server or a response from a backdoor server in your network accompanied by malicious follow-up activity on the server hosting the backdoor - this may indicate that a malicious backdoor service is likely running in your network. Malicious follow-up activity includes excessive scans, denied packets, installation of malicious services, local buffer overflow attacks etc. Backdoors such as Unix rootkits or Trojan horses are malicious programs that offer extensive remote control of a host and may be left by an attacker on a compromised host to maintain future remote access.											
Offset	Open (	Source IP	Destination IP	Service Name	Event	Device	Reported User	Keyword	Severity	Count	)	Close	Operation		
1	(	ANY	SAME, \$TARGET01, ANY	ANY	Penetrate/Backdoor/Rootkit/Connect, Penetrate/Backdoor/Trojan/Connect, Penetrate/Backdoor/Trojan/SYN, Penetrate/Backdoor/CommandShell, Penetrate/Backdoor/RemoteControlApp/Connect	ANY	None	ANY	ANY	1			OR		
2		SAME, \$TARGET01, ANY	ANY	ANY	Penetrate/Backdoor/RemoteControlApp/Response, Penetrate/Backdoor/Trojan/Response, Penetrate/Backdoor/Trojan/SYN-ACK	ANY	None	ANY	ANY	1	)		FOLLOWED-BY		
3	((	SAME, \$TARGET01, ANY	DISTINCT, ANY	SAME_ANY_DEST_PORT	AttacksProtected, FirewallPolicyViolation/ACL, FirewallPolicyViolation/NAT	ANY	None	ANY	ANY	25			OR		
4		SAME, \$TARGET01, ANY	ANY	ANY	DoS/Network/TCP, DoS/Network/UDP, DoS/Network/ICMP, DoS/Network/Misc, DoS/Distributed, Probe/HostInfo/All, Propagate/CopyFiles, Propagate/Worm, Penetrate/Backdoor/CovertChannel	ANY	None	ANY	ANY	1	)		OR		
5		ANY	SAME, \$TARGET01, ANY	ANY	Persist/All, Penetrate/Backdoor/CovertChannel	ANY	None	ANY	ANY	1	)				

143411

Table 21-1 Rule Fields and Arguments

Rule Field	Field Description and Arguments	Argument Descriptions
Offset	The row number.	
Open (	Identifies the open of a clause. Clauses are used to compare one or more compound conditions in a rule.	Displays the open braces you create a clauses.

Table 21-1 Rule Fields and Arguments

Rule Field	Field Description and Arguments	Argument Descriptions
Source IP	IP address of the packet originator.	
	Variables	<i>ANY</i>—(Default). Signifies that the IP address for each count is any IP address. <i>SAME</i>—Signifies that the IP address for each count is the same IP address. This variable is local to its offset. <i>DISTINCT</i>— Signifies that the IP address for each count is a unique IP address. This variable is local to its offset. <i>\$Target01 to \$Target20</i>—The same variable in another field or offset signifies that the IP address for each count is the same IP address.
	Network Groups	<i>Defined network groups</i> —Topologically valid network groups as defined under Management > IP Management.
	Networks	Topologically valid network groups as defined under Management > IP Management.
	Devices	The hosts and reporting devices present in the system.
	IP addresses	IP addresses present on devices in the system or user entered dotted quads.
	IP ranges	The range of addresses between two dotted quads.

Table 21-1 **Rule Fields and Arguments**

Rule Field	Field Description and Arguments	Argument Descriptions
Destination IP	IP address of the packet destination. Often referred to as the target.	
	Variables	<i>ANY</i>—(Default). Signifies that the IP address for each count is any IP address. <i>SAME</i>—Signifies that the IP address for each count is the same IP address. This variable is local to its offset. <i>DISTINCT</i>—Signifies that the IP address for each count is a unique IP address. This variable is local to its offset. <i>\$Target01 to \$Target20</i>—The same variable in another field or offset signifies that the IP address for each count is the same IP address.
	Network Groups—	<i>Defined network groups</i> —Topologically valid network groups as defined under Management > IP Management.
	Networks—	Topologically valid network groups as defined under Management > IP Management.
	Devices —The hosts and reporting devices present in the system.	The hosts and reporting devices present in the system.
	IP addresses—	IP addresses present on devices in the system or user entered dotted quads.
	IP ranges —The range of addresses between two dotted quads.	The range of addresses between two dotted quads.
Service Name	A TCP/IP-based network service, identified by protocol and port, defined within the packet.	

Table 21-1 **Rule Fields and Arguments**

Rule Field	Field Description and Arguments	Argument Descriptions
	Variables	ANY—(Default) No constraint is placed on the source or destination ports or protocol or port. SAME type variables signify that the specified destination port, source port and protocol are the same for each count. These variables are local to the offset. • SAME_ANY_DEST_PORT SAME_TCP_DEST_PORT SAME_UDP_DEST_PORT • SAME_ANY_SRC_PORT SAME_TCP_SRC_PORT SAME_UDP_SRC_PORT DISTINCT type variables signify that the specified destination port, source port and protocol are unique for each count. These variables are local to the offset. • DISTINCT_ANY_DEST_PORT DISTINCT_TCP_DEST_PORT DISTINCT_UDP_DEST_PORT Identical variables in different fields or offsets signify that the specified port and protocol for each count are identical to each other. • \$ANY_BOTH_PORT5 • \$ANY_DEST_PORT1 to ANY_DEST_PORT5 • \$ANY_SRC_PORT1 • \$TCP_BOTH_PORT1, \$TCP_BOTH_PORT2 • \$TCP_DEST_PORT1 to \$TCP_DEST_PORT5 • \$TCP_SRC_PORT1, \$TCP_SRC_PORT2 • \$UDP_BOTH_PORT1, \$UDP_BOTH_PORT2 • \$UDP_DEST_PORT1 to \$UDP_DEST_PORT5 • \$UDP_SRC_PORT1, \$UDP_SRC_PORT2

Table 21-1 **Rule Fields and Arguments**

Rule Field	Field Description and Arguments	Argument Descriptions
	Defined services —One or more services defined under Management > Service Management.	
	Service groups —One or more service groups defined under Management > Service Management.	• Backdoor • Instant Messaging • Mail Retrieval • Online Game • P2P • Recent Backdoor • TCP-highport • UDP-highport • vulnerable-protocols
Event	Identifies one or more event types. An event type indicates some type of network activity or condition. Sometimes, events reported from different devices and different device types identify the same activity or condition, and therefore, they map to the same event type within MARS. Event types are sorted into event groups, such as “Probe/PortSweep/Stealth”, to catch any of the network conditions identified by the group.	
	Variables —Signify any single event type defined under Management > Event Management, only useful for lines in tandem with the same variable.	• ANY—Any of the active event types can match this rule. • SAME • DISTINCT • \$EVENT_TYPE01, \$EVENT_TYPE10
	Event types —Events that have been merged into types.	• ANY • SAME • DISTINCT • All events
	Event type groups —Groups of event types.	• ANY • SAME • DISTINCT
	Red Severity Event Types—Displays all severe event types	
	Yellow Severity Event Types—Displays all yellow event types	

Table 21-1 Rule Fields and Arguments

Rule Field	Field Description and Arguments	Argument Descriptions
	Green Severity Event Types—Displays all green event types	
Device	The value of this condition can be one of the following:	
	Variables —Signify any single device defined under Admin > System Management > Security and Monitor Devices, only useful for lines in tandem with the same variable.	• ANY—(Default) Specifies that this rule is applied to events generated by any of the reporting devices defined in MARS. • SAME • DISTINCT • Unknown Reporting Device—Specifies that this rule is applied to events generated by any reporting device that is not defined in MARS. • \$DEVICE01 to \$DEVICE10
	• Reporting Devices—Identifies one or more hosts or reporting devices for which events are inspected. Valid values are one or more devices as defined under Admin > System Setup > Security and Monitor Devices.	
	Defined Device Types—	
Reported User	Identifies the active user on the host when this event was recorded. Not all events include this data. The value of this condition can be one of the following:	• ANY—No constraint is placed on the reported user. • NONE—(Default) Specifies that this condition should not be used to match this rule. • Variables—Signify any single user, only useful for lines in tandem with the same variable. • Invalid User Name—Specifies that this condition is met when the user name reported is invalid.

Table 21-1 **Rule Fields and Arguments**

Rule Field	Field Description and Arguments	Argument Descriptions
Severity	The value of this condition can be one of the following:	• ANY—(Default) Specifies that this rule is applied to events of all severity levels. • Green—Restricts this rule to firing against low-severity events. • Yellow—Restricts this rule to firing against medium-severity events. • Red—Restricts this rule to firing against high-severity events.
Count	Identifies the number of items the event must occur before the condition is met. The value for this condition is a whole number ranging between 1 and 100. The default value is 1.  Note Events of the same event type occurring in the same session in a three-second period increment the active count by one. This inherent threshold ensures that a event floods of the same type does not increase the active count arbitrarily and incorrectly fire the rule.	<i>Example usage:</i> When a backdoor rootkit install is detected, the count should be 1 as it is only going to be reported once and it is not something you expect to ever see on your network. However, if you are using deny messages to detect infected hosts, you may want the count value to be higher. For example, you may want to allow for several common mistakes, such as password failures, before firing a rule for the event. People accidentally mistype passwords, they don't accidentally install a rootkit.
Close	Identifies the close of a clause.	

Table 21-1 **Rule Fields and Arguments**

Rule Field	Field Description and Arguments	Argument Descriptions
Operation	The value of this field can be one of the following:	• None—(Default) Defines a single-line rule or a simple condition. • AND—A boolean “and” used to construct a compound condition (two or more lines). This line and the next line must both be satisfied before the compound condition is met. • OR—A boolean “or” used to construct a compound condition (two or more lines). Either this line or the next line can be satisfied to meet the compound condition. • FOLLOWED-BY—Identifies a compound condition (two or more lines). specifically a sequential order of occurrence. Also referred to as a time conditional rule (e.g., Y must happen after X).The condition of this line must be met, and then the condition of the next line must be met before the compound condition is met.

Table 21-1 **Rule Fields and Arguments**

Rule Field	Field Description and Arguments	Argument Descriptions
Time Range	Identifies the period of time over which the count value is augmented. For rules that have a Count value greater than one, the Time Range value determines how long the period should be before the count value is reset. For example, you can assume that if no more than three login attempts have occurred over a 10-minute period that counter can be reset.	Usage Guideline: The Time Range value combined with the Count value can affect the operation of your MARS. Each time an event is captured that satisfied a unique instance of an inspection rule, a monitoring session is constructed to track possible future occurrences until either the Count value is reached or the time period expires.

Table 21-1 **Rule Fields and Arguments**

Rule Field	Field Description and Arguments	Argument Descriptions
Action	Identifies the action that MARS will take when the rule is fired. Actions are user-defined alerts that include an action name and description, which also doubles as the message text provided in the alert. Each action can combine alert techniques, such as email and syslog. Each alert technique can have multiple values. For example, an action can generate two emails, a page, and a SNMP trap. Each rule can have multiple such actions. Alerts can be constructed using one or more of the following techniques:  Note You will see the column Action/Operation. In this case, you can select either one of the following actions or one of the operators .	• NONE—(Default) This action states that no further action will be taken. When NONE value is selected, the firing of the rule causes an event record to be created and stored in MARS. Regardless of the selected action, this record is always created. • Email—Identifies the list of administrators to whom an alert should be sent. An e-mail address must be defined for the selected administrators. • Syslog—Identifies the list of hosts to whom an alert should be sent. You can select any number of devices to which you want a syslog message sent. • Page—Identifies the list of administrators to whom an alert should be sent. The message format is text. A pager number must be defined for the selected administrators. • SNMP—Lists the hosts to which a Simple Network Management Protocol (SNMP) alert can be sent. • SMS—List of users to receive notification by Short Message Service (SMS). The message can be up to 160 characters. An SMS number must be ten numbers and a domain name, for example, 1234567890@provider.com. • Distributed Threat Mitigation (DTM)— Lists the Cisco IOS Intrusion Prevention System (IPS) devices to which an IPS alert action can be sent (alarm, alarm and drop, or alarm and reset if it is a TCP session.) See the Technology Preview: Configuring Distributed Threat Mitigation with Intrusion Prevention System in Cisco Security MARS, page 1 document for DTM configuration information.

Working Examples

The examples in this section demonstrate the use of variables, in particular, how to use variables to detect Deny patterns.



Note

We recommend that you study the system inspection rules for more complex examples.



Note

For a single offset rule, the variables SAME and SAME_ANY_DEST_PORT can be substituted in any of the examples for \$TARGET01 and \$ANY_DEST_PORT1, respectively. The “ANY” in \$ANY_DEST_PORT1 means either UDP or TCP protocol.

Example A: Excessive Denies to a Particular Port on the Same Host

Figure 21-3 Rule for Excessive Denies to a Particular Port on the Same Host

☐	Rule Name:		Example A							Status:		Active
Action:									Time Range:		0hh:0mm:10ss	
Description:		Excessive denies to a particular port on the same host.										
Offset	Open (	Source IP	Destination IP	Service Name	Event	Device	Severity	Counts	Zone	) Close	Operation	
1		ANY	\$TARGET01	\$ANY_DEST_PORT1	FirewallPolicyViolation/ACL	ANY	ANY	100	Training			

In this example, the rule fires when 100 of the specified events occur from any source IP address to the same destination IP address, and the destination port numbers are identical.

Example B: Same Source Causing Excessive Denies on a Particular Port

Figure 21-4 Rule for Same Source Doing Excessive Denies on a Particular Port

☐	Rule Name:		Example B							Status:		Active
Action:									Time Range:		0hh:0mm:10ss	
Description:		Same source doing excessive denies on a particular port.										
Offset	Open (	Source IP	Destination IP	Service Name	Event	Device	Severity	Counts	Zone	) Close	Operation	
1		\$TARGET01	ANY	\$ANY_DEST_PORT1	FirewallPolicyViolation/ACL	ANY	ANY	100	Training			

In this example, the rule fires when 100 of the specified events occur that have the source IP address, any Destination IP address, and identical destination port numbers.

Example C: Same Host, Same Destination, Same Port Denied

Figure 21-5 Rule for Same Host, Destination, Same Port Denied

☐	Rule Name:		Example C						Status:		Active		
Action:										Time Range:		0hh:0mm:10ss	
Description:		Same host, destination, same port getting denied.											
Offset	Open (	Source IP	Destination IP	Service Name	Event	Device	Severity	Counts	Zone	) Close	Operation		
1		\$TARGET01	\$TARGET02	\$ANY_DEST_PORT1	FirewallPolicyViolation/ACL	ANY	ANY	20	Training				

In this example, the rule fires when 20 of the specified events occur that have the same source and destination addresses, and identical destination port numbers.

Working with System and User Inspection Rules

Navigate to the Rules page by clicking the **Rules** tab. From this tab, you can access the **System Inspection Rules** and **Drop Rules** tabs.

You can perform the following actions with System Inspection rules:

- Change the Source IP, Destination IP and Device arguments of system rules
- Duplicate system rules
- Create (Add), delete and edit your own rules (user rules)
- Toggle the state of the rules between active and inactive
- Add, edit, or delete a rule group

**Note**

Upgrade the MARS software regularly to obtain new and updated System Inspection rules. For more information, see the *Install and Setup Guide for Cisco Security Monitoring, Analysis, and Response System*.

**Note**

When you edit a rule, you must click **Activate** to enable the changes.

Change Rule Status—Active and Inactive

The CS-MARS correlation engine continuously tests only active rule criteria against incoming events to identify incidents. Inactive rules do not consume resources used for realtime operations.

**Note**

A rule cannot be deleted, it can be made active or inactive.

To change the status of a rule, follow these steps:

- Step 1** Navigate to the **Rules > Inspection Rules** page.
- Step 2** Select the checkbox of the rule (or rules) to change.
- Step 3** Click **Change Status**.
The selected rules are made inactive if active, and active if inactive and displayed on a different page.
- Step 4** To display inactive rules, select **Inactive** from the View dropdown list. To display active rules, select **Active**.

Duplicate a Rule

Duplicating a rule creates a new rule that is a copy of an existing system or user inspection rule. You can edit all of the fields of a duplicate rule, but only the Source IP, Destination IP, and Device fields of a system inspection rule. The original rule is left unchanged after duplication.

**Note**

You cannot delete a rule after it is created by **Duplicate** or **Add**.

To duplicate a rule, follow these steps:

Step 1 Select the checkbox of the rule to duplicate.

Step 2 Click **Duplicate**.

The name of duplicated rule is the name of the original rule extended with a timestamp of when the original was duplicated (for example, System Rule: Client Exploit - Sasser Worm Copied: 05.10.05/16:54:21). The name can be changed by editing the duplicate rule.

Edit a Rule

You can edit rules with inline editing, or with the rule wizard. To edit inline, you click the argument to edit. The rule wizard is invoked by selecting a rule to edit then clicking **Edit**. The rule wizard begins with the Rule Name field and progress through each subsequent field.

**Note**

You only edit the Source IP, Destination IP, and Device fields of a system inspection rule. See [Duplicate a Rule, page 21-17](#) for further information on modifying system inspection rules.

**Note**

A rule cannot be deleted, it can be made active or inactive.

Edit a Rule with Inline Editing

You can perform inline editing to rules from the Incidents Detail page, or from the Inspections Rules page. To edit a rule with the Inline Editing, follow these steps:

Step 1 Click the Rule argument that you want to edit.
The edit page for the selected field appears.

Step 2 Change the argument, then click **Apply**.

Step 3 Repeat [Step 1](#) as required.

Step 4 Add Open and Close parentheses as required then click **Submit**.
If no parentheses are required, just click **Submit**.

Step 5 Click **Activate** to include the rule in event correlation processing.

Edit a Rule with the Rule Wizard

The Rule Wizard can only be invoked from the Inspections Rule page.

To edit a rule with the Rule Wizard, follow these steps:

Step 1 Select the check box of the rule to edit.

- Step 2** Click **Edit**.
The rule wizard page appears for the Rule Name field.
- Step 3** Do one of the following actions:
- Change the argument of the field, then click **Apply**. Proceed to [Step 6](#).
 - Change the argument, then click **Next** to proceed to the next field.
 - Click **Next** to proceed to the next field without changing the argument.
 - Click **Previous** to go back to the previous field.
Previous does not appear for the Rule Name page.
- Step 4** Repeat • as required.
- Step 5** Click **Apply** after making all edits.
-  **Tip** To skip to the end, click the Count argument, after which, only the **Action**, and **Time Range** fields must be reviewed.
- Step 6** Add Open and Close parentheses as required then click **Submit**.
If no parentheses are required, just click **Submit**.
- Step 7** Click **Activate** to include the rule in event correlation processing.

Add an Inspection Rule



Note Rules that you add are called User Inspection Rules.

- Step 1** Navigate to the Inspection Rules page.
- Step 2** Click **Add**.
- Step 3** Enter a name and description for the rule, then click **Next**.
- Step 4** Select Source IP address .

Figure 21-6 User Inspection Rule Creation Form


The following numbers correspond to the numbers shown in Figure 21-6.

1. Check the boxes next to the items in the **Sources Selected** field to select them, and click the **Toggle Equal** button to change them between equal and not equal.
 2. Click the **Select All** button to select all items in the **Sources Selected** field. Items selected in the Sources Selected field are deselected when you click **Select All**.
 3. Use the **Equal** and **Not Equal** buttons to bring highlighted items from the **Sources Available** field into the **Sources Selected** field.
 4. Filter sources from this drop-down list.
 5. Enter search text, and click **Enter** to move items that match the search criteria from the **Sources Available** field to the **Sources Selected** field.
 6. To add a new item to the sources, click the **Add** button. To edit or delete an existing source, click the **Edit** or **Delete** button.
 7. Click an item or items in the Sources Selected field, and use the **Remove** button.
 8. To move IP values up into the Sources Selected field, click the **Equal** ☐ **==** up icon, or the **Not Equal** ☐ **!=** up icon.
 9. Check the radio button next to **IP** or **Range**, and enter an IP address or a range of IP addresses into their respective fields.
 10. Select items in the Sources Selected field by clicking them. Enter a group name, and click the **Grouped As** button to group them.
- Step 5** Follow the wizard, and select the values for the rule, clicking the **Next** button to progress to the next step.
- Step 6** When you are asked, “Are you done defining the rule conditions,” you can:
- Click the **Yes** button for a single line rule. Continue to add repetition requirements (counts), alert information, and valid time ranges for each line.

- Click the **No** button, to create a multi-line rule that uses an operator (OR, AND, or FOLLOWED BY). Return to [Step 4](#) and continue to make your selections. Continue to add rule information, and click **Submit** when finished.
- Click the **Submit** button when finished.

Step 7 When the rule is complete, you need to activate it by clicking the **Activate** button.

Figure 21-7 Clicking the Activate button



Note

If you are creating or editing several rules, it is better for the system to click the **Activate** button for several changes rather than for each individual change.

Working with Drop Rules

Navigate to the Drop Rules page by clicking the **Rules > Drop Rules** tabs.

Drop rules instruct the MARS to either drop a false positive completely from the appliance, or to keep it in the database. On the Drop Rules page, you add, edit, duplicate, activate an inactive rule, or inactivate an active rule. Inactive rules do not fire.

While working with drop rules is similar to working with inspection rules, it is not identical.

Change Drop Rule Status— Active and Inactive

-
- Step 1** Check the box next to the rule.
- Step 2** Click **Change Status**.
- When you change the status to inactive, the rule displays only on the inactive rules page.
- Step 3** To display inactive Drop Rules, select **Inactive** from the **View** dropdown list.
-

Duplicate a Drop Rule

-
- Step 1** Check the box next to the rule.
- Step 2** Click the **Duplicate**.
- After duplicating a rule, you can edit the duplicat without altering the original.
-

Edit a Drop Rule

- Step 1** Check the box next to the rule.
- Step 2** Click **Edit** on the field that you want to change.
- Step 3** Follow the rule's wizard and complete any other changes to the rule.
- Step 4** Click **Submit**.



Note When the rule or rules are complete, click **Activate**.

Add a Drop Rule

- Step 1** Click **Add**.
- Step 2** Enter a name and description for the rule, and click **Next**.
- Step 3** Select your sources.


Figure 21-8 Drop Rule Creation Form

The following numbers correspond to the numbers in the [Drop Rule Creation Form](#) as shown in [Figure 21-8](#):

1. Check the boxes next to the items in the **Sources Selected** field to select them, and click the **Toggle Equal** button to change them between equal and not equal.

2. Click the **Select All** button to select all items in the **Sources Selected** field. (Note: if you have items highlighted in the Sources Selected field, clicking **Select All** will de-select them.)
 3. Use the **Equal** and **Not Equal** buttons to bring highlighted items from the **Sources Available** field into the **Sources Selected** field.
 4. Filter sources from this drop-down list.
 5. Enter search text, and click **Enter** to move items that match the search criteria from the **Sources Available** field to the **Sources Selected** field.
 6. To add a new item to the sources, click the **Add** button. To edit or delete an existing source, click the **Edit** or **Delete** button. See [IP Management, page 23-3](#) for more information.
 7. Click an item or items in the Sources Selected field, and use the **Remove** button.
 8. To move IP values up into the Sources Selected field, click the **Equal**  (Up) icon, or the **Not Equal**  (Up) icon.
 9. Check the radio button next to **IP** or **Range**, and enter an IP address or a range of IP addresses into their respective fields.
 10. Select items in the Sources Selected field by clicking them. Enter a group name, and click the **Grouped As** button to group them.
- Step 4** Follow the wizard, and select the values for the rule, clicking the **Next** button to progress to the next step.
- Step 5** When you are asked, “Are you done defining the rule conditions,” click the **Submit** button.
- Step 6** When the rule is complete, you need to activate it by clicking the **Activate** button.

Figure 21-9 Clicking the Activate button



Note

If you are creating or editing several rules, it is better for the system to click the **Activate** button for several changes rather than for each individual change.

Setting Alerts

You have two options for learning about rules that have fired: you can log in and view the appropriate pages in the HTML interface or you can have MARS send alerts to external devices and users. Actions provide instructions to MARS on the second method.

Using Rules, you can alert a person if a rule has fired. The roles and groups you can choose are determined by the information you have entered in User Management. for more information on adding users into the Local Controller.

Configure an Alert for an Existing Rule

-
- Step 1** Click on a rule argument.
- Step 2** Click **Next** until the Action/Operation column is selected.
- Step 3** Click the **Add** button to add users for an alert.
- Step 4** Enter a **Name** and **Description** for the notification.
- Step 5** Check the box next to the type of notification that you want to send. Your choices are:
- **Email** – select the roles or groups that you want to receive an email.
 - **Syslog** – select the systems that you want to receive the syslogs.
 - **Page** – select the roles or groups that you want to receive an electronic page on their pagers or cellular telephones.
 - **SNMP** – select the systems that you want to receive the SNMP trap information.



Note

For SNMP and Syslog, you need to configure the receiving systems for this feature to work.

- Step 6** Click the **Change Recipient** button to add or edit recipients for alerts for that notification type (email, syslog, page, or SNMP).
- Step 7** Check the box next to the role, group, or system that you want to receive alerts.
- Click the **Add** button to select recipients (to move them into the left field.)
 - To remove recipients, click their names to highlight them (in the left field) and click the **Remove** button.
- Step 8** Repeat steps 5 - 7 for all the alert selections that you want to include.
- Step 9** Click the **Submit** button.
- Step 10** Click the **Apply** button.



Note

If a user adds an alert to a rule created on the Global Controller, and the rule is pushed down and fired on the Local Controller, the designated user receives the alert from the Local Controller and not the Global Controller

Rule and Report Groups

This section contains the following subsections:

- [Rule and Report Group Overview, page 21-25](#)
- [Global Controller and Local Controller Restrictions for Rule and Report Groups, page 21-26](#)
- [Add, Modify, and Delete a Rule Group, page 21-26](#)
- [Add, Modify, and Delete a Report Group, page 21-29](#)
- [Display Incidents Related to a Rule Group, page 21-31](#)

- [Create Query Criteria with Report Groups, page 21-32](#)
- [Using Rule Groups in Query Criteria, page 21-33](#)

Rule and Report Group Overview

Rule and report groups help you manage rules and reports by speeding access to those rules and reports relevant to your task at hand. You can create groups, or use the groups provided with CS-MARS (System groups). Groups act as filters to limit the display of rules, reports, and incidents in the CS-MARS HTML interface. All groups can be modified or deleted.

CS-MARS provides over 100 system rules and 150 system reports. More can be added by creating custom rules and reports, and by performing periodic software updates. A rule or report group contains a subset of these rules or reports as members. Usually rules or reports within the same group have related functions (such as, reconnaissance activities, server attack, etc.). When you select a group from a dropdown filter, only those rules and reports that are members are displayed on the page. When you select a rule group on the Incidents page, only those incidents related to the rules of the selected group display. Report and rule groups can also be used when constructing queries.

For instance, there are at least 16 system rules that detect suspicious network access events and incidents, and 15 system reports to report this information. CS-MARS provides a system rule group and a system report group named “Access” that can filter the Inspection Rules, Incidents, and Report pages to display only those rules and reports related to monitoring access event (such as password attacks), thereby eliminating the need to search for the pertinent rules and reports within the complete rule and report pages or dropdown lists. CS-MARS provides system rule and report groups as listed in [Table 21-2](#).

Table 21-2 **Predefined Rule and Report Groups**

System Report Groups	Corresponding System Rule Groups
System: Access	System: Access
System: All Events - Aggregate View	—
System: All Exploits - Aggregate View	—
System: COBIT DS3.3 - Monitoring and Reporting	—
System: COBIT DS5.10: Security Violations	—
System: COBIT DS5.19: Malicious software	—
System: COBIT DS5.20: Firewall control	—
System: COBIT DS5.2: Authentication and Access	—
System: COBIT DS5.4: User Account Changes	—
System: COBIT DS5.7: Security Surveillance	—
System: COBIT DS9.4: Configuraton Control	—
System: COBIT DS9.5: Unauthorized Software	—
System: CS-MARS Distributed Threat Mitigation (Cisco DTM)	System: CS-MARS Distributed Threat Mitigation (Cisco DTM)
System: CS-MARS Incident Response	System: CS-MARS Incident Response
System: CS-MARS Issue	

Table 21-2 *Predefined Rule and Report Groups (continued)*

System Report Groups	Corresponding System Rule Groups
System: Client Exploits, Virus, Worm and Malware	System: Client Exploits, Virus, Worm and Malware
System: Configuration Changes	—
System: Configuration Issue	System: Configuration Issue
System: Database Server Activity	System: Database Server Activity
System: Host Activity	System: Host Activity
System: Network Attacks and DoS	System: Network Attacks and DoS
System: New Malware Outbreak (Cisco ICS)	System: New Malware Outbreak (Cisco ICS)
System: Operational Issue	System: Operational Issue
System: Reconnaissance	System: Reconnaissance
System: Resource Issue	System: Resource Issue
System: Resource Usage	—
System: Restricted Network Traffic	System: Restricted Network Traffic
System: SOX 302(a)(4)(A)	—
System: SOX 302(a)(4)(D)	—
System: Security Posture Compliance (Cisco NAC)	System: Security Posture Compliance (Cisco NAC)
System: Server Exploits	System: Server Exploits

Global Controller and Local Controller Restrictions for Rule and Report Groups

Global Controller and Local Controller rule and report groups have the following restrictions:

- Rule and report groups created on the Global Controller are pushed to all the Local Controllers.
- Rule groups created on a Local Controller are local to the Local Controller. They are not copied to the Global Controller or to other Local Controllers.
- Local Controller account holders can edit only the Source IP, Destination IP, and Device fields of a rule group created on a Global Controller.
- Local Controller account holders cannot edit Global Controller report groups.
- Local Controller account holders cannot delete Global Controller rule and report groups.



Note

The procedures described in this section are valid for both the Local and Global Controllers, except that the Case Bar does not appear on the Global Controller HTML interface.

Add, Modify, and Delete a Rule Group

Adding a New Rule Group

To add a rule group follow these steps:

Step 1 Navigate to the Inspection Rules page, as shown in [Figure 21-10](#).

Figure 21-10 Inspection Rules Page

Step 2 Click **Add Group**.
The Add Group dialog box appears, as shown in [Figure 21-11](#).

Figure 21-11 Add Group Dialog Box

Step 3 Enter the new group name in the **Name** field.

Step 4 Click the checkboxes of the rules to be added to the new rule group.




Tip The dropdown list above the list of rules can limit the display of rules to active system rules, active user rules, or inactive rules. The search function displays only those rules that match a search string (for example, “New Malware Traffic Match.”). The asterisk wildcard character (*) is supported.

Step 5 Click **Add**.

The selected rules appear in the lefthand pane of the dialog box. To remove a rule from the group, highlight the item in the lefthand pane and click **Remove**.

Step 6 Click **Submit**.

The new rule group name appears in the **Group** dropdown filter on the Inspection Rules page, as shown in [Figure 21-12](#). In this example, the new rule group name is “Example Group.” Because it is a user-created rule group, the rule group name appears without the prefix “System.” You can also click **Cancel** to return to the Inspection Rules page without creating a new rule group.

Figure 21-12 New Rule Group Appears on the Dropdown List of the Inspections Rules Page

Modifying a Rule Group

To edit a rule group, follow these steps:

- Step 1** Navigate to the Inspection Rules page, as shown in [Figure 21-10](#).
- Step 2** Select the rule group to edit in the **Group** pulldown filter.
- Step 3** Click **Edit Group**.
The Add Group dialog box appears, as shown in [Figure 21-11](#). The rule group name appears in the **Name** field, and the included rules appear as selected rules in the lefthand pane of the dialog box.
- Step 4** To add additional rules, click the checkbox of all the rules to be added to the group, then click **Add**. To remove rules, highlight the items in the lefthand pane to remove, then click **Remove**.
- Step 5** Click **Submit**.

Deleting a Rule Group

- Step 1** Navigate to the Inspection Rules page, as shown in [Figure 21-10](#).
- Step 2** Select the rule group to delete in the **Group** pulldown filter.
- Step 3** Click **Delete Group**.
The Delete Group dialog box appears listing the rules in the group to be deleted. You are prompted to confirm deletion.
- Step 4** Click **Yes**.
The rule group no longer appears in the **Group** dropdown filters on the Incident and Inspection Rules pages.

Add, Modify, and Delete a Report Group

Adding a New Report Group

To add a report group follow these steps:

Step 1 Navigate to the Report page, as shown in [Figure 21-13](#).

Figure 21-13 Reports Page

Step 2 Click **Add Group**.
The Add Group dialog box appears, as shown in [Figure 21-14](#).

Figure 21-14 Add Report Group Dialog Box

Step 3 Enter the new report group name in the **Name** field.

Step 4 Click the checkboxes of the reports to be added to the new report group.



Tip

The dropdown filter above the list of reports can filter the display of reports to display system reports, user reports, or all reports. The search function displays only those reports that match a search string (for example, “Spy” for Spyware). The asterisk wildcard character (*) is supported.


Step 5 Click **Add**.

The selected reports appear in the lefthand pane of the dialog box. To remove a report from the group, highlight the item in the lefthand pane and click **Remove**.

Step 6 Click **Submit**.

The new report group name appears in the **Group** dropdown list display filter on the Report page, as shown in [Figure 21-15](#), and on the Query Page. Because it is a user-created report group, the report group name appears without the prefix “system.” You can also click **Cancel** to return to the Report page without creating a new report group.

Figure 21-15 The New Report Group Appears on the Dropdown Filter of the Report Page



Modifying a Report Group

To edit a report group, follow these steps:

-
- Step 1** Navigate to the Reports page, as shown in [Figure 21-13](#).
 - Step 2** Select the report group to edit from the **Group** pull-down list.
 - Step 3** Click **Edit Group**.
The Add Report Group dialog box appears, as shown in [Figure 21-14](#). The report group name appears in the **Name** field, and the reports that comprise the report group appear in the lefthand pane of the dialog box.
 - Step 4** To add additional reports, click the checkboxes of the reports to be added to the group, then click **Add**. To remove reports, highlight the items to remove in the lefthand pane, then click **Remove**.
 - Step 5** Click **Submit**.
-

Deleting a Report Group

-
- Step 1** Navigate to the Reports page, as shown in [Figure 21-13](#).
 - Step 2** Select the report group to delete in the **Group** pulldown filter.

- Step 3** Click **Delete Group**.
The Delete Report Group dialog box appears listing the reports in the group to delete. You are prompted to verify deletion.
- Step 4** Click **Yes**.
The report group no longer appears in the report group dropdown lists on the Report and Query pages.

Display Incidents Related to a Rule Group

To display incidents that occur from the firing of rules in a specific rule group, follow these steps:

- Step 1** Navigate to the Incidents page.
- Step 2** Select the rule group in the dropdown filter above the Matched Rules column, as shown in [Figure 21-16](#). The Incidents page will display only those incidents that occurred from rules firing in the selected rule group.

Figure 21-16 Rule Group on Incidents Page

The screenshot shows the Cisco Systems Incidents page. At the top, there is a navigation bar with 'SUMMARY' and 'INCIDENTS' tabs. Below this, there are tabs for 'Incidents', 'False Positives', and 'Cases'. The 'Incidents' tab is selected, and the page title is 'INCIDENTS | CS-MARS Standalone: LC20-Doc v4.1'. Below the title, there is a 'Select Case:' dropdown menu with 'No Case Selected...' selected. Below this, there is a 'Recent Incidents' section with a 'View' button. A dropdown menu is open, showing a list of rule groups. The 'All Rules' option is selected. The list of rule groups includes: 'All Rules', '--Rule Groups--', 'Example Group', 'System Group: Access', 'System Group: CS-MARS Incident Response', 'System Group: Client Exploits, Virus, Worm and Malware', 'System Group: Configuration Issue', 'System Group: Database Server Activity', 'System Group: Host Activity', 'System Group: Network Attacks and DoS', and 'System Group: Operational Issue'. The table below the dropdown menu shows incident details.

Incident ID	Event Type	Action	TI
I:10985516	Inactive CS-MARS reporting device	Alert	Se
I:10985515	Inactive CS-MARS reporting device		Se
I:10985514	Inactive CS-MARS reporting device	Alert	Se
I:10985513	Inactive CS-MARS reporting device		Se
I:10985512	Inactive CS-MARS reporting device	Alert	Se

Create Query Criteria with Report Groups

To create queries from report groups, follow these steps:

- Step 1** Navigate to the Query page.
- Step 2** Select a report group in the **Load Report as On-Demand Query with Filter** dropdown filter, as shown in Figure 21-17.
- Only the reports that comprise the report group can now display in the Select Report dropdown list, as shown in Figure 21-18.

Figure 21-17 Selecting A Report Group to Make a Query





Figure 21-18 Selecting a Report Within the Report Group to Make a Query



- Step 3** Select the report in the secondary dropdown list.
- The **Query** criteria are automatically populated per the selected report.

Using Rule Groups in Query Criteria

To populate the Rule field of the **Query Event Data** bar using rule groups, follow these steps:

- Step 1** Navigate to the Query page.
- Step 2** Click **Any** in the **Rules** field of the **Query Event Data** bar.
The Filter by Rule dialog box appears as shown in [Figure 21-19](#).
- Step 3** Select the rule group in the dropdown list above the list of rules, as shown in [Figure 21-14](#).
The list of rules will display only those rules in the selected rule group.

Figure 21-19 Rule Group Used to Populate Rule Criterion in Query

- Step 4** Click the checkboxes of the rules to include in the query.
- Step 5** Click **Add**. The selected items appear in the lefthand pane of the Query dialog box.
To remove rules, highlight the items to remove in the lefthand pane, then click **Remove**.
- Step 6** Click **Apply**.
The selected rules appear in the **Rules** field of the **Query Event Data** bar.



Sending Alerts and Incident Notifications

A Cisco Systems MARS alert action is a signal transmitted to people or devices as notification that a MARS rule has fired, and that an incident has been logged. Alert actions can only be configured through the Action parameter of a rule. An alert action determines which alert notification types are sent to which MARS user accounts or user groups. MARS can transmit alerts by the methods listed in [Table 22-1](#).

Table 22-1 MARS Incident Notification Methods

Alert Notification Type	Description
Sent in Human-Readable Format	
- E-mail - XML Notification - Short Message Service (SMS) - Pager	E-mail, SMS, and pager alerts send the incident ID, matched rule name, severity, and incident time in email, SMS and pager formats respectively. You must login to the MARS to view all the incident details. XML notification sends an email notification of an incident with an attached XML data file (see Example 22-2). The XML data file contains the same incident details that can be viewed from the GUI, except for path and mitigation information. The XML data file can be sent as a plain-text file or as a compressed gzip file. The XML data filename is constructed with the incident ID number, for example <code>CS-MARS-Incident-13725095.xml</code>. You can parse and extract data from the XML file with a custom application. For example, you can integrate the XML data with trouble ticketing software. See Appendix A, “Cisco Security MARS XML API Reference,” for further information on the MARS XML notification schema and usage guidelines. MARS SMS text message notifications can be up to 160 characters in length. Because the MARS SMS incident notification exceeds 160 characters, it is sent in three segments. Pager messages are sent through the MARS internal modem. MARS dials a carrier’s IXO/TAP number and uses SNPP to transmit the alpha-numeric page. Pager notifications are still possible when the network is down. Pagers can often receive messages in places where mobile phones are inoperative or forbidden (for instance, hospitals).
Sent to a Device	
- SNMP trap - Syslog - Distributed Threat Mitigation	These alerts send the incident ID, matched rule severity, and incident time to devices or applications, all of which must be properly configured within the MARS device administration pages. See the section, Reporting and Mitigation Devices Overview, page 2-1 for information on configuring individual devices to work with MARS.

Table 22-2 provides links and description of related Alert Action configuration procedures. Although some of these procedures are documented elsewhere in this user guide, they are duplicated here for your convenience.

Table 22-2 Alert Notification Procedures

Alert Related Procedures	Description
Configure the E-mail Server Settings	To send Email, SMS, and XML notifications, MARS requires that you configure the E-mail Server settings.
Configure a Rule to Send an Alert Action	Complete this procedure to create or modify an alert action.
Create a New User—Role, Identity, Password, and Notification Information	Alert notifications can be sent only to user accounts configured on MARS. A new user account can be configured from the User Management tab, or when creating an alert action for a rule. This is where you enter the service provider phone numbers and email addresses for E-mail, SMS, Pager, and
Create a Custom User Group	Complete this procedure to create a MARS user group other than the default MARS user groups. Unlike default user groups, custom groups can be edited.
Add a User to a Custom User Group	Complete this procedure to include a newly created user account into a MARS user group.

Example 22-1 shows a typical email alert notification. Example 22-2 shows an XML notification with its attached XML data file. When compression is configured, the XML data file arrives as a GZIP compressed file.



Note

Alert notifications cannot be customized.

Example 22-1 MARS Notification by Email

-----Original Message-----

From: notifier.Latest@serviceprovider.cisco.com [<mailto:notifier.MyLatest@cisco.com>]

Sent: Monday, May 15, 2006 8:48 AM

To: Naliza Mahda (Nalmah)

Subject: Incident Notification (green, Rule Name: System Rule: CS-MARS Database Partition Usage)

The following incident occurred:

Start time: Mon May 15 08:47:26 2006

End time: Mon May 15 08:47:26 2006

Fired Rule Id: 134473

Fired Rule: System Rule: CS-MARS Database Partition Usage

Incident Id: 597842933

For more details about this incident, please go to:


https://MyLatest/Incidents/IncidentDetails.jsp?Incident_Id=597842933

https://MyLatest.cisco.com/Incidents/IncidentDetails.jsp?Incident_Id=597842933
https://10.2.3.7/Incidents/IncidentDetails.jsp?Incident_Id=597842933
https://192.168.1.101/Incidents/IncidentDetails.jsp?Incident_Id=597842933

For all recent incidents, please go to:

<https://MyLatest/Incidents/>
<https://MyLatest.cisco.com/Incidents/>
<https://10.2.3.7/Incidents/>
<https://192.168.1.101/Incidents/>

Example 22-2 MARS XML Notification Email Attachment



Configure the E-mail Server Settings

To send alert actions, MARS must be configured to communicate with an e-mail server. To configure the e-mail server settings, follow these steps:

Step 1 Click **Admin > Configuration Information**.

The Device Configuration window appears, as shown in [Figure 22-1](#).

Figure 22-1 MARS Device Configuration Window

CS-MARS Device Config

→ Name:

Interface Name	IP Address								Net Mask			Default Gateway			
eth0	10	89	149	151	255	255	255	128	10	89	149	254			
eth1	192	168	1	100	255	255	255	0							

→ Mail Gateway:

IP:Port :

Email domain name: (ex: Enter 'domain1' for user@domain1)

- Step 2** In the **IP:Port** field of the **Mail Gateway** section, enter the IP address and **Email Domain Name** of your Mail Gateway server.
- Step 3** Click the **Update** button at the bottom of the page to update the MARS configuration.

Configure a Rule to Send an Alert Action

To send alert notifications to individual users or groups of users, configure the Action parameters of a rule to create an alert action. This procedure configures alerts for pre-existing rules. When you create a rule, the Action parameters are configured after the count number parameter.




Note

Drop rules do not have Action parameters and cannot trigger alerts.

To modify or create an alert for an existing rule, follow these steps:

- Step 1** Click the **RULES** tab to navigate to the Inspection Rules page.
- Step 2** Identify the Rule to configure, and click the value displayed in the **Action** field.

The Action Selection dialog box, as shown in [Figure 22-2](#), appears below the rule description table. All previously defined alert actions are listed in the right-hand area of the Action dialog box. An alert action determines which alert notifications are sent to which users or user groups when the rule fires. You can edit or delete existing alert actions or create a new one.

Figure 22-2 Action Selection Dialog

Step 3 Do one of the following five actions:

1. Remove an alert action currently applied to the rule.
 - In the left-hand area, pick the alert actions to remove with Ctrl+Click, then click **Remove >>**.
The alert action is deleted from the left-hand area.
 - Proceed to Step 13 to complete the procedure.
- Apply an existing alert action to the rule.
 - In the right-hand area, click the check boxes of the alert actions you require, then click <<== .
The alert action appears in the left-hand area.
 - Proceed to Step 13 to complete the procedure.
- Delete an existing alert action from MARS.
 - Click the check box of the alert action in the right-hand area, then click **Delete**.
A delete verification window appears.
 - Click **Yes**.
The alert action is deleted from the right-hand area.
 - Proceed to Step 13 to complete the procedure.
- Edit an existing alert action.
 - Click the check box of the alert action in the right-hand area, then click **Edit**.
The Alert recipients page appears in a new window, as shown in Figure 22-3.
 - Proceed to Step 4 to complete the procedure.
- Create a new alert action.
 - Click **Add**.
The Alert recipients page appears in a new window, as shown in Figure 22-3.

- Proceed to Step 4 to complete the procedure.

Figure 22-3 *Alert Recipients Window*

Name:
 Description:

☐ Email

☐ Syslog

☐ Page

☐ SNMP

☒ SMS

☐ Distributed Threat Mitigation

☐ Alarm ☐ Drop ☐ Reset
☐ Deny Attacker ☐ Deny Flow

☒ XML Email

☐ Compress

Lucre, Phremeus

Lucre, Phremeus

143790

Step 4 For a new alert enter a name and description in the **Name** and **Description** fields. If editing an existing alert, you can modify the name or description.

Step 5 Click the check box of a notification type to select or deselect it.

Recipients for the notification types are as follows:

- **E-mail**—Users or user groups can receive an e-mail.
- **Page**—Users or user groups can receive an alpha-numeric electronic page on their pagers or pager-enabled mobile telephones.
- **SMS**—Users or groups can receive a text message on their SMS-enabled mobile telephones.


- **XML Email**—Users or groups can receive an email message with incident details appended in an XML data file. Click the **Compress** check box to send the XML data file as a compressed gzip file. For more information on this feature, see [Appendix A, “Cisco Security MARS XML API Reference.”](#)
- **Syslog**—Specified devices can receive syslog messages.
- **SNMP**—Specified devices can receive SNMP trap information.
- **Distributed Threat Mitigation**—For more information on this feature, see [Technology Preview: Configuring Distributed Threat Mitigation with Intrusion Prevention System in Cisco Security MARS](#), page 1.

**Note**

For SNMP and Syslog, you must configure the receiving systems to receive notifications.

- Step 6** Click the **Change Recipient** button to add or remove a recipient for a notification type. For E-Mail, Page, SMS, and XML Email, the **Select** (recipient) dialog box appears, as shown in [Figure 22-4](#).

Figure 22-4 **Select Recipient Dialog Box**



143782

For Syslog and SNMP, the **Select** (device) dialog box appears, as shown in [Figure 22-5](#).

Figure 22-5 **Device Selection Page**



For Distributed Threat Management notification, the Select (IOS-IPS Devices) dialog box appears (not shown).



Tip

If you do not know the group to which a user or device belongs, select **All** from the dropdown list to view all users or devices.

- Step 7** Click the check box next to the users or device you want to receive the notification, then click << **Add**. Your selections appear in the left-hand area. To remove items, Ctrl+click the items in the left-hand area, then click **Remove**. The items are then deleted from the left-hand area.
- Step 8** If you are not adding a user, skip to [Step 9](#). To add a new user, do the following substeps:
- Click **Add**.
The User Configuration page appears in a separate window, as shown in [Figure 22-6](#).
 - Enter the User Configuration information then click **Submit**.
You are returned to the [Select Recipient Dialog Box](#).
For reference on user configuration fields, see the section, “[Create a New User—Role, Identity, Password, and Notification Information](#)”
 - Add the new user to the recipient list as described in [Step 7](#).
- Step 9** Click **Submit**.
You are returned to the [Alert Recipients Window](#).
- Step 10** Repeat [Step 6](#) through [Step 9](#) until you have assigned recipients to all the notification types you have selected.
- Step 11** Click **Submit**.

You are returned to the [Action Selection Dialog](#). Any newly-created or edited action alert appears in the right-hand area.

Step 12 Click the check boxes next to the action alerts to be sent when the rule fires. Click << **Add**.

Your selections appear in the left-hand area.

Step 13 Click **Next**.

The Time Range dialog may or may not appear.

Step 14 Click **Next** if the Time Range dialog appears.

The Rule Summary table appears.

Step 15 Click **Submit** to save your changes to the rule.

Step 16 Verify that the alert actions you selected appear in the Action field of the rule description.



Note An inactive rule is made active by applying an alert action. To inactivate a rule, select the rule and click **Change Status**.

This ends the [Configure a Rule to Send an Alert Action](#) procedure.

Create a New User—Role, Identity, Password, and Notification Information

To create a new MARS user, complete the following steps:

New user accounts and user groups are created on the **Management > User Management** tab, or as a substep in creating an alert notification recipient (with the **Add** button on the Select [user] dialog).

Step 1 Navigate to the User Management page by either of the following methods:

- Click **Add** on the **Management > User Management** tab.
- Click **Add** on the Select (user) dialog box when creating an alert notification. See [“Configure a Rule to Send an Alert Action” section on page 22-5](#).

The User Configuration page appears, as shown in [Figure 22-6](#).

Figure 22-6 *User Configuration Page*

Role: Admin

Login: pnadmin

Password:

Re-enter password:

First Name:

Last Name:

Organization:

Email:

SMS:

Work Phone:

Home Phone:

Fax:

Pager: (Cell phone or pager number e.g: 4082345678)

Service Provider:

143791

Step 2 From the **Role** field, select a **Role** for the user.

- **Admin:** has full use of the MARS.
- **Notification Only:** for a non-user of the MARS appliance, use this to send alerts to people who are not administrators, security analysts, or operators.
- **Operator:** has read-only privileges.
- **Security Analyst:** has full use of the MARS, except cannot access the Admin tab

Step 3 Create or change the user's password if necessary.

Step 4 Enter the user's credentials and personal information, which may include any of the following:

- First name
- Last name
- Organization name
- Email address
- Short Message Service (SMS) number—for example, 8885551212@servprov.com
- Work telephone number
- Home telephone number
- FAX number
- Pager number or ID— may also be a mobile telephone number, for example, 5552345678

Step 5 If you are not creating a notification by pager, go to [Step 10](#).

Step 6 For notification by pager, you must specify a service provider (cell phone or pager company). From the Service Provider field, select **New Provider**.

This pull-down menu is populated as you add new providers.

Additional service provider information fields appear on the same page, as shown in [Figure 22-7](#).

Figure 22-7 Service Provider Fields to Add or Change a Service Provider

Step 7 In the **Provider Name** field, enter the name of the service provider.

Step 8 In the **Provider Phone No** field, enter the service provider's telephone number.

This is the number the service provider requires for accepting alpha-numeric messages using the IXO/TAP protocol. The format is like a regular phone number, such as: 18001234567. The format of 1-800-1234567 is also acceptable. If dialing "9" is required to access a number outside your private branch exchange, type a "9," before the full telephone number (for example, 9,1-800-1234567).

Step 9 In the **Provider Baudrate** field, enter the baud rate specified by the provider.

This is the baud rate the service provider requires for the specified phone number. Common values are 1200, 2400, 4800, and 9600.

Step 10 Click **Submit** to close the User Configuration page and return to the **User Management** tab.

This ends the [Create a New User—Role, Identity, Password, and Notification Information](#) procedure.

Create a Custom User Group

To create a custom user group in addition to the default groups created by MARS, complete the following procedure:

Step 1 Navigate to the **Management > User Management** tab.

Step 2 Click **Add Group**.

Step 3 In the **Name** field, enter a name for the group.

Step 4 To add users to the group, click the check box of users from the list on the right-hand area. Click **Add**.

The checked names appear in the left-hand side of the dialog box.

To remove users from the group, pick the users from the left-hand side with Ctrl+click. Click **Remove**.

The selected names appear in the right-hand side of the dialog box.

Step 5 Click **Submit**.

You are returned to the User Management tab.

This ends the [Create a Custom User Group](#) procedure.

Add a User to a Custom User Group

To include a user in a custom User Group, complete the following steps:

**Note**

The user is automatically added to the User Group that corresponds to their role. Admin, Operator, Notification, and Security Analyst are system groups and cannot be edited.

-
- Step 1** Navigate to the **Management > User Management** tab.
- Step 2** Select the User Group to edit from the **Select Group** dropdown list.
The members of the group are displayed.
- Step 3** Click **Edit Group**. The User Group dialog box appears.
- Step 4** Check the users to add to the group from the list on the right hand side. Click **Add**. The checked names move to the left-hand area of the dialog box.
- Step 5** Click **Submit**.
You are returned to the **User Management** tab.
This ends the [Add a User to a Custom User Group](#) procedure.
-



Management Tab Overview

Use the management features in the Local Controller to assign: event, addressing, service, and user information. This information is used in rules, queries, and to determine false positives.


Activating

In general, you need to activate changes in the Management tabs if the changes are part of a rule.

To activate a set of management additions or changes

Step 1 When changes (or additions) are complete, activate them by clicking **Activate**.

Figure 23-1 Clicking the Activate Button



Event Management

To open the Event Management sub-tab, click the **Management > Event Management** tabs.

On the Event Management page, you can search and filter events and event groups, and work with groups of events.

Search for an Event Description or CVE Names

You can search for partial matches of event descriptions or Common Vulnerabilities and Exposures (CVE) names.

-
- Step 1** Enter the text that you want to search for in the **Search** field.
- Step 2** Click **Search**.
-

To view a list of all currently supported CVEs

-
- Step 1** Enter CVE into the **Search** field.
- Step 2** Click **Search**.
-

Event Groups

Using and creating event groups is one of the most powerful ways to leverage rules. You can take any of the events presented here, group them, and then use them with rules to concentrate your searches for attacks.

To filter by event groups or severity

From the appropriate list, select the group or severity.

Edit a Group of Events

**Note**

You can not edit system-defined groups.

-
- Step 1** Select the group in the **Select Group** list.
- Step 2** Click **Edit Group**.
- Step 3** Click each group in the Chosen and Available fields to highlight it. Click it again to de-highlight it.
- Step 4** Click **Add** or **Remove** to move highlighted items as needed.
- Step 5** Click **Submit**.
-

Add a Group

-
- Step 1** Click **Add**.
- Step 2** In the **Name** field, enter a name for the group.
- Step 3** In the **Available** field, click each group that you want to add to highlight it. Click it again to de-highlight it.

Step 4 Click **Add**.

Step 5 Click **Submit**.

IP Management

The IP Management page, accessed by clicking **Management > IP Management**, enables the definition of network assets that you use as building blocks for inspection rules, drop rules, reports and queries, topology discovery schedules, and in defining reporting devices and mitigation devices. You can define assets as networks, IP ranges, or hosts. You can also defined named variables for use within inspection rules.

The vulnerability assessment information that you define for a host, specifically the operating system type and patch level and the known services that run on the host, assists MARS in determining false positives.



Tip

You can filter the list of objects displayed by the View list box. This selection allows you to filter to hosts, networks, IP ranges, or variables.

Search for an Address, Network, Variable, or Host

Step 1 Enter the text that you want to search for in the **Search** field.

Step 2 Click **Search**.

Filter by Groups

From the **Select Group** list, select the group.

Edit a Group

Step 1 Select **Management > IP Management**.

The IP Management page appears.

Step 2 Select the group in the **Select Group** list.

Step 3 Click **Edit Group**.

Step 4 Click each group in the **Chosen** and **Available** fields to highlight it. Click it again to de-highlight it.

Step 5 Click **Add** or **Remove** to move highlighted items as needed.

Step 6 Click **Submit**.

Add a Group

-
- Step 1** Select **Management > IP Management**.
The IP Management page appears.
- Step 2** Click **Add Group**.
- Step 3** In the **Name** field, enter a name for the group.
- Step 4** In the **Available** field, click a group to highlight it. To de-highlight an item, click it again.
- Step 5** Click **Add** to move the selected Event Type Groups into the **Chosen** field.
- Step 6** Click **Submit**.
-

Add a Network, IP Range, or Variable

-
- Step 1** Select **Management > IP Management**.
The IP Management page appears.

Figure 23-2 Add a Network, IP Range, or Variable

The screenshot shows a web form titled 'Add a Network, IP Range, or Variable'. At the top, there is a 'Type:' dropdown menu with a blue arrow pointing down. The dropdown menu is open, showing three options: 'Network' (highlighted in blue), 'IP Range', and 'Variables'. Below the dropdown, there is a large yellow horizontal bar. Underneath the bar, there are two sets of input fields. The first set is labeled 'Network IP:' and consists of four small input boxes separated by dots. The second set is labeled 'IP Mask:' and also consists of four small input boxes separated by dots. At the bottom right of the form, there are two buttons: 'Cancel' and 'Submit'. A vertical number '143375' is visible on the right side of the form.

- Step 2** Click **Add**.
- Step 3** In the **Type** list select: network, IP range, or variable.
- Step 4** For each type enter the appropriate information.
- Network: name, network IP, network mask
 - IP range: name and range
 - Variable: variable name
- Step 5** Click **Submit**.
-

Add a Host

Within MARS, a host is manually or automatically defined as the result of one of the following options:

- A reporting device or mitigation device defined under the Admin > Security and Monitoring Devices tab.
- A host managed by a reporting device defined under the Admin > Security and Monitoring Devices tab, such as a host running Cisco Security Agent and discovered by MARS when processing the logs provided by the CSA Management Console.
- An asset that you want to identify for the purpose of actively interacting with that host from the MARS system, such as third-party syslog sever to which you want to forward syslog messages using alerts.
- A host that is discovered by the system as part of topology discovery. For example, when processing the ARP cache table on a Cisco Catalyst Switch.
- A host involved in a session that, at one time or another, was considered suspicious, such as a potential target of an attack. In this case, MARS will have performed a Nessus and nmap port sweep of the host to identify whether it was likely breached.

Because of these various options, you can have a large number of hosts defined on the IP Management page in the web interface. If you do not have a vulnerability assessment package that is compatible with MARS, you should consider providing as much information as possible about these hosts. For more information, see [Define Vulnerability Assessment Information, page 10-11](#).

**Note**

If you are attempting to add a host and you are detecting a conflict with a previously defined host, see [Delete a Device, page 2-19](#) for additional troubleshooting information.

To manually add a host, follow these steps:

-
- Step 1** Select **Management > IP Management**.
The IP Management page appears.
- Step 2** Click **Add**.
- Step 3** In the Type list select **host**.

Figure 23-3 General Information for a Host

Type:

↓

General	Vulnerability Assessment Info								
→ *Device Name: → Access IP: → Operating System: Windows → NetBIOS Name: Enter interface information: Add Interface Remove Interface/IP <table> <thead> <tr> <th>Name:</th> <th>IP Address:</th> <th>Network Mask:</th> <th></th> </tr> </thead> <tbody> <tr> <td>☐ ether0</td> <td> </td> <td> </td> <td> Add IP/Network Mask </td> </tr> </tbody> </table>		Name:	IP Address:	Network Mask:		☐ ether0			Add IP/Network Mask
Name:	IP Address:	Network Mask:							
☐ ether0			Add IP/Network Mask						

143370

- Step 4** In the Name field, enter the host's name.
- Step 5** In the Access IP field, identify the address used to pull log events from this host or used to connect to when performing dynamic vulnerability assessments while investigating detected attacks.
- Step 6** If the host is running a variety of Windows, Solaris, or Linux, select the corresponding value in the Operating System field. Otherwise, verify that Generic is selected.
- Step 7** If you are running NetBIOS on your network, enter the name associated with this host.
NetBIOS provides name registration and resolution services. MARS uses this setting to provide attack path analysis and address resolution.
- Step 8** Add as many IP address and masks to the interface by clicking **Add IP/Mask**.
- Step 9** Under Enter Interface Information, enter the values for the interface name, IP address, and network mask.
- Step 10** If you have a dual-homed host, you can add additional interfaces by clicking **Add Interface**.
- Step 11** To specify vulnerability assessment information, continue with [Define Vulnerability Assessment Information](#), page 10-11.

Edit Host Information

- Step 1** Select **Management > IP Management**.
- Step 2** Check the box next to the host that you want to edit.
- Step 3** If you are editing interface or IP mask information, make your changes here and click **Submit**.

- Step 4** If you need to edit the host's properties, click **Properties**.
 - Step 5** Make changes to the operating system as necessary, and click **Next**.
 - Step 6** To make changes to service or application, remove the old service by select its radio button, and click **Delete**.
 - Step 7** Click **Add Service**, and continue with Step 3.
-

Service Management

To open the Service Management sub-tab, click the **Management > Service Management** tabs.

Service is a combination of source port, destination port and protocol. The Service Management page displays services and their descriptions, ports and protocols. On the Service Management page, you can work with the services on your networks.

Search for a Service

-
- Step 1** Enter the text that you want to search for in the **Search** field.
 - Step 2** Click **Search**.
To filter by service groups
From the appropriate list, select the group.
-

Add a Group of Services

-
- Step 1** Click **Add**.
 - Step 2** In the **Name** field, enter a name for the group.
 - Step 3** In the **Available** field, click items to select them, and click them again to de-select them.
 - Step 4** Click **Add**.
 - Step 5** Click **Submit**.
-

Edit a Group of Services



Note You can not edit system-defined groups.

-
- Step 1** Select the group in the **Select Group** list.
 - Step 2** Click **Edit Group**.

- Step 3** Click each group in the **Chosen** and **Available** fields to highlight it. Click it again to de-highlight it.
- Step 4** Click **Add** or **Remove** to move the highlighted items as needed.
- Step 5** Click **Submit**.
-

Add a Service

- Step 1** Click **Add**.
- Step 2** Enter the service's details.
- Step 3** Click **Submit**.
-

Edit a Service

- Step 1** Check the box next to the service.
- Step 2** Click **Edit**.
- Step 3** Make your changes, and click **Submit**.
-

Delete a Service

- Step 1** Check the box next to the service.
- Step 2** Click **Delete**.
- Step 3** On the confirmation page, click **Yes**.
-

User Management

The User Management page allows you to manage users and administrators of the MARS system, including the roles and groups to which those users belong. On this page, you can define new user accounts, enabling access to specific features of the web interface. You can define user-specific notification settings for the user, such as a valid e-mail address or pager number. Some system-wide settings, such as pager and cell phone service provider settings, are also accessible exclusively through this page. To access the User Management page, click either **Management > User Management** or **Admin > User Management**.

In MARS, four separate user roles exist that can be assigned to any user who needs to access the web interface:

- *Admin* has full read/write privileges. Users in this role can define new users with any desired role. Users in the role can change the password settings of the accounts in any user role.

- *Security Analyst* has full read privileges but is restricted to write for reports privileges. Users in this role can only define new users (and change passwords of users) with the Notifications Only role.
- *Operator* has read only privileges. Users in this role cannot define new users or change passwords, even of their own user account.
- *Notifications Only*. This user role has no permissions to access to the MARS web interface; use this role to identify users who will receive notifications, such as e-mail, SMS, or pager notifications.

While roles are system defined, you can define, edit, and delete user groups. For more information, see [Create a User Group, page 23-12](#) and [Add or Remove a User from a User Group, page 23-12](#).

Good security practices suggest strong passwords for use with the MARS Appliances. When defining user names and password, keep the following guidelines in mind:

Login names and passwords:

- can be alphanumeric characters
- can contain special characters (!, @, #, etc.)
- *cannot* contain single or double quotes ('or ")
- are case sensitive

Login names can have up to 20 characters. Passwords can have up to 64 characters.

Add a New User

Defining a new user involves specifying the user name, password, role, contact information, and notification information.

To add a new user, follow these steps:

-
- Step 1** From the **Management > User Management** tab, click **Add**. The User Configuration page appears, as shown in [Figure 23-4](#).

Figure 23-4 User Configuration Page

Role: Admin

Login: padmin

Password:

Re-enter password:

First Name:

Last Name:

Organization:

Email:

SMS:

Work Phone:

Home Phone:

Fax:

Pager: (Cell phone or pager number e.g: 4082345678)

Service Provider:

Step 2 From the **Role** field, select a **Role** for the user.

- **Admin:** has full use of Local Controller.
- **Notification Only:** for a non-user of the Local Controller appliance, use this to send alerts to people who are not admins, security analysts, or operators.
- **Operator:** has read-only privileges.
- **Security Analyst:** has full use of Local Controller, except cannot access the Admin tab

Step 3 Create or change the user's password if necessary.

Step 4 Enter the user's credentials and personal information.

The information can include the following:

- First name
- Last name
- Organization name
- Email address
- Short Message Service (SMS) number—for example, 8885551212@servprov.com
- Work telephone number
- Home telephone number
- FAX number
- Pager number— may also be a mobile telephone number, for example, 5552345678

- Step 5** If you are creating a notification by pager, go to the next section, “[Add a Service Provider \(Cell phone/Pager\)](#)”, otherwise click **Submit** to complete the procedure for adding a user.

Add a Service Provider (Cell phone/Pager)

When configuring a notification by pager, add a service provider (cell phone or pager company) by completing the following procedure:

- Step 1** From the **Service Provider** field, select **New Provider**. Additional fields appear, as shown in [Figure 23-5](#).

The pull-down menu is populated as you add new service providers.

Figure 23-5 Select a New Provider and Provide Contact Details

- Step 2** In the **Provider Name** field, enter the name of the service provider.
- Step 3** In the **Provider Phone No** field, enter the service provider’s telephone number.
- This is the number the service provider uses for accepting alpha-numeric messages using the IXO/TAP protocol. The format is like a regular phone number, such as: 18001234567. The format of 1-800-1234567 is also acceptable. If dialing “9” is required to access a number outside your private branch exchange, type a “9,” before the full telephone number (for example, 9,1-800-1234567).
- Step 4** In the **Provider Baudrate** field, enter the baud rate specified by the provider.
- This is the baud rate the service provider requires for the specified phone number. Common values are 1200, 2400, 4800, and 9600.
- Consult your service provider’s website for more information on their baud rates.
- Step 5** Click **Submit** to close the User Configuration page and return to the **User Management** tab.

Search for a User

- Step 1** Enter the text that you want to search for in the **Search** field.
- Step 2** Click **Search**.

Edit or Remove a User

-
- Step 1** Form the **Management User tab**, check the box next to the user's name.
 - Step 2** Click **Delete** to delete the user.
 - Step 3** Click **Edit** to change the user's configuration information.
The User Configuration page appears.
 - Step 4** Edit the User Configuration page.
 - Step 5** Click **Submit**.
-

Create a User Group

-
- Step 1** Click **Add Group**.
 - Step 2** In the **Name** field, enter a name for the group.
 - Step 3** To add to the group, check the users from the list on the right hand side. Click **Add**.
The checked names move to the lefthand side of the dialog box.
 - Step 4** To remove users from the group, select the users from the left hand side with Ctrl+click . Click **Remove**.
The selected names move to the righthand side of the dialog box.
 - Step 5** Click **Submit**.
-

Add or Remove a User from a User Group

To add or remove a user from a custom User Group, do the following steps:

**Note**

Admin, Operator, Notification, and Security Analyst are system groups and cannot be edited. The user is automatically added to the User Group that corresponds to their role.

-
- Step 1** Select the User Group from the **Select Group** field. The members of the group are displayed.
 - Step 2** Click **Edit Group**. The User Group dialog box appears.
 - Step 3** To add to the group, check the users from the list on the right hand side. Click **Add**.
The checked names move to the lefthand side of the dialog box.
 - Step 4** To remove users from the group, select the users from the left hand side with Ctrl+click . Click **Remove**.
The selected names move to the righthand side of the dialog box.
 - Step 5** Click **Submit**. You are returned to the **User Management** tab.
-

Filter by Groups

From the **Select Group** list, select the group. Only the members of the group are displayed.



System Maintenance

Much of the system maintenance information for the MARS Appliance is provided exclusively in the *Install and Setup Guide for Cisco Security Monitoring, Analysis, and Response System*.

The MARS Appliance requires little maintenance. To perform maintenance tasks, you can use the CLI or the web interface as needed. Some hardware maintenance tasks require physical access to the MARS Appliance.

This chapter contains the following sections:

- [Setting Runtime Logging Levels, page 24-1](#)
- [Viewing the Appliance's Log Files, page 24-2](#)
- [Viewing the Audit Trail, page 24-3](#)
- [Retrieving Raw Messages, page 24-3](#)
- [Hard Drives, page 24-7](#)
- [Replacing the Lithium Cell CMOS Battery, page 24-8](#)
- [Change the Default Password of the Administrator Account, page 24-8](#)

For information about upgrading, backing up, and restoring data on the MARS Appliance, see the following sections of the *Install and Setup Guide for Cisco Security Monitoring, Analysis, and Response System*:

- [Performing Command Line Administration Tasks, page 6-1](#)
- [Checklist for Upgrading the Appliance Software, page 6-7](#)
- [Configuring and Performing Appliance Data Backups, page 6-18](#)
- [Recovery Management, page 6-27](#)

Setting Runtime Logging Levels

To set the appliance's runtime logging levels, navigate to **Admin > System Maintenance > Set Runtime Logging Levels**. For typical use, it is best to leave this page set to its defaults.

When you have made your selections, click the **Change Logging Levels** button.

The following log levels are available:

- **Fatal.** Enables fatal logging messages. Fatal messages record very severe error events that will likely lead the application to abort.

- **Error.** Enables error and fatal logging messages. Error messages record error events that might still allow the application to continue running.
- **Warn.** Enables warning, error, and fatal logging messages. Warning messages record potentially harmful situations.
- **Info.** Enables informational, warning, error, and fatal logging messages. Informational messages highlight the progress of the application at coarse-grained level.
- **Debug.** Enables debug, informational, warning, error, and fatal logging messages. Debug messages record fine-grained informational events that are most useful to debug an application.
- **Trace.** Enables trace, debug, information, warning, error, and fatal logging messages. Trace messages record finer-grained informational events than debug messages.

Viewing the Appliance's Log Files

To view the appliance's log files or to change their levels or source, navigate to **Admin > System Maintenance > View Log Files**.

Figure 24-1 Back-end log viewing options

View Backend Log

You can view the appliance's back-end logs either by selecting a number of days, hours, and minutes or you can view logs by selecting a start and ending date and time.

You can select the levels of logs that you want. Your choices are: All, Fatal, Error, Warn, Info, and Debug.

You can also choose the source of the files that you want to view. Select either Backend or GUI.

View the Back-end Log

-
- Step 1** Click the appropriate radio button:
- **Last:** The present time minus the number of days, hours, and minutes entered.
 - **Start/End:** Absolute literal time ranges defined by the date to the minute.
- Step 2** Select user, group, etc.
- Step 3** Select the source.
- Step 4** Click **Submit**.
-

Viewing the Audit Trail

You can track the activities of the appliance's users by analyzing the appliance's log files. To set the appliance's audit trail logs, navigate to **Admin > System Maintenance > View Audit Trail**. For typical use, it is best to leave this page set to its defaults.

You can view the user audit trails either by selecting a number of days, hours, and minutes, or you can view a specific interval by selecting a start and ending date and time.

View an Audit Trail

-
Step 1	Click the appropriate radio button: • Last: DD-HH-MM • Start/End: YY-MM-DD-HH-MM
Step 2	From the list, select the user or user group.
Step 3	Click Submit .
-

Retrieving Raw Messages

Using this feature, you can retrieve raw messages from either an archive server (see [Configuring and Performing Appliance Data Backups, page 6-18](#)) or from the database running on the Local Controller. These two methods offer different advantages:

- **Archive server.** Retrieving raw messages, or event data, from an archive server is much faster than retrieving from the database. Therefore, it is the recommended option if it is available and it covers the time period you are investigating. However, this option is only available if you have enabled data archiving and waited the requisite time for the initial archival operation to occur; it is a scheduled operation that runs nightly around 2:00 a.m. Once the initial archive is performed, the event data is written to the archive server every hour. That data is not archived in real-time identifies another limitation to this option, and that is the historical period that can be studied. If you need to view data that is more current than an hour old, you should select the Database option to ensure that correct data is retrieved. For all other periods, the archive server option is recommended. To enable archiving, see [Configuring and Performing Appliance Data Backups, page 6-18](#).
- **Database.** Retrieving event data from the database provides slower performance than the archive server. However, it provides access to the most current data received. When you select this option, you can specify where you want the retrieved records to be written: in the default local directory or the a remote server, if one is available.

This section contains the following topics:

- [Retrieve Raw Messages From Archive Server, page 24-3](#)
- [Retrieve Raw Messages From the Database of a Local Controller, page 24-5](#)

Retrieve Raw Messages From Archive Server

Use this selection if archiving is enabled.

To retrieve event data from an archive server, follow these steps:

Step 1 Click **Admin > System Maintenance > Retrieve Raw Messages**.

Retrieve Raw Messages:

Specify Time Range:

Start: 2005 October 7 7 Hrs 12 Mins 15 Secs
 End: 2005 October 7 7 Hrs 22 Mins 15 Secs

☒ Retrieve Data From Archived Files

☐ Retrieve Data From DB

☐ Save To Local ☐ Save To Remote

☒ Force Generate Files Maximum No. of Files: 10

Select Reporting Device:

All Devices

143783

Step 2 Specify the time range by specifying values in the Start and End fields.

Step 3 Verify that **Retrieve Data From Archived Files** is selected.

The data will be retrieved from the server identified under Admin > System Maintenance > Data Archiving.

Step 4 Click **Submit**.



Note

While MARS is generating your files, you can still use the system for other tasks.

Result: The Retrieving Progress 0% screen appears. When the operation is complete, the Raw Message Files screen appears, identifying a new Gzip archive file with a filename based on specified time range.

Raw Message Files	Download
2005-10-07-06-14-28_2005-10-08-06-24-28.gz	Click Here to Download

143797

Step 5 To download and view the generated raw message file, click Click Here to Download next to the filename.

The filename adheres to the following syntax:
 YYYY-MM-DD-HH-MM-SS_YYYY-MM-DD-HH-MM-SS.gz.

Step 6 Use WinZip or another archive expansion program to extract the contents of the Gzip archive file.

Step 7 Once the textfile is extracted from the GNU Zip archive format, its contents resemble the following:

```
33750>Wed Jul 27 16:16:06 PDT 2005>BR-FW-1>10.4.1.1 Mon Jan 6 11:05:34 2003 <134>Jan 06
2003 11:03:53: %PIX-6-302001: Built inbound TCP connection 21000 for faddr 10.1.2.4/9000
gaddr 10.1.5.20/80 laddr 10.1.5.20/80
```

where it reads: *device ID>>date>>device name>>raw message*.

**Note**

If you see Chinese or other unfamiliar characters in the resulting text file, please use Microsoft Internet Explorer to view the file and verify that the Western European ISO or Western European Windows encoding value is selected (View > Encoding). The “»” sign appears correctly as a separator when a compatible encoding is selected.

Retrieve Raw Messages From the Database of a Local Controller

Use this selection if archiving is not enabled or if you need to view event data that was received within the past hour.

To retrieve event data from the database, follow these steps:

Step 1 Click **Admin > System Maintenance > Retrieve Raw Messages**.

Retrieve Raw Messages:

Specify Time Range:

Start: 2005 October 7 7 Hrs 12 Mins 15 Secs
 End: 2005 October 7 7 Hrs 22 Mins 15 Secs

☐ Retrieve Data From Archived Files

☒ Retrieve Data From DB

☐ Save To Local ☒ Save To Remote

☒ Force Generate Files Maximum No. of Files: 10

Select Reporting Device:

All Devices

143784

Step 2 Specify the time range by specifying values in the Start and End fields.

Step 3 Select **Retrieve from Database**.

Step 4 Select one of the following options:

- **Save to Local.** This option retrieves the data from the database and stores it on the local appliance.
- **Save to Remote.** This option retrieves the data from the database and stores it on the archive server, as identified under Admin > System Maintenance > Data Archiving.

Step 5 Review the Cached Files time range information, and then do one of the following:

- If you want data from within this time range, you do not need for Force Generate Files.
- If you want data that does not fall within the Cached Files time range, select the **Force Generate Files** check box.
- If there is no cached file information, select the **Force Generate Files** check box.

If no cached file data is shown, then no previous queries have been performed and stored. For example, if you preform three separate queries, using time range A, from the database using the time range, saving the files to the local MARS Appliance. If you later specify the same time range A and do the retrieval

again but you do not clear the Force generate files check box, the system performs the query, generating the file again. However, if you have already retrieved and stored some data before, you can specify to retrieve them from those saved files by clearing the Force generate files check box.

Step 6 Enter the maximum number of retrieved files to retain in the Maximum No. of Files field.

This value refers to the maximum number of event files to be generated for this query.



Note Requesting large numbers of files can take some time.

Step 7 Select the list of devices for which you want to pull event data in the Reporting Devices list.


You can select a specific device by name or All Devices.

Step 8 Click **Submit**.



Note While MARS is generating your files, you can still use the system for other tasks.

Result: The Retrieving Progress 0% screen appears. When the operation is complete, the Raw Message Files screen appears, identifying a new Gzip archive file with a filename based on specified time range.



Step 9 To download and view the generated raw message file, click Click Here to Download next to the filename.

The filename adheres to the following syntax:

YYYY-MM-DD-HH-MM-SS_YYYY-MM-DD-HH-MM-SS.gz.

Step 10 Use WinZip or another archive expansion program to extract the contents of the Gzip archive file.

Step 11 Once the textfile is extracted from the GNU Zip archive format, its contents resemble the following:

```
33750>Wed Jul 27 16:16:06 PDT 2005>BR-FW-1>10.4.1.1 Mon Jan 6 11:05:34 2003 <134>Jan 06
2003 11:03:53: %PIX-6-302001: Built inbound TCP connection 21000 for faddr 10.1.2.4/9000
gaddr 10.1.5.20/80 laddr 10.1.5.20/80
```

where it reads: *device ID>>date>>device name>>raw message*.



Note If you see Chinese or other unfamiliar characters in the resulting text file, please use Microsoft Internet Explorer to view the file and verify that the Western European ISO or Western European Windows encoding value is selected (View > Encoding). The “»” sign appears correctly as a separator when a compatible encoding is selected.

Hard Drives

Status Lights

Depending on the model of the appliance, each hard drive has two status lights under or next to the drive. The following states can be determined based on the status lights:

- A steady green light indicates that the drive is functioning normally.
- A blinking orange light indicates that the drive is performing I/O operations.
- No light indicates that the disk has no power.

Partition Checking

The appliance automatically runs checks on different partitions of the hard drive after the system has been re-booted 25 - 30 times, or if the appliance has not been re-booted in 180 days.

Hotswapping Hard Drives

If a hard drive fails in the MARS 50, 100, 100e, 200, GC, or GCm appliance models, the MARS administrator receives an e-mail notification. The notification identifies the drive number of the failed hard drive.

Remove a Hard Drive

-
Step 1	Log in through the CLI tool.
Step 2	Enter the hotswap command.
Step 3	Remove the hard drive.
Step 4	Replace a Hard Drive, page 24-7
-

Replace a Hard Drive

-
Step 1	Use the door key to open the chassis door.
Step 2	Use the drive bay key to unlock the drive bay you want to change.
Step 3	Pull out the drive.
Step 4	Use a screwdriver to remove the drive bay holder from the hard drive.
Step 5	Put the new hard drive in the drive bay holder, screwing it into place.
Step 6	Gently push the drive into the place you removed it from.
Step 7	Lock the drive bay back into place.
Step 8	Close the bay door, and re-lock it.

**Note**

You can only swap one hard drive at a time. Make sure to verify the system has completed the initializing of the new hard drive before swapping another hard drive.

Replacing the Lithium Cell CMOS Battery

**Caution**

A risk of explosion exists if you replace the lithium cell cmos battery with the incorrect type. Never try to replace the Lithium Cell CMOS battery. If this battery needs replacement, contact Cisco for more information.

Replace the Lithium Cell CMOS Battery

**Note**

Take proper electrostatic discharge (ESD) measures before physically touching the appliance.

If the CMOS battery needs replacement, follow these steps:

- Step 1** Turn the appliance's power off.
- Step 2** Unplug the appliance from the wall electrical socket.
- Step 3** Locate the lithium cell CMOS battery.
- Step 4** Remove it.
- Step 5** Set the new battery in its place.
- Step 6** Plug the appliance into the electrical socket in the wall.
- Step 7** Turn the appliance's power on.

**Note**

The lithium battery can be harmful to the environment. Contact your local waste disposal service for information about safely disposing of the battery.

Change the Default Password of the Administrator Account

Good security practices require that you change the default password. We recommend using strong passwords for the MARS Appliance appliances.

Login names and passwords:

- can be alphanumeric characters
- are case sensitive

- can contain special characters (!, @, #, etc.)
- **cannot** contain single or double quotes ('or ")

Login names can contain up to 20 characters. Passwords can contain up to 64 characters.

To change the default password and setup administrator notification, follow these steps:

-
Step 1	Click the Management > User Management tab.
Step 2	Check the box next to Administrator, and click Edit .
Step 3	Enter the new Administrator password and the Administrator e-mail address.
Step 4	Click Submit .
-

Change the Default Password of the Administrator Account



Cisco Security MARS XML API Reference

This appendix provides resources for creating XML applications that integrate Cisco Security MARS XML data into third-party applications.

XML Overview

The XML schema are written in conformance with the standard World Wide Web Consortium (W3C) XML schema language. A schema by definition, describes all data and data structures required to create your application. Many XML development environments provide enough capability to view the schema in a way that you can identify all components, their relationships, constraints, attributes, annotations, and usage guidelines at a glance. Some applications generate hyperlinked reference documentation. By providing sufficient documentation and annotation tags within the schemas, Cisco supports such documentation generating applications.

[Table A-1](#) lists resources for XML development. The list is not exhaustive nor an endorsement by Cisco of any particular product.

Table A-1 XML Resources

Resource Description	URL
Latest Cisco Security MARS Schemas	http://www.cisco.com/cgi-bin/tablebuild.pl/cs-mars
W3C XML Schema standards forum with resource links	http://www.w3.org/XML/Schema
General XML description with resource links	http://en.wikipedia.org/wiki/XML
Online XML Tutorials	http://www.w3schools.com/xml/default.asp
XML Documentation Generators (generates hyperlinked command references from any schema)	http://lists.w3.org/Archives/Public/xmlschema-dev/2006Feb/0050.html http://www.stylusstudio.com/xml_schema_doc_gen.html
XMLSpy® XML development software	http://www.altova.com/products.html

XML Incident Notification Data File and Schema

XML incident notification sends an email notification of an incident with an attached XML data file. The XML data file contains all incident details that can be viewed on the GUI except for Path/Mitigation data. The XML data file can be sent as a plain-text file or as a compressed gzip file. The filename is constructed with the incident ID number, for example `CS-MARS-Incident-13725095.xml`. The compressed version of the same data file would be `CS-MARS-Incident-13725095.xml.gz`.

An XML application can be written to parse and extract data from the XML incident notification data file for integration into third-party software, such as a trouble ticketing system, or helpdesk software.

[Table A-2](#) lists the documentation for the Cisco Security MARS XML incident notification feature.

Table A-2 *Related XML Incident Notification Documents*

Resource Description	Resource Location
Configuring XML incident notification on MARS	Chapter 22, “Sending Alerts and Incident Notifications”
A ZIP file containing the XML incident notification schema	TBD
A hyper-linked component reference, generated from the XML incident notification schema	TBD
Sample XML incident notification data generated by MARS	Appendix A, “Example A-1”

XML Incident Notification Data File Sample Output

[Example A-1](#) is XML incident notification data generated by the events that trigger the rule “CS-MARS Database Partition Usage.”

Example A-1 *XML Incident Notification Data File Contents*

```
<?xml version="1.0" encoding="UTF-8"?>
<CSMARS-NOTIFICATION>
  <Header>
    <Version>1.0</Version>
    <GenTimeStamp>May 15, 2006 8:48:02 AM PDT</GenTimeStamp>
    <CSMARSHostIpAddr_eth0>10.2.3.7</CSMARSHostIpAddr_eth0>
    <CSMARSHostIpAddr_eth1>192.168.1.101</CSMARSHostIpAddr_eth1>
    <CSMARSHostName>MyLatest</CSMARSHostName>
    <CSMARSThreatName />
    <CSMARSThreatVersion>4.2.1</CSMARSThreatVersion>
  </Header>
  <Data>
    <Incident id="597842933">
      <StartTime>May 15, 2006 8:47:26 AM PDT</StartTime>
      <EndTime>May 15, 2006 8:47:26 AM PDT</EndTime>
      <Severity>LOW</Severity>
      <Session id="597744001">
        <Instance>0</Instance>
        <RuleOffset>1</RuleOffset>
        <SessionEndPoints>
          <Source ipaddress="0.0.0.0" />
          <Destination ipaddress="10.2.3.7" />
          <SourcePort>0</SourcePort>
        </SessionEndPoints>
      </Session>
    </Incident>
  </Data>
</CSMARS-NOTIFICATION>
```

```

        <DestinationPort>0</DestinationPort>
        <Protocol>-1</Protocol>
    </SessionEndPoints>
    <Event id="597744001">
        <EventType id="125755" />
        <TimeStamp>May 15, 2006 8:47:26 AM PDT</TimeStamp>
        <ReportingDevice id="50" />
        <RawMessage>Mon May 15 08:47:26 PDT 2006 &lt;13&gt;%MARS-3-100026 CS-MARS
MyLatest : Current database partition pn_event_session_8 utilization has reached 75%;
next database partition pn_event_session_9 containing data between Thu Apr 20 11:59:13 PDT
2006 and Fri Apr 21 11:32:17 PDT 2006 will be purged approximately at Mon May 15 11:56:02
PDT 2006.</RawMessage>
        <FalsePositiveType>NOT_AVAILABLE</FalsePositiveType>
    <EventEndPoints>
        <Source ipaddress="0.0.0.0" />
        <Destination ipaddress="10.2.3.7" />
        <SourcePort>0</SourcePort>
        <DestinationPort>0</DestinationPort>
        <Protocol>-1</Protocol>
    </EventEndPoints>
    <NATtedEndPoints>
        <Source ipaddress="0.0.0.0" />
        <Destination ipaddress="10.2.3.7" />
        <SourcePort>0</SourcePort>
        <DestinationPort>0</DestinationPort>
        <Protocol>-1</Protocol>
    </NATtedEndPoints>
    <FiringEventFlag>true</FiringEventFlag>
</Event>
</Session>
<Rule id="134473">
    <Name>System Rule: CS-MARS Database Partition Usage</Name>
    <Description>This rule indicates that the current CS-MARS database partition
filled up to 75% of its capacity and the next database partition will be purged soon to
create space for new events. The estimated purge times are in the event message. This is
normal CS-MARS activity and will result in old events and incidents to purged from CS-MARS
database. Users are urged to archive CS-MARS data to prevent permanent data
loss.</Description>
</Rule>
<NetworkAddressObj id="0">
    <IPAddress>0.0.0.0</IPAddress>
    <MAC />
    <DNSName />
    <DynamicInfo>
        <HostName />
        <MACAddress />
        <AAAUUser />
        <EnforcementDeviceAndPort />
        <ReportingDevice />
        <StartTime>Dec 31, 1969 4:00:00 PM PST</StartTime>
        <EndTime>Dec 31, 1969 4:00:00 PM PST</EndTime>
        <UpdateTime>Dec 31, 1969 4:00:00 PM PST</UpdateTime>
    </DynamicInfo>
</NetworkAddressObj>
<NetworkAddressObj id="167904007">
    <IPAddress>10.2.3.7</IPAddress>
    <MAC>
        <MACAddress>00:30:48:83:25:d9</MACAddress>
        <LastUpdateTime>May 15, 2006 6:59:09 AM PDT</LastUpdateTime>
    </MAC>
    <DNSName>MyLatest</DNSName>
    <Device id="50" />
    <DynamicInfo>
        <HostName />

```

```

        <MACAddress />
        <AAAUUser />
        <EnforcementDeviceAndPort />
        <ReportingDevice />
        <StartTime>Dec 31, 1969 4:00:00 PM PST</StartTime>
        <EndTime>Dec 31, 1969 4:00:00 PM PST</EndTime>
        <UpdateTime>Dec 31, 1969 4:00:00 PM PST</UpdateTime>
    </DynamicInfo>
</NetworkAddressObj>
<EventTypeObj id="125755">
    <Name>1000029</Name>
    <Description>CS-MARS DB partition filling up causing the next partition to be
purged soon</Description>
    <Severity>LOW</Severity>
    <CVE />
</EventTypeObj>
<DeviceObj id="50">
    <Name>MyLatest</Name>
    <NetBiosName />
    <DefaultGateway>10.2.3.1</DefaultGateway>
    <OperatingSystem id="0" />
    <InterfaceAddressObj id="117924">
        <Name>eth0</Name>
        <IPAddress>10.2.3.7</IPAddress>
        <MAC>
            <MACAddress>00:30:48:83:25:d9</MACAddress>
            <LastUpdateTime>May 15, 2006 6:59:09 AM PDT</LastUpdateTime>
        </MAC>
    </InterfaceAddressObj>
    <InterfaceAddressObj id="123040">
        <Name>eth1</Name>
        <IPAddress>192.168.1.101</IPAddress>
        <MAC />
    </InterfaceAddressObj>
</DeviceObj>
</Incident>
</Data>
</CSMARS-NOTIFICATION>

```

XML Incident Notification Schema

The XML incident notification schema document (csmars-incident-notification.xsd) can be downloaded from the the following URL:

<http://www.cisco.com/TBD>

Usage Guidelines and Conventions for XML Incident Notification

All XML incident notification elements are defined in the XML incident notification schema. A component reference document generated from the schema is available for your convenience at the following URL:

[URL TBD](#)

You can generate a similar document with the application of your choice, or view components, their relationships, constraints, attributes, annotations, and usage guidelines within your XML development environment.

MARS uses a best effort approach to create XML incident notification data. If an error occurs during data compilation, MARS does not stop the process, but sends the data, even if it is partial. Validating the data file against the schema would result in errors for these cases.

The following conventions are observed for XML incident notification data:

- Character encoding is Unicode Transformation Format 8 (UTF-8)
- The reported time zone would be the time zone of the local controller reporting the incident
- Raw messages from reporting devices are XML-escaped in the data file. Your XML parser should be able to unescape XML data.
- If there is no value for an element available from MARS, the element is included in the data file as an empty node. For instance, a DNS name may not be available for a device.
- All date formats are **Mmm dd, yyyy hh:mm:ss AM TZD**
 - **Mmm** is the month (Jan, Feb, Mar. . . Dec)
 - **dd** is the day (1–9, 10–31)
 - **yyyy** is the year (0000–9999)
 - **hh:mm:ss** is hours, minutes, seconds
 - hh** are 1–9, 10–12
 - mm** are 00–60
 - ss** are 00–60
 - **AM** or **PM**
 - **TZD** is time zone designator (PDT, PST, MDT, MST, etc.)



Regular Expression Reference

- [PCRE Regular Expression Details, page B-1](#)
- [Backslash, page B-2](#)
- [Circumflex and Dollar, page B-7](#)
- [Full Stop \(Period, Dot\), page B-8](#)
- [Matching a Single Byte, page B-8](#)
- [Square Brackets and Character Classes, page B-8](#)
- [Posix Character Classes, page B-9](#)
- [Vertical Bar, page B-10](#)
- [Internal Option Setting, page B-10](#)
- [Subpatterns, page B-11](#)
- [Named Subpatterns, page B-12](#)
- [Repetition, page B-12](#)
- [Atomic Grouping and Possessive Quantifiers, page B-14](#)
- [Back References, page B-15](#)
- [Assertions, page B-16](#)
- [Conditional Subpatterns, page B-19](#)
- [Comments, page B-20](#)
- [Recursive Patterns, page B-20](#)
- [Subpatterns as Subroutines, page B-21](#)
- [Callouts, page B-22](#)

PCRE Regular Expression Details

The syntax and semantics of the regular expressions supported by PCRE are described below. Regular expressions are also described in the Perl documentation and in a number of books, some of which have copious examples. Jeffrey Friedl's "Mastering Regular Expressions", published by O'Reilly, covers regular expressions in great detail. This description of PCRE's regular expressions is intended as reference material.

The original operation of PCRE was on strings of one-byte characters. However, there is now also support for UTF-8 character strings. To use this, you must build PCRE to include UTF-8 support, and then call `pcre_compile()` with the `PCRE_UTF8` option. How this affects pattern matching is mentioned in several places below. There is also a summary of UTF-8 features in the section on UTF-8 support in the main PCRE page.

A regular expression is a pattern that is matched against a subject string from left to right. Most characters stand for themselves in a pattern, and match the corresponding characters in the subject. As a trivial example, the pattern

```
The quick brown fox
```

matches a portion of a subject string that is identical to itself. The power of regular expressions comes from the ability to include alternatives and repetitions in the pattern. These are encoded in the pattern by the use of *metacharacters*, which do not stand for themselves but instead are interpreted in some special way.

There are two different sets of metacharacters: those that are recognized anywhere in the pattern except within square brackets, and those that are recognized in square brackets. Outside square brackets, the metacharacters are as follows:

```
\      general escape character with several uses
^      assert start of string (or line, in multiline mode)
$      assert end of string (or line, in multiline mode)
.      match any character except newline (by default)
[      start character class definition
|      start of alternative branch
(      start subpattern
)      end subpattern
?      extends the meaning of (
        also 0 or 1 quantifier
        also quantifier minimizer
*      0 or more quantifier
+      1 or more quantifier
        also "possessive quantifier"
{      start min/max quantifier
```

Part of a pattern that is in square brackets is called a "character class". In a character class the only metacharacters are:

```
\      general escape character
^      negate the class, but only if the first character
-      indicates character range
[      POSIX character class (only if followed by POSIX syntax)
]      terminates the character class
```

The following sections describe the use of each of the metacharacters.

Backslash

The backslash character has several uses. Firstly, if it is followed by a non-alphanumeric character, it takes away any special meaning that character may have. This use of backslash as an escape character applies both inside and outside character classes.

For example, if you want to match a `*` character, you write `\*` in the pattern. This escaping action applies whether or not the following character would otherwise be interpreted as a metacharacter, so it is always safe to precede a non-alphanumeric with backslash to specify that it stands for itself. In particular, if you want to match a backslash, you write `\\`.

If a pattern is compiled with the `PCRE_EXTENDED` option, whitespace in the pattern (other than in a character class) and characters between a `#` outside a character class and the next newline character are ignored. An escaping backslash can be used to include a whitespace or `#` character as part of the pattern.

If you want to remove the special meaning from a sequence of characters, you can do so by putting them between `\Q` and `\E`. This is different from Perl in that `$` and `@` are handled as literals in `\Q...\E` sequences in PCRE, whereas in Perl, `$` and `@` cause variable interpolation. Note the following examples:

Pattern	PCRE matches	Perl matches
<code>\Qabc\$xyz\E</code>	<code>abc\$xyz</code>	<code>abc</code> followed by the contents of <code>\$xyz</code>
<code>\Qabc\ \$xyz\E</code>	<code>abc\ \$xyz</code>	<code>abc\ \$xyz</code>
<code>\Qabc\E\ \$\Qxyz\E</code>	<code>abc\$xyz</code>	<code>abc\$xyz</code>

The `\Q...\E` sequence is recognized both inside and outside character classes.

Non-printing Characters

A second use of backslash provides a way of encoding non-printing characters in patterns in a visible manner. There is no restriction on the appearance of non-printing characters, apart from the binary zero that terminates a pattern, but when a pattern is being prepared by text editing, it is usually easier to use one of the following escape sequences than the binary character it represents:

<code>\a</code>	alarm, that is, the BEL character (hex 07)
<code>\cx</code>	"control-x", where x is any character
<code>\e</code>	escape (hex 1B)
<code>\f</code>	formfeed (hex 0C)
<code>\n</code>	newline (hex 0A)
<code>\r</code>	carriage return (hex 0D)
<code>\t</code>	tab (hex 09)
<code>\ddd</code>	character with octal code ddd, or backreference
<code>\xhh</code>	character with hex code hh
<code>\x{hhh...}</code>	character with hex code hhh... (UTF-8 mode only)

The precise effect of `\cx` is as follows: if x is a lower case letter, it is converted to upper case. Then bit 6 of the character (hex 40) is inverted. Thus `\cz` becomes hex 1A, but `\c{` becomes hex 3B, while `\c;` becomes hex 7B.

After `\x`, from zero to two hexadecimal digits are read (letters can be in upper or lower case). In UTF-8 mode, any number of hexadecimal digits may appear between `\x{` and `}`, but the value of the character code must be less than 2^{31} (that is, the maximum hexadecimal value is 7FFFFFFF). If characters other than hexadecimal digits appear between `\x{` and `}`, or if there is no terminating `}`, this form of escape is not recognized. Instead, the initial `\x` will be interpreted as a basic hexadecimal escape, with no following digits, giving a character whose value is zero.

Characters whose value is less than 256 can be defined by either of the two syntaxes for `\x` when PCRE is in UTF-8 mode. There is no difference in the way they are handled. For example, `\xdc` is exactly the same as `\x{dc}`.

After `\0` up to two further octal digits are read. In both cases, if there are fewer than two digits, just those that are present are used. Thus the sequence `\0x\07` specifies two binary zeros followed by a BEL character (code value 7). Make sure you supply two digits after the initial zero if the pattern character that follows is itself an octal digit.

The handling of a backslash followed by a digit other than 0 is complicated. Outside a character class, PCRE reads it and any following digits as a decimal number. If the number is less than 10, or if there have been at least that many previous capturing left parentheses in the expression, the entire sequence is taken as a back reference. A description of how this works is given later, following the discussion of parenthesized subpatterns.

Inside a character class, or if the decimal number is greater than 9 and there have not been that many capturing subpatterns, PCRE re-reads up to three octal digits following the backslash, and generates a single byte from the least significant 8 bits of the value. Any subsequent digits stand for themselves. For example:

```
\040  is another way of writing a space
\40   is the same, provided there are fewer than 40 previous capturing subpatterns
\7    is always a back reference
\11   might be a back reference, or another way of writing a tab
\011  is always a tab
\0113 is a tab followed by the character "3"
\113  might be a back reference, otherwise the character with octal code 113
\377  might be a back reference, otherwise the byte consisting entirely of 1 bits
\81   is either a back reference, or a binary zero followed by the two characters
"8" and "1"
```

Note that octal values of 100 or greater must not be introduced by a leading zero, because no more than three octal digits are ever read.

All the sequences that define a single byte value or a single UTF-8 character (in UTF-8 mode) can be used both inside and outside character classes. In addition, inside a character class, the sequence `\b` is interpreted as the backspace character (hex 08), and the sequence `\X` is interpreted as the character "X". Outside a character class, these sequences have different meanings (see [Unicode Character Properties](#), page B-5).

Generic Character Types

The third use of backslash is for specifying generic character types. The following are always recognized:

```
\d    any decimal digit
\D    any character that is not a decimal digit
\s    any whitespace character
\S    any character that is not a whitespace character
\w    any "word" character
\W    any "non-word" character
```

Each pair of escape sequences partitions the complete set of characters into two disjoint sets. Any given character matches one, and only one, of each pair.

These character type sequences can appear both inside and outside character classes. They each match one character of the appropriate type. If the current matching point is at the end of the subject string, all of them fail, since there is no character to match.

For compatibility with Perl, `\s` does not match the VT character (code 11). This makes it different from the POSIX "space" class. The `\s` characters are HT (9), LF (10), FF (12), CR (13), and space (32).

A "word" character is an underscore or any character less than 256 that is a letter or digit. The definition of letters and digits is controlled by PCRE's low-valued character tables, and may vary if locale-specific matching is taking place (see "Locale support" in the **pcreapi** page). For example, in the "fr_FR" (French) locale, some character codes greater than 128 are used for accented letters, and these are matched by `\w`.

In UTF-8 mode, characters with values greater than 128 never match `\d`, `\s`, or `\w`, and always match `\D`, `\S`, and `\W`. This is true even when Unicode character property support is available.

Unicode Character Properties

When PCRE is built with Unicode character property support, three additional escape sequences to match generic character types are available when UTF-8 mode is selected. They are:

```
\p{xx}  a character with the xx property
\P{xx}  a character without the xx property
\X      an extended Unicode sequence
```

The property names represented by `xx` above are limited to the Unicode general category properties. Each character has exactly one such property, specified by a two-letter abbreviation. For compatibility with Perl, negation can be specified by including a circumflex between the opening brace and the property name. For example, `\p{^Lu}` is the same as `\P{Lu}`.

If only one letter is specified with `\p` or `\P`, it includes all the properties that start with that letter. In this case, in the absence of negation, the curly brackets in the escape sequence are optional; these two examples have the same effect:

```
\p{L}
\pL
```

The following property codes are supported:

```
C      Other
Cc     Control
Cf     Format
Cn     Unassigned
Co     Private use
Cs     Surrogate

L      Letter
Ll     Lower case letter
Lm     Modifier letter
Lo     Other letter
Lt     Title case letter
Lu     Upper case letter

M      Mark
Mc     Spacing mark
Me     Enclosing mark
Mn     Non-spacing mark

N      Number
Nd     Decimal number
Nl     Letter number
No     Other number

P      Punctuation
Pc     Connector punctuation
```

Pd	Dash punctuation
Pe	Close punctuation
Pf	Final punctuation
Pi	Initial punctuation
Po	Other punctuation
Ps	Open punctuation
S	Symbol
Sc	Currency symbol
Sk	Modifier symbol
Sm	Mathematical symbol
So	Other symbol
Z	Separator
Zl	Line separator
Zp	Paragraph separator
Zs	Space separator

Extended properties such as "Greek" or "InMusicalSymbols" are not supported by PCRE.

Specifying caseless matching does not affect these escape sequences. For example, `\p{Lu}` always matches only upper case letters.

The `\X` escape matches any number of Unicode characters that form an extended Unicode sequence. `\X` is equivalent to

```
(?>\PM\pM*)
```

That is, it matches a character without the "mark" property, followed by zero or more characters with the "mark" property, and treats the sequence as an atomic group (see below). Characters with the "mark" property are typically accents that affect the preceding character.

Matching characters by Unicode property is not fast, because PCRE has to search a structure that contains data for over fifteen thousand characters. That is why the traditional escape sequences such as `\d` and `\w` do not use Unicode properties in PCRE.

Simple Assertions

The fourth use of backslash is for certain simple assertions. An assertion specifies a condition that has to be met at a particular point in a match, without consuming any characters from the subject string. The use of subpatterns for more complicated assertions is described below. The backslashed assertions are:

<code>\b</code>	matches at a word boundary
<code>\B</code>	matches when not at a word boundary
<code>\A</code>	matches at start of subject
<code>\Z</code>	matches at end of subject or before newline at end
<code>\z</code>	matches at end of subject
<code>\G</code>	matches at first matching position in subject

These assertions may not appear in character classes (but note that `\b` has a different meaning, namely the backspace character, inside a character class).

A word boundary is a position in the subject string where the current character and the previous character do not both match `\w` or `\W` (i.e. one matches `\w` and the other matches `\W`), or the start or end of the string if the first or last character matches `\w`, respectively.

The `\A`, `\Z`, and `\z` assertions differ from the traditional circumflex and dollar (described in the next section) in that they only ever match at the very start and end of the subject string, whatever options are set. Thus, they are independent of multiline mode. These three assertions are not affected by the

PCRE_NOTBOL or PCRE_NOTEOL options, which affect only the behaviour of the circumflex and dollar metacharacters. However, if the *startoffset* argument of **pcre_exec()** is non-zero, indicating that matching is to start at a point other than the beginning of the subject, `\A` can never match. The difference between `\Z` and `\z` is that `\Z` matches before a newline that is the last character of the string as well as at the end of the string, whereas `\z` matches only at the end.

The `\G` assertion is true only when the current matching position is at the start point of the match, as specified by the *startoffset* argument of **pcre_exec()**. It differs from `\A` when the value of *startoffset* is non-zero. By calling **pcre_exec()** multiple times with appropriate arguments, you can mimic Perl's `/g` option, and it is in this kind of implementation where `\G` can be useful.

Note, however, that PCRE's interpretation of `\G`, as the start of the current match, is subtly different from Perl's, which defines it as the end of the previous match. In Perl, these can be different when the previously matched string was empty. Because PCRE does just one match at a time, it cannot reproduce this behaviour.

If all the alternatives of a pattern begin with `\G`, the expression is anchored to the starting match position, and the "anchored" flag is set in the compiled regular expression.

Circumflex and Dollar

Outside a character class, in the default matching mode, the circumflex character is an assertion that is true only if the current matching point is at the start of the subject string. If the *startoffset* argument of **pcre_exec()** is non-zero, circumflex can never match if the PCRE_MULTILINE option is unset. Inside a character class, circumflex has an entirely different meaning (see [Square Brackets and Character Classes](#), page B-8 and [Posix Character Classes](#), page B-9).

Circumflex need not be the first character of the pattern if a number of alternatives are involved, but it should be the first thing in each alternative in which it appears if the pattern is ever to match that branch. If all possible alternatives start with a circumflex, that is, if the pattern is constrained to match only at the start of the subject, it is said to be an "anchored" pattern. (There are also other constructs that can cause a pattern to be anchored.)

A dollar character is an assertion that is true only if the current matching point is at the end of the subject string, or immediately before a newline character that is the last character in the string (by default). Dollar need not be the last character of the pattern if a number of alternatives are involved, but it should be the last item in any branch in which it appears. Dollar has no special meaning in a character class.

The meaning of dollar can be changed so that it matches only at the very end of the string, by setting the PCRE_DOLLAR_ENDONLY option at compile time. This does not affect the `\Z` assertion.

The meanings of the circumflex and dollar characters are changed if the PCRE_MULTILINE option is set. When this is the case, they match immediately after and immediately before an internal newline character, respectively, in addition to matching at the start and end of the subject string. For example, the pattern `/^abc$/` matches the subject string "def\nabc" (where `\n` represents a newline character) in multiline mode, but not otherwise. Consequently, patterns that are anchored in single line mode because all branches start with `^` are not anchored in multiline mode, and a match for circumflex is possible when the *startoffset* argument of **pcre_exec()** is non-zero. The PCRE_DOLLAR_ENDONLY option is ignored if PCRE_MULTILINE is set.

Note that the sequences `\A`, `\Z`, and `\z` can be used to match the start and end of the subject in both modes, and if all branches of a pattern start with `\A` it is always anchored, whether PCRE_MULTILINE is set or not.

Full Stop (Period, Dot)

Outside a character class, a dot in the pattern matches any one character in the subject, including a non-printing character, but not (by default) newline. In UTF-8 mode, a dot matches any UTF-8 character, which might be more than one byte long, except (by default) newline. If the `PCRE_DOTALL` option is set, dots match newlines as well. The handling of dot is entirely independent of the handling of circumflex and dollar, the only relationship being that they both involve newline characters. Dot has no special meaning in a character class.

Matching a Single Byte

Outside a character class, the escape sequence `\C` matches any one byte, both in and out of UTF-8 mode. Unlike a dot, it can match a newline. The feature is provided in Perl in order to match individual bytes in UTF-8 mode. Because it breaks up UTF-8 characters into individual bytes, what remains in the string may be a malformed UTF-8 string. For this reason, the `\C` escape sequence is best avoided.

PCRE does not allow `\C` to appear in lookbehind assertions (described below), because in UTF-8 mode this would make it impossible to calculate the length of the lookbehind.

Square Brackets and Character Classes

An opening square bracket introduces a character class, terminated by a closing square bracket. A closing square bracket on its own is not special. If a closing square bracket is required as a member of the class, it should be the first data character in the class (after an initial circumflex, if present) or escaped with a backslash.

A character class matches a single character in the subject. In UTF-8 mode, the character may occupy more than one byte. A matched character must be in the set of characters defined by the class, unless the first character in the class definition is a circumflex, in which case the subject character must not be in the set defined by the class. If a circumflex is actually required as a member of the class, ensure it is not the first character, or escape it with a backslash.

For example, the character class `[aeiou]` matches any lower case vowel, while `[^aeiou]` matches any character that is not a lower case vowel. Note that a circumflex is just a convenient notation for specifying the characters that are in the class by enumerating those that are not. A class that starts with a circumflex is not an assertion: it still consumes a character from the subject string, and therefore it fails if the current pointer is at the end of the string.

In UTF-8 mode, characters with values greater than 255 can be included in a class as a literal string of bytes, or by using the `\x{ }` escaping mechanism.

When caseless matching is set, any letters in a class represent both their upper case and lower case versions, so for example, a caseless `[aeiou]` matches "A" as well as "a", and a caseless `[^aeiou]` does not match "A", whereas a careful version would. When running in UTF-8 mode, PCRE supports the concept of case for characters with values greater than 128 only when it is compiled with Unicode property support.

The newline character is never treated in any special way in character classes, whatever the setting of the `PCRE_DOTALL` or `PCRE_MULTILINE` options is. A class such as `[^a]` will always match a newline.

The minus (hyphen) character can be used to specify a range of characters in a character class. For example, `[d-m]` matches any letter between d and m, inclusive. If a minus character is required in a class, it must be escaped with a backslash or appear in a position where it cannot be interpreted as indicating a range, typically as the first or last character in the class.

It is not possible to have the literal character "]" as the end character of a range. A pattern such as `[W-]46]` is interpreted as a class of two characters ("W" and "-") followed by a literal string "46]", so it would match "W46]" or "-46]". However, if the "]" is escaped with a backslash it is interpreted as the end of range, so `[W-\\]46]` is interpreted as a class containing a range followed by two other characters. The octal or hexadecimal representation of "]" can also be used to end a range.

Ranges operate in the collating sequence of character values. They can also be used for characters specified numerically, for example `[\000-\037]`. In UTF-8 mode, ranges can include characters whose values are greater than 255, for example `[\x{100}-\x{2ff}]`.

If a range that includes letters is used when caseless matching is set, it matches the letters in either case. For example, `[W-c]` is equivalent to `[\\^_`wxyzabc]`, matched caselessly, and in non-UTF-8 mode, if character tables for the "fr\_FR" locale are in use, `[\xc8-\xcb]` matches accented E characters in both cases. In UTF-8 mode, PCRE supports the concept of case for characters with values greater than 128 only when it is compiled with Unicode property support.

The character types `\d`, `\D`, `\p`, `\P`, `\s`, `\S`, `\w`, and `\W` may also appear in a character class, and add the characters that they match to the class. For example, `[\dABCDEF]` matches any hexadecimal digit. A circumflex can conveniently be used with the upper case character types to specify a more restricted set of characters than the matching lower case type. For example, the class `^[^W_]` matches any letter or digit, but not underscore.

The only metacharacters that are recognized in character classes are backslash, hyphen (only where it can be interpreted as specifying a range), circumflex (only at the start), opening square bracket (only when it can be interpreted as introducing a POSIX class name - see the next section), and the terminating closing square bracket. However, escaping other non-alphanumeric characters does no harm.

Posix Character Classes

Perl supports the POSIX notation for character classes. This uses names enclosed by `[:` and `:]` within the enclosing square brackets. PCRE also supports this notation. For example,

```
[01[:alpha:]]
```

matches "0", "1", any alphabetic character, or "%". The supported class names are

<code>alnum</code>	letters and digits
<code>alpha</code>	letters
<code>ascii</code>	character codes 0 - 127
<code>blank</code>	space or tab only
<code>cntrl</code>	control characters
<code>digit</code>	decimal digits (same as <code>\d</code>)
<code>graph</code>	printing characters, excluding space
<code>lower</code>	lower case letters
<code>print</code>	printing characters, including space
<code>punct</code>	printing characters, excluding letters and digits
<code>space</code>	white space (not quite the same as <code>\s</code>)
<code>upper</code>	upper case letters
<code>word</code>	"word" characters (same as <code>\w</code>)
<code>xdigit</code>	hexadecimal digits

The "space" characters are HT (9), LF (10), VT (11), FF (12), CR (13), and space (32). Notice that this list includes the VT character (code 11). This makes "space" different to `\s`, which does not include VT (for Perl compatibility).

The name "word" is a Perl extension, and "blank" is a GNU extension from Perl 5.8. Another Perl extension is negation, which is indicated by a `^` character after the colon. For example,

```
[12[:^digit:]]
```

matches "1", "2", or any non-digit. PCRE (and Perl) also recognize the POSIX syntax `[.ch.]` and `[=ch=]` where "ch" is a "collating element", but these are not supported, and an error is given if they are encountered.

In UTF-8 mode, characters with values greater than 128 do not match any of the POSIX character classes.

Vertical Bar

Vertical bar characters are used to separate alternative patterns. For example, the pattern

```
gilbert|sullivan
```

matches either "gilbert" or "sullivan". Any number of alternatives may appear, and an empty alternative is permitted (matching the empty string). The matching process tries each alternative in turn, from left to right, and the first one that succeeds is used. If the alternatives are within a subpattern ([defined below](#)), "succeeds" means matching the rest of the main pattern as well as the alternative in the subpattern.

Internal Option Setting

The settings of the `PCRE_CASELESS`, `PCRE_MULTILINE`, `PCRE_DOTALL`, and `PCRE_EXTENDED` options can be changed from within the pattern by a sequence of Perl option letters enclosed between `"(?"` and `)"`. The option letters are

```
i for PCRE_CASELESS
m for PCRE_MULTILINE
s for PCRE_DOTALL
x for PCRE_EXTENDED
```

For example, `(?im)` sets caseless, multiline matching. It is also possible to unset these options by preceding the letter with a hyphen, and a combined setting and unsetting such as `(?im-sx)`, which sets `PCRE_CASELESS` and `PCRE_MULTILINE` while unsetting `PCRE_DOTALL` and `PCRE_EXTENDED`, is also permitted. If a letter appears both before and after the hyphen, the option is unset.

When an option change occurs at top level (that is, not inside subpattern parentheses), the change applies to the remainder of the pattern that follows. If the change is placed right at the start of a pattern, PCRE extracts it into the global options (and it will therefore show up in data extracted by the `pcre_fullinfo()` function).

An option change within a subpattern affects only that part of the current pattern that follows it, so

```
(a(?i)b)c
```


matches `abc` and `aBc` and no other strings (assuming `PCRE_CASELESS` is not used). By this means, options can be made to have different settings in different parts of the pattern. Any changes made in one alternative do carry on into subsequent branches within the same subpattern. For example,

```
(a(?:i)b|c)
```

matches `"ab"`, `"aB"`, `"c"`, and `"C"`, even though when matching `"C"` the first branch is abandoned before the option setting. This is because the effects of option settings happen at compile time. There would be some very weird behaviour otherwise.

The PCRE-specific options `PCRE_UNGREEDY` and `PCRE_EXTRA` can be changed in the same way as the Perl-compatible options by using the characters `U` and `X` respectively. The `(?X)` flag setting is special in that it must always occur earlier in the pattern than any of the additional features it turns on, even when it is at top level. It is best to put it at the start.

Subpatterns

Subpatterns are delimited by parentheses (round brackets), which can be nested. Turning part of a pattern into a subpattern does two things:

Step 1 It localizes a set of alternatives. For example, the pattern :

```
cat(aract|erpillar|)
```

matches one of the words `"cat"`, `"cactaract"`, or `"caterpillar"`. Without the parentheses, it would match `"cactaract"`, `"erpillar"` or the empty string.

Step 2 It sets up the subpattern as a capturing subpattern. This means that, when the whole pattern matches, that portion of the subject string that matched the subpattern is passed back to the caller via the *ovector* argument of `pcre_exec()`. Opening parentheses are counted from left to right (starting from 1) to obtain numbers for the capturing subpatterns.

For example, if the string `"the red king"` is matched against the pattern

```
the ((red|white) (king|queen))
```

the captured substrings are `"red king"`, `"red"`, and `"king"`, and are numbered 1, 2, and 3, respectively.

The fact that plain parentheses fulfil two functions is not always helpful. There are often times when a grouping subpattern is required without a capturing requirement. If an opening parenthesis is followed by a question mark and a colon, the subpattern does not do any capturing, and is not counted when computing the number of any subsequent capturing subpatterns. For example, if the string `"the white queen"` is matched against the pattern

```
the ((?:red|white) (king|queen))
```

the captured substrings are `"white queen"` and `"queen"`, and are numbered 1 and 2. The maximum number of capturing subpatterns is 65535, and the maximum depth of nesting of all subpatterns, both capturing and non-capturing, is 200.

As a convenient shorthand, if any option settings are required at the start of a non-capturing subpattern, the option letters may appear between the `"?"` and the `":"`. Thus the two patterns

```
(?:saturday|sunday)
(?: (?:i)saturday|sunday)
```

match exactly the same set of strings. Because alternative branches are tried from left to right, and options are not reset until the end of the subpattern is reached, an option setting in one branch does affect subsequent branches, so the above patterns match "SUNDAY" as well as "Saturday".

Named Subpatterns

Identifying capturing parentheses by number is simple, but it can be very hard to keep track of the numbers in complicated regular expressions. Furthermore, if an expression is modified, the numbers may change. To help with this difficulty, PCRE supports the naming of subpatterns, something that Perl does not provide. The Python syntax (*?P<name>...*) is used. Names consist of alphanumeric characters and underscores, and must be unique within a pattern.

Named capturing parentheses are still allocated numbers as well as names. The PCRE API provides function calls for extracting the name-to-number translation table from a compiled pattern. There is also a convenience function for extracting a captured substring by name. For further details see the *pcreapi* documentation.

Repetition

Repetition is specified by quantifiers, which can follow any of the following items:

- a literal data character
- the `.` metacharacter
- the `\C` escape sequence
- the `\X` escape sequence (in UTF-8 mode with Unicode properties)
- an escape such as `\d` that matches a single character
- a character class
- a back reference (see next section)
- a parenthesized subpattern (unless it is an assertion)

The general repetition quantifier specifies a minimum and maximum number of permitted matches, by giving the two numbers in curly brackets (braces), separated by a comma. The numbers must be less than 65536, and the first must be less than or equal to the second. For example:

```
z{2,4}
```

matches "zz", "zzz", or "zzzz". A closing brace on its own is not a special character. If the second number is omitted, but the comma is present, there is no upper limit; if the second number and the comma are both omitted, the quantifier specifies an exact number of required matches. Thus

```
[aeiou]{3,}
```

matches at least 3 successive vowels, but may match many more, while

```
\d{8}
```

matches exactly 8 digits. An opening curly bracket that appears in a position where a quantifier is not allowed, or one that does not match the syntax of a quantifier, is taken as a literal character. For example, `{,6}` is not a quantifier, but a literal string of four characters.

In UTF-8 mode, quantifiers apply to UTF-8 characters rather than to individual bytes. Thus, for example, `\x{100}{2}` matches two UTF-8 characters, each of which is represented by a two-byte sequence. Similarly, when Unicode property support is available, `\X{3}` matches three Unicode extended sequences, each of which may be several bytes long (and they may be of different lengths).

The quantifier `{0}` is permitted, causing the expression to behave as if the previous item and the quantifier were not present.

For convenience (and historical compatibility) the three most common quantifiers have single-character abbreviations:

```
*    is equivalent to {0,}
+    is equivalent to {1,}
?    is equivalent to {0,1}
```

It is possible to construct infinite loops by following a subpattern that can match no characters with a quantifier that has no upper limit, for example:

```
(a?)*
```

Earlier versions of Perl and PCRE used to give an error at compile time for such patterns. However, because there are cases where this can be useful, such patterns are now accepted, but if any repetition of the subpattern does in fact match no characters, the loop is forcibly broken.

By default, the quantifiers are "greedy", that is, they match as much as possible (up to the maximum number of permitted times), without causing the rest of the pattern to fail. The classic example of where this gives problems is in trying to match comments in C programs. These appear between `/*` and `*/` and within the comment, individual `*` and `/` characters may appear. An attempt to match C comments by applying the pattern

```
/\*. *\*/
```

to the string

```
/* first comment */ not comment /* second comment */
```

fails, because it matches the entire string owing to the greediness of the `.*` item.

However, if a quantifier is followed by a question mark, it ceases to be greedy, and instead matches the minimum number of times possible, so the pattern

```
/\*. *?\*/
```

does the right thing with the C comments. The meaning of the various quantifiers is not otherwise changed, just the preferred number of matches. Do not confuse this use of question mark with its use as a quantifier in its own right. Because it has two uses, it can sometimes appear doubled, as in

```
\d??\d
```

which matches one digit by preference, but can match two if that is the only way the rest of the pattern matches.

If the PCRE_UNGREEDY option is set (an option which is not available in Perl), the quantifiers are not greedy by default, but individual ones can be made greedy by following them with a question mark. In other words, it inverts the default behaviour.

When a parenthesized subpattern is quantified with a minimum repeat count that is greater than 1 or with a limited maximum, more memory is required for the compiled pattern, in proportion to the size of the minimum or maximum.

If a pattern starts with `.*` or `{0,}` and the `PCRE_DOTALL` option (equivalent to Perl's `/s`) is set, thus allowing the `.` to match newlines, the pattern is implicitly anchored, because whatever follows will be tried against every character position in the subject string, so there is no point in retrying the overall match at any position after the first. PCRE normally treats such a pattern as though it were preceded by `\A`.

In cases where it is known that the subject string contains no newlines, it is worth setting `PCRE_DOTALL` in order to obtain this optimization, or alternatively using `^` to indicate anchoring explicitly.

However, there is one situation where the optimization cannot be used. When `.*` is inside capturing parentheses that are the subject of a backreference elsewhere in the pattern, a match at the start may fail, and a later one succeed. Consider, for example:

```
(.*)abc\1
```

If the subject is "xyz123abc123" the match point is the fourth character. For this reason, such a pattern is not implicitly anchored.

When a capturing subpattern is repeated, the value captured is the substring that matched the final iteration. For example, after

```
(tweedle[dume]{3}\s*)+
```

has matched "tweedledum tweedledee" the value of the captured substring is "tweedledee". However, if there are nested capturing subpatterns, the corresponding captured values may have been set in previous iterations. For example, after

```
/(a|(b))+/
```

matches "aba" the value of the second captured substring is "b".

Atomic Grouping and Possessive Quantifiers

With both maximizing and minimizing repetition, failure of what follows normally causes the repeated item to be re-evaluated to see if a different number of repeats allows the rest of the pattern to match. Sometimes it is useful to prevent this, either to change the nature of the match, or to cause it fail earlier than it otherwise might, when the author of the pattern knows there is no point in carrying on.

Consider, for example, the pattern `\d+foo` when applied to the subject line

```
123456bar
```

After matching all 6 digits and then failing to match "foo", the normal action of the matcher is to try again with only 5 digits matching the `\d+` item, and then with 4, and so on, before ultimately failing. "Atomic grouping" (a term taken from Jeffrey Friedl's book) provides the means for specifying that once a subpattern has matched, it is not to be re-evaluated in this way.

If we use atomic grouping for the previous example, the matcher would give up immediately on failing to match "foo" the first time. The notation is a kind of special parenthesis, starting with `(?>` as in this example:

```
(?>\d+)foo
```

This kind of parenthesis "locks up" the part of the pattern it contains once it has matched, and a failure further into the pattern is prevented from backtracking into it. Backtracking past it to previous items, however, works as normal.

An alternative description is that a subpattern of this type matches the string of characters that an identical standalone pattern would match, if anchored at the current point in the subject string.

Atomic grouping subpatterns are not capturing subpatterns. Simple cases such as the above example can be thought of as a maximizing repeat that must swallow everything it can. So, while both `\d+` and `\d+?` are prepared to adjust the number of digits they match in order to make the rest of the pattern match, `(?>\d+)` can only match an entire sequence of digits.

Atomic groups in general can of course contain arbitrarily complicated subpatterns, and can be nested. However, when the subpattern for an atomic group is just a single repeated item, as in the example above, a simpler notation, called a "possessive quantifier" can be used. This consists of an additional `+` character following a quantifier. Using this notation, the previous example can be rewritten as

```
\d++foo
```

Possessive quantifiers are always greedy; the setting of the `PCRE_UNGREEDY` option is ignored. They are a convenient notation for the simpler forms of atomic group. However, there is no difference in the meaning or processing of a possessive quantifier and the equivalent atomic group.

The possessive quantifier syntax is an extension to the Perl syntax. It originates in Sun's Java package.

When a pattern contains an unlimited repeat inside a subpattern that can itself be repeated an unlimited number of times, the use of an atomic group is the only way to avoid some failing matches taking a very long time indeed. The pattern

```
(\D+|<\d+>)*[!?] 
```

matches an unlimited number of substrings that either consist of non-digits, or digits enclosed in `<>`, followed by either `!` or `?`. When it matches, it runs quickly. However, if it is applied to

```
aaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaa
```

it takes a long time before reporting failure. This is because the string can be divided between the internal `\D+` repeat and the external `*` repeat in a large number of ways, and all have to be tried. (The example uses `[!?]` rather than a single character at the end, because both PCRE and Perl have an optimization that allows for fast failure when a single character is used. They remember the last single character that is required for a match, and fail early if it is not present in the string.) If the pattern is changed so that it uses an atomic group, like this:

```
((?>\D+)|<\d+>)*[!?] 
```

sequences of non-digits cannot be broken, and failure happens quickly.

Back References

Outside a character class, a backslash followed by a digit greater than 0 (and possibly further digits) is a back reference to a capturing subpattern earlier (that is, to its left) in the pattern, provided there have been that many previous capturing left parentheses.

However, if the decimal number following the backslash is less than 10, it is always taken as a back reference, and causes an error only if there are not that many capturing left parentheses in the entire pattern. In other words, the parentheses that are referenced need not be to the left of the reference for numbers less than 10. See [Non-printing Characters, page B-3](#) for further details of the handling of digits following a backslash.

A back reference matches whatever actually matched the capturing subpattern in the current subject string, rather than anything matching the subpattern itself (see [Subpatterns as Subroutines, page B-21](#) for a way of doing that). So the pattern

```
(sens|respons)e and \1ibility
```

matches "sense and sensibility" and "response and responsibility", but not "sense and responsibility". If careful matching is in force at the time of the back reference, the case of letters is relevant. For example,

```
((?i)rah)\s+\1
```

matches "rah rah" and "RAH RAH", but not "RAH rah", even though the original capturing subpattern is matched caselessly.

Back references to named subpatterns use the Python syntax (?P=name). We could rewrite the above example as follows:

```
(?<p1>( ?i)rah)\s+(?P=p1)
```

There may be more than one back reference to the same subpattern. If a subpattern has not actually been used in a particular match, any back references to it always fail. For example, the pattern

```
(a|(bc))\2
```

always fails if it starts to match "a" rather than "bc". Because there may be many capturing parentheses in a pattern, all digits following the backslash are taken as part of a potential back reference number. If the pattern continues with a digit character, some delimiter must be used to terminate the back reference. If the PCRE_EXTENDED option is set, this can be whitespace. Otherwise an empty comment (see [Comments, page B-20](#)) can be used.

A back reference that occurs inside the parentheses to which it refers fails when the subpattern is first used, so, for example, (a\1) never matches. However, such references can be useful inside repeated subpatterns. For example, the pattern

```
(a|b\1)+
```

matches any number of "a"s and also "aba", "ababbaa" etc. At each iteration of the subpattern, the back reference matches the character string corresponding to the previous iteration. In order for this to work, the pattern must be such that the first iteration does not need to match the back reference. This can be done using alternation, as in the example above, or by a quantifier with a minimum of zero.

Assertions

An assertion is a test on the characters following or preceding the current matching point that does not actually consume any characters. The simple assertions coded as \b, \B, \A, \G, \Z, \z, ^ and \$ are described [above](#).

More complicated assertions are coded as subpatterns. There are two kinds: those that look ahead of the current position in the subject string, and those that look behind it. An assertion subpattern is matched in the normal way, except that it does not cause the current matching position to be changed.

Assertion subpatterns are not capturing subpatterns, and may not be repeated, because it makes no sense to assert the same thing several times. If any kind of assertion contains capturing subpatterns within it, these are counted for the purposes of numbering the capturing subpatterns in the whole pattern. However, substring capturing is carried out only for positive assertions, because it does not make sense for negative assertions.

Lookahead Assertions

Lookahead assertions start with `(?=` for positive assertions and `(?!` for negative assertions. For example,

```
\w+ (?:;)
```

matches a word followed by a semicolon, but does not include the semicolon in the match, and

```
foo(?!bar)
```

matches any occurrence of "foo" that is not followed by "bar". Note that the apparently similar pattern

```
(?!foo)bar
```

does not find an occurrence of "bar" that is preceded by something other than "foo"; it finds any occurrence of "bar" whatsoever, because the assertion `(?!foo)` is always true when the next three characters are "bar". A lookbehind assertion is needed to achieve the other effect.

If you want to force a matching failure at some point in a pattern, the most convenient way to do it is with `(?!)` because an empty string always matches, so an assertion that requires there not to be an empty string must always fail.

Lookbehind Assertions

Lookbehind assertions start with `(?<=` for positive assertions and `(?<!` for negative assertions. For example,

```
(?<!foo)bar
```

does find an occurrence of "bar" that is not preceded by "foo". The contents of a lookbehind assertion are restricted such that all the strings it matches must have a fixed length. However, if there are several alternatives, they do not all have to have the same fixed length. Thus

```
(?<=bullock|donkey)
```

is permitted, but

```
(?<!dogs?|cats?)
```

causes an error at compile time. Branches that match different length strings are permitted only at the top level of a lookbehind assertion. This is an extension compared with Perl (at least for 5.8), which requires all branches to match the same length of string. An assertion such as

```
(?<=ab(c|de))
```

is not permitted, because its single top-level branch can match two different lengths, but it is acceptable if rewritten to use two top-level branches:

```
(?<=abc|abde)
```

The implementation of lookbehind assertions is, for each alternative, to temporarily move the current position back by the fixed width and then try to match. If there are insufficient characters before the current position, the match is deemed to fail.

PCRE does not allow the `\C` escape (which matches a single byte in UTF-8 mode) to appear in lookbehind assertions, because it makes it impossible to calculate the length of the lookbehind. The `\X` escape, which can match different numbers of bytes, is also not permitted.

Atomic groups can be used in conjunction with lookbehind assertions to specify efficient matching at the end of the subject string. Consider a simple pattern such as

```
abcd$
```

when applied to a long string that does not match. Because matching proceeds from left to right, PCRE will look for each "a" in the subject and then see if what follows matches the rest of the pattern. If the pattern is specified as

```
^.*abcd$
```

the initial `.*` matches the entire string at first, but when this fails (because there is no following "a"), it backtracks to match all but the last character, then all but the last two characters, and so on. Once again the search for "a" covers the entire string, from right to left, so we are no better off. However, if the pattern is written as

```
^(?>.*)(?<=abcd)
```

or, equivalently, using the possessive quantifier syntax,

```
^.*+(?<=abcd)
```

there can be no backtracking for the `.*` item; it can match only the entire string. The subsequent lookbehind assertion does a single test on the last four characters. If it fails, the match fails immediately. For long strings, this approach makes a significant difference to the processing time.

Using Multiple Assertions

Several assertions (of any sort) may occur in succession. For example,

```
(?<=\d{3})(?!999)foo
```

matches "foo" preceded by three digits that are not "999". Notice that each of the assertions is applied independently at the same point in the subject string. First there is a check that the previous three characters are all digits, and then there is a check that the same three characters are not "999". This pattern does *not* match "foo" preceded by six characters, the first of which are digits and the last three of which are not "999". For example, it doesn't match "123abcfoo". A pattern to do that is


```
(?<=\d{3}\.)(?!999)foo
```

This time the first assertion looks at the preceding six characters, checking that the first three are digits, and then the second assertion checks that the preceding three characters are not "999".

Assertions can be nested in any combination. For example,

```
(?<=(?!foo)bar)baz
```

matches an occurrence of "baz" that is preceded by "bar" which in turn is not preceded by "foo", while

```
(?<=\d{3}(?!999)\.)(?!999)foo
```

is another pattern that matches "foo" preceded by three digits and any three characters that are not "999".

Conditional Subpatterns

It is possible to cause the matching process to obey a subpattern conditionally or to choose between two alternative subpatterns, depending on the result of an assertion, or whether a previous capturing subpattern matched or not. The two possible forms of conditional subpattern are

```
(?(condition)yes-pattern)
(?(condition)yes-pattern|no-pattern)
```

If the condition is satisfied, the yes-pattern is used; otherwise the no-pattern (if present) is used. If there are more than two alternatives in the subpattern, a compile-time error occurs.

There are three kinds of condition. If the text between the parentheses consists of a sequence of digits, the condition is satisfied if the capturing subpattern of that number has previously matched. The number must be greater than zero. Consider the following pattern, which contains non-significant white space to make it more readable (assume the PCRE_EXTENDED option) and to divide it into three parts for ease of discussion:

```
( \ ( ) ?    [ ^ ( ) ] +    ( ? ( 1 ) \ ) )
```

The first part matches an optional opening parenthesis, and if that character is present, sets it as the first captured substring. The second part matches one or more characters that are not parentheses. The third part is a conditional subpattern that tests whether the first set of parentheses matched or not. If they did, that is, if subject started with an opening parenthesis, the condition is true, and so the yes-pattern is executed and a closing parenthesis is required. Otherwise, since no-pattern is not present, the subpattern matches nothing. In other words, this pattern matches a sequence of non-parentheses, optionally enclosed in parentheses.

If the condition is the string (R), it is satisfied if a recursive call to the pattern or subpattern has been made. At "top level", the condition is false. This is a PCRE extension. Recursive patterns are described in the next section.

If the condition is not a sequence of digits or (R), it must be an assertion. This may be a positive or negative lookahead or lookbehind assertion. Consider this pattern, again containing non-significant white space, and with the two alternatives on the second line:

```
(? ( ? = [ ^ a - z ] * [ a - z ] )
  \ d { 2 } - [ a - z ] { 3 } - \ d { 2 }    |    \ d { 2 } - \ d { 2 } - \ d { 2 } )
```

The condition is a positive lookahead assertion that matches an optional sequence of non-letters followed by a letter. In other words, it tests for the presence of at least one letter in the subject. If a letter is found, the subject is matched against the first alternative; otherwise it is matched against the second. This pattern matches strings in one of the two forms dd-aaa-dd or dd-dd-dd, where aaa are letters and dd are digits.

Comments

The sequence (?# marks the start of a comment that continues up to the next closing parenthesis. Nested parentheses are not permitted. The characters that make up a comment play no part in the pattern matching at all.

If the PCRE_EXTENDED option is set, an unescaped # character outside a character class introduces a comment that continues up to the next newline character in the pattern.

Recursive Patterns

Consider the problem of matching a string in parentheses, allowing for unlimited nested parentheses. Without the use of recursion, the best that can be done is to use a pattern that matches up to some fixed depth of nesting. It is not possible to handle an arbitrary nesting depth. Perl provides a facility that allows regular expressions to recurse (amongst other things). It does this by interpolating Perl code in the expression at run time, and the code can refer to the expression itself. A Perl pattern to solve the parentheses problem can be created like this:

```
$re = qr{\( (? : (?>[^\(]+) | (?p{$re}) ) * \)}x;
```

The (?p{...}) item interpolates Perl code at run time, and in this case refers recursively to the pattern in which it appears. Obviously, PCRE cannot support the interpolation of Perl code. Instead, it supports some special syntax for recursion of the entire pattern, and also for individual subpattern recursion.

The special item that consists of (? followed by a number greater than zero and a closing parenthesis is a recursive call of the subpattern of the given number, provided that it occurs inside that subpattern. (If not, it is a "subroutine" call, which is described in the next section.) The special item (?R) is a recursive call of the entire regular expression.

For example, this PCRE pattern solves the nested parentheses problem (assume the PCRE_EXTENDED option is set so that white space is ignored):

```
\( ( (?>[^\(]+) | (?R) ) * \)
```

First it matches an opening parenthesis. Then it matches any number of substrings which can either be a sequence of non-parentheses, or a recursive match of the pattern itself (that is a correctly parenthesized substring). Finally there is a closing parenthesis.

If this were part of a larger pattern, you would not want to recurse the entire pattern, so instead you could use this:

```
( \ ( ( (?>[^\(]+) | (?1) ) * \ ) )
```

We have put the pattern into parentheses, and caused the recursion to refer to them instead of the whole pattern. In a larger pattern, keeping track of parenthesis numbers can be tricky. It may be more convenient to use named parentheses instead. For this, PCRE uses `(?P>name)`, which is an extension to the Python syntax that PCRE uses for named parentheses (Perl does not provide named parentheses). We could rewrite the above example as follows:

```
(?P<pn> \ ( ( ?> [ ^ ( ) ] + ) | ( ?P>pn ) * \ ) )
```

This particular example pattern contains nested unlimited repeats, and so the use of atomic grouping for matching strings of non-parentheses is important when applying the pattern to strings that do not match. For example, when this pattern is applied to

```
(aaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaa ( )
```

it yields "no match" quickly. However, if atomic grouping is not used, the match runs for a very long time indeed because there are so many different ways the `+` and `*` repeats can carve up the subject, and all have to be tested before failure can be reported.

At the end of a match, the values set for any capturing subpatterns are those from the outermost level of the recursion at which the subpattern value is set. If you want to obtain intermediate values, a callout function can be used (see [Subpatterns as Subroutines](#), page B-21 and the **pcrecallout** documentation). If the pattern above is matched against

```
(ab(cd)ef)
```

the value for the capturing parentheses is "ef", which is the last value taken on at the top level. If additional parentheses are added, giving

```
\ ( ( ( ?> [ ^ ( ) ] + ) | ( ?R ) * ) \ )
    ^               ^
    ^               ^
```

the string they capture is "ab(cd)ef", the contents of the top level parentheses. If there are more than 15 capturing parentheses in a pattern, PCRE has

to obtain extra memory to store data during a recursion, which it does by using **pcre_malloc**, freeing it via **pcre_free** afterwards. If no memory can be obtained, the match fails with the **PCRE_ERROR_NOMEMORY** error.

Do not confuse the `(?R)` item with the condition `(R)`, which tests for recursion. Consider this pattern, which matches text in angle brackets, allowing for arbitrary nesting. Only digits are allowed in nested brackets (that is, when recursing), whereas any characters are permitted at the outer level.

```
< (?: (?(R) \d++ | [ ^<> ] * ) | ( ?R ) ) * >
```

In this pattern, `(?R)` is the start of a conditional subpattern, with two different alternatives for the recursive and non-recursive cases. The `(?R)` item is the actual recursive call.

Subpatterns as Subroutines

If the syntax for a recursive subpattern reference (either by number or by name) is used outside the parentheses to which it refers, it operates like a subroutine in a programming language. An earlier example pointed out that the pattern

```
(sens|respons)e and \libility
```

matches "sense and sensibility" and "response and responsibility", but not "sense and responsibility". If instead the pattern

```
(sens|respons)e and (?1)ibility
```

is used, it does match "sense and responsibility" as well as the other two strings. Such references must, however, follow the subpattern to which they refer.

Callouts

Perl has a feature whereby using the sequence `{...}` causes arbitrary Perl code to be obeyed in the middle of matching a regular expression. This makes it possible, amongst other things, to extract different substrings that match the same pair of parentheses when there is a repetition.

PCRE provides a similar feature, but of course it cannot obey arbitrary Perl code. The feature is called "callout". The caller of PCRE provides an external function by putting its entry point in the global variable `pcre_callout`. By default, this variable contains NULL, which disables all calling out.

Within a regular expression, `(?C)` indicates the points at which the external function is to be called. If you want to identify different callout points, you can put a number less than 256 after the letter C. The default value is zero. For example, this pattern has two callout points:

```
(?C1) \dabc (?C2) def
```

If the `PCRE_AUTO_CALLOUT` flag is passed to `pcre_compile()`, callouts are automatically installed before each item in the pattern. They are all numbered 255.

During matching, when PCRE reaches a callout point (and `pcre_callout` is set), the external function is called. It is provided with the number of the callout, the position in the pattern, and, optionally, one item of data originally supplied by the caller of `pcre_exec()`. The callout function may cause matching to proceed, to backtrack, or to fail altogether. A complete description of the interface to the callout function is given in the **pcrecallout** documentation.

Last updated: 09 September 2004

Copyright © 1997-2004 University of Cambridge.



Date/Time Format Specification

The date/time field parsing is supported using the Unix `strptime()` standard C library function.

The **`strptime()`** function is the converse function to **`strftime()`** and converts the character string pointed to by *s* to values which are stored in the *tm* structure pointed to by *tm*, using the format specified by *format*. Here *format* is a character string that consists of field descriptors and text characters, reminiscent of `scanf(3)`. Each field descriptor consists of a `%` character followed by another character that specifies the replacement for the field descriptor. All other characters in the *format* string must have a matching character in the input string, except for whitespace, which matches zero or more whitespace characters in the input string.

The **`strptime()`** function processes the input string from left to right. Each of the three possible input elements (whitespace, literal, or format) are handled one after the other. If the input cannot be matched to the format string the function stops. The remainder of the format and input strings are not processed.

The supported input field descriptors are listed below. In case a text string (such as a weekday or month name) is to be matched, the comparison is case insensitive. In case a number is to be matched, leading zeros are permitted but not required.

`%%`

The `%` character.

`%a` or `%A`

The weekday name according to the current locale, in abbreviated form or the full name.

`%b` or `%B` or `%h`

The month name according to the current locale, in abbreviated form or the full name.

`%c`

The date and time representation for the current locale.

`%C`

The century number (0-99).

`%d` or `%e`

The day of month (1-31).

`%D`

Equivalent to `%m/%d/%y`. (This is the American style date, very confusing to non-Americans, especially since `%d/%m/%y` is widely used in Europe. The ISO 8601 standard format is `%Y-%m-%d`.)

`%H`

The hour (0-23).

%I

The hour on a 12-hour clock (1-12).

%j

The day number in the year (1-366).

%m

The month number (1-12).

%M

The minute (0-59).

%n or **%t**

Arbitrary whitespace.

%p

The locale's equivalent of AM or PM. (Note: there may be none.)

%r

The 12-hour clock time (using the locale's AM or PM). In the POSIX locale equivalent to **%I:%M:%S %p**. If *t_fmt_ampm* is empty in the LC_TIME part of the current locale then the behaviour is undefined.

%R

Equivalent to **%H:%M**.

%S

The second (0-60; 60 may occur for leap seconds; earlier also 61 was allowed).

%T

Equivalent to **%H:%M:%S**.

%U

The week number with Sunday the first day of the week (0-53). The first Sunday of January is the first day of week 1.

%w

The weekday number (0-6) with Sunday = 0.

%W

The week number with Monday the first day of the week (0-53). The first Monday of January is the first day of week 1.

%x

The date, using the locale's date format.

%X

The time, using the locale's time format.

%y

The year within century (0-99). When a century is not otherwise specified, values in the range 69-99 refer to years in the twentieth century (1969-1999); values in the range 00-68 refer to years in the twenty-first century (2000-2068).

%Y

The year, including century (for example, 1991).

Some field descriptors can be modified by the E or O modifier characters to indicate that an alternative format or specification should be used. If the alternative format or specification does not exist in the current locale, the unmodified field descriptor is used.

The E modifier specifies that the input string may contain alternative locale-dependent versions of the date and time representation:

%Ec

The locale's alternative date and time representation.

%EC

The name of the base year (period) in the locale's alternative representation.

%Ex

The locale's alternative date representation.

%EX

The locale's alternative time representation.

%Ey

The offset from %EC (year only) in the locale's alternative representation.

%EY

The full alternative year representation.

The O modifier specifies that the numerical input may be in an alternative locale-dependent format:

%Od or %Oe

The day of the month using the locale's alternative numeric symbols; leading zeros are permitted but not required.

%OH

The hour (24-hour clock) using the locale's alternative numeric symbols.

%OI

The hour (12-hour clock) using the locale's alternative numeric symbols.

%Om

The month using the locale's alternative numeric symbols.

%OM

The minutes using the locale's alternative numeric symbols.

%OS

The seconds using the locale's alternative numeric symbols.

%OU

The week number of the year (Sunday as the first day of the week) using the locale's alternative numeric symbols.

%Ow

The number of the weekday (Sunday=0) using the locale's alternative numeric symbols.

%OW

The week number of the year (Monday as the first day of the week) using the locale's alternative numeric symbols.

%Oy

The year (offset from %C) using the locale's alternative numeric symbols.

%F

Equivalent to %Y-%m-%d, the ISO 8601 date format.

%g

The year corresponding to the ISO week number, but without the century (0-99).

%G

The year corresponding to the ISO week number. (For example, 1991.)

%u

The day of the week as a decimal number (1-7, where Monday = 1).

%V

The ISO 8601:1988 week number as a decimal number (1-53). If the week (starting on Monday) containing 1 January has four or more days in the new year, then it is considered week 1. Otherwise, it is the last week of the previous year, and the next week is week 1.

%z

An RFC-822/ISO 8601 standard time zone specification.

%Z

The timezone name.

Similarly, because of GNU extensions to *strftime*, %k is accepted as a synonym for %H, and %l should be accepted as a synonym for %I, and %P is accepted as a synonym for %p. Finally

%s

The number of seconds since the epoch, i.e., since 1970-01-01 00:00:00 UTC. Leap seconds are not counted unless leap second support is available.



GLOSSARY

#

5-tuple (Quintuple) The five pieces of data found within all IP-based network packets: source IP address, source port, destination IP address, destination port, and protocol. You can define inspection rules, queries, and reports using the data found in the 5-tuple.

A

(\)

Access IP Address This is the IP address that MARS uses to connect to the device and to get its configuration information. MARS needs this address for NAT-related session correlation, attack path calculation, and mitigation enter access information.

Activate Making changes or edits known to the MARS after submitting changes.

D

Devices The hosts and reporting devices present in the system.

Discovery The act of identifying, either automatically or manually, devices in networks.

Dynamic Vulnerability Scanning The MARS STM probes selected networks, and their components, for vulnerabilities.

E

Event A security event reported to the MARS STM appliance. Events have: types, sources, destinations, reporting devices, etc.

Event Types Groups of similar security events. An event type is the normalized signature from a reporting device.

F

False Positive An event that resembles a valid security threat, but is not.

Firing Events An event that contributed to a rule firing.

I

Incident	Incidents are collections of events and sessions that meet the criteria for a rule, having helped to cause it to fire.
Incident Instances	An instance of an incident.

M

MI B	management information base
mitigate	To stop a detected attack or anomaly. The method of mitigation varies based on network composition and configuration.

O

Offset	The offset of a firing event is the line number of the rule criteria that this firing event matches.

P

Pre NAT Source Address	Session endpoints.
Post NAT Source Address	The source as appearing at the destination.
Post NAT Destination Address	Session endpoints.
Pre NAT Destination Address	The destination as appearing at the source.

Q

Query	A user-defined request to the database for information.

R

Report	A user-defined request to the database on an automatic or on-demand basis.
Reporting Device	A discovered device that reports information – usually in the form of logs – to a MARS STM appliance.

Reporting IP Address This is the IP address as it appears to MARS. This address is where the logs (syslog, SNMP traps, LEA) come from.

Rule The sub-set of events that contributed to the incidents of the specified rules firing.

S

Service A protocol and range of IP addresses.

Session A session is a collection of events that all share a common source and destination, which were reported within a given time window. For example, usually the events in a session map well to the events generated between the opening and closing of a TCP/IP connection.

Sessionize Combining event data from multiple reporting devices to reconstruct the occurrence of a session. Sessionizing takes two forms: reconstructing a session-oriented protocol, such as TCP, where the initial handshake and the session tear down and reconstructing a sessionless protocol, such as UDP, where the initial start and session end times are defined more based on first and last packets tracked within a restricted time period. In other words, packets that fall outside of the time period are considered part of different sessions.

T

True Positive A valid security threat.

U

Unreported device A device from which the MARS Appliance receives events, such as syslog messages, SNMP notifications, or NetFlow events, but the device is not defined in the appliance. Without a definition, MARS is unable to correlate events correctly as it needs to know which message format to use in parsing.

T

True Positive A valid security threat.



Numerics

802.1x, logging in Cisco Secure ACS [14-5](#)

A

AAA devices [14-1](#)

Action [19-3](#)

Activate button [21-18, 21-19, 21-21, 21-23, 23-1](#)

adding

cell phone number [22-11, 23-11](#)

CSV file [2-20](#)

devices [2-18](#)

manually [2-18](#)

seed file [2-20](#)

drop rules [21-22](#)

event groups [23-2](#)

inspection rules [21-19](#)

pager number [22-11, 23-11](#)

seed file [2-20](#)

service [23-8](#)

user [22-10, 23-9](#)

user group [23-12](#)

adding IP groups [23-4](#)

adding service provider [22-11, 23-11](#)

admin roles, see user management [23-8](#)

Adobe SVG [17-10](#)

alert

action [21-15](#)

Distributed Threat Management [21-15](#)

Email [21-15](#)

NONE [21-15](#)

Page [21-15](#)

SMS [21-15](#)

SNMP [21-15](#)

Syslog [21-15](#)

alerts [22-1](#)

all matching event raw messages [20-7](#)

all matching events [20-7](#)

all matching sessions [20-7](#)

anomaly detection, see NetFlow [2-31](#)

attack diagram [17-9](#)

attack paths

L2 [19-5](#)

L3 [19-5](#)

audit trail [24-3](#)

B

bootstrap

devices [1-5](#)

bytes transmitted [20-8](#)

C

cell phone paging [22-11, 23-11](#)

changing

drop rule status [21-21](#)

inspection rule status [21-17](#)

Cisco Adaptive Security Appliance, see Cisco ASA [4-1](#)

Cisco ASA

add to MARS [4-5](#)

bootstrapping [4-2](#)

security context

add discovered [4-10](#)

define reporting options for [4-11](#)

- make MARS aware of [4-8](#)
- Cisco Firewall Services Modules, see Cisco FWSM [4-1](#)
- Cisco FWSM
 - add to MARS [4-5](#)
 - bootstrapping [4-2](#)
 - security context
 - add discovered [4-10](#)
 - define reporting options for [4-11](#)
 - make MARS aware of [4-8](#)
- Cisco Secure ACS, 802.1x feature support [14-5](#)
- Cisco Secure ACS, 802.1x support [14-1](#)
- Cisco Secure ACS, audit logs required by MARS [14-3](#)
- Cisco Secure ACS, bootstrap [14-2](#)
- Cisco Secure ACS, event logs studied by MARS [14-1](#)
- Cisco Secure ACS, MARS agent [14-7](#)
- Cisco Secure ACS, NAC support [14-1](#)
- Cisco Secure ACS, representing in MARS [14-12](#)
- Cisco Secure ACS, sever support [14-2](#)
- Cisco Secure ACS, solution engine support [14-2](#)
- Cisco Secure ACS, supported versions [14-1](#)
- Cisco Secure ACS, TACACS+ command
 - authorization [14-6](#)
- Collapse All [19-5](#)
- columns
 - seed file [2-22](#)
- Common Vulnerabilities and Exposures [23-2](#)
- community strings [2-37](#)
- configuration
 - NetFlow [2-30](#)
- creating
 - report [20-24](#)
- CSV files [2-20](#)
- CVE [23-2](#)

D

- data reduction [17-9](#)
- default password
 - change [24-9](#)

- deleting service [23-8](#)
- destination IP address ranking [20-6](#)
- destination network group ranking [20-6](#)
- destination network ranking [20-6](#)
- destination ranking [20-6](#)
- device,re-add [2-19](#)
- devices
 - bootstrap overview [1-5](#)
 - define
 - overview [1-6, 16-10](#)
 - deleting [2-19](#)
 - deleting all displayed [2-20](#)
 - edit [2-18](#)
- diagrams
 - attack [17-9](#)
- discovering networks
 - automatic [2-39](#)
- discovery
 - scheduling [2-39](#)
 - updating [2-39](#)
- display format
 - query [20-5](#)
- distributed threat mitigation, taskflow order [1-7](#)
- drop rule
 - activate and inactive [21-21](#)
- drop rules
 - adding [21-22](#)
 - editing [21-22](#)
- drop rule status
 - changing [21-21](#)
- DTM, See distributed threat mitigation. [1-7](#)
- dynamic information [19-10](#)
- dynamic vulnerability scanning [2-29](#)

E

- editing
 - drop rules [21-22](#)
 - host information [23-6](#)

- inspection rules [21-18](#)
- IP groups [23-3](#)
- service [23-8](#)
- user [23-12](#)
- event groups [23-2](#)
- event log
 - changing pulling time interval for Windows [10-10](#)
- event management [23-1](#)
 - editing [23-2](#)
- Event Type [19-3](#)
- event type group ranking [20-6](#)
- event type ranking [20-5](#)
- Expand All [19-5](#)

F

- false positive
 - system determined [19-8](#)
 - unconfirmed [19-8](#)
 - user confirmed
 - false positive [19-8](#)
 - positive [19-8](#)
- false positives
 - tuning [19-5](#)

H

- hosts
 - adding [23-4](#)
 - editing [23-6](#)
- Hot Spot Graph [17-9](#)

I

- incident count [20-8](#)
- Incident Details page [19-4](#)
- Incident ID [19-3](#)
- Incident Path [19-3](#)

- incidents [17-8](#)
 - action [19-3](#)
 - event type [19-3](#)
 - incident ID [19-3](#)
 - incident path [19-3](#)
 - incident vector [19-3](#)
 - instances [19-6](#)
 - matched rule [19-3](#)
 - severity [19-3](#)
 - time [19-3](#)
 - time ranges [19-4](#)
- incidents table
 - navigation [19-3](#)
- incident table [19-5](#)
- Incident Vector [19-3](#)
- inspection rule
 - activate and inactive [21-17](#)
- inspection rules
 - adding [21-19](#)
 - editing [21-18](#)
- inspection rule status
 - changing [21-17](#)
- instances
 - incidents [19-6](#)
- IP groups
 - adding [23-4](#)
 - editing [23-3](#)
- IP management [23-3](#)
 - adding
 - hosts [23-4](#)
 - IP range [23-4](#)
 - network [23-4](#)
 - variable [23-4](#)

L

- L2 attack path [19-5](#)
- L3 attack path [19-5](#)
- Linux host, bootstrap [10-2](#)

loading

MARS

seed file [2-24](#)

log files [24-2](#)

M

MAC address report [20-7](#)

management

events [23-1](#)

IP [23-3](#)

service [23-7](#)

user [23-8](#)

matched incident ranking [20-7](#)

Matched Rule [19-3](#)

matched rule ranking [20-7](#)

Microsoft Windows host, bootstrap [10-4](#)

mitigate [19-5](#)

mitigation policy

suggested content [1-1](#)

monitoring policy

suggested content [1-1](#)

N

NAC, AAA server support [14-1](#)

NAT connection report [20-7](#)

NetFlow, enable processing [2-34](#)

NetFlow [2-30](#)

configuration [2-30](#)

Global NetFlow UPD Port [2-35](#)

NetFlow, bootstrap reporting devices [2-32](#)

NetFlow,enable processing [2-35](#)

NetFlow,examined networks [2-35](#)

NetFlow,guidelines [2-32](#)

NetFlow,how it is used [2-31](#)

NetFlow,performance tuning [2-35](#)

NetFlow,supported versions [2-31](#)

network group ranking [20-6](#)

network ranking [20-6](#)

Network Status tab

Incidents [17-12](#)

Top Destinations [17-13](#)

Top Event Types [17-12](#)

Top Sources [17-13](#)

O

Order/Rank By [20-7](#)

order by [20-7](#)

bytes transmitted [20-8](#)

incident count [20-8](#)

session count [20-7](#)

time [20-8](#)

P

pager [22-11, 23-11](#)

PIX

add to MARS [4-5](#)

bootstrapping [4-2](#)

security context

add discovered [4-10](#)

define reporting options for [4-11](#)

make MARS aware of [4-8](#)

PIX Security Appliance, see PIX [4-1](#)

PN Log agent [14-7](#)

PN Log Agent, error messages [14-10](#)

PN MARS

audit trail [24-3](#)

log files [24-2](#)

seed file columns [2-22](#)

post NAT destination addresses [20-11](#)

post NAT source addresses [20-10](#)

pre NAT destination addresses [20-11](#)

pre NAT source addresses [20-10](#)

protocol ranking [20-6](#)
 public networks [2-38](#)

Q

queries

action

ANY [20-12](#)

actions [20-12](#)

destination IP [20-11](#)

ANY [20-11](#)

devices [20-11](#)

IP addresses [20-11](#)

IP ranges [20-11](#)

networks [20-11](#)

post NAT destination addresses [20-11](#)

pre NAT destination addresses [20-11](#)

devices [20-11](#)

display format

all matching event raw messages [20-7](#)

all matching events [20-7](#)

all matching sessions [20-7](#)

destination IP address ranking [20-6](#)

destination ranking [20-6](#)

event type group ranking [20-6](#)

MAC address report [20-7](#)

matched incident ranking [20-7](#)

matched rule ranking [20-7](#)

NAT connection report [20-7](#)

protocol ranking [20-6](#)

reporting device ranking [20-7](#)

reporting device type ranking [20-7](#)

source IP address ranking [20-6](#)

source port ranking [20-6](#)

unknown event report [20-7](#)

use only firing events [20-8](#)

event type grouping [20-11](#)

event types [20-11](#)

ANY [20-11](#)

operation

AND [20-12, 21-13](#)

FOLLOWED-BY [20-12, 21-13](#)

none [20-12, 21-13](#)

OR [20-12, 21-13](#)

result format

destination network group ranking [20-6](#)

destination network ranking [20-6](#)

event type ranking [20-5](#)

network group ranking [20-6](#)

network ranking [20-6](#)

reported user ranking [20-7](#)

source network group ranking [20-6](#)

source network ranking [20-6](#)

rule [20-12](#)

ANY [20-12](#)

save as

reports [20-13](#)

rules [20-13](#)

service

ANY [20-11](#)

defined services [20-11](#)

service variables [20-11](#)

severity

ANY [20-12](#)

green [20-12](#)

red [20-12](#)

yellow [20-12](#)

source IP

ANY [20-10](#)

devices [20-10](#)

IP addresses [20-10](#)

IP ranges [20-10](#)

networks [20-10](#)

post NAT source addresses [20-10](#)

pre NAT source addresses [20-10](#)

variables [20-10](#)

time range

last [20-8](#)

- start and end times [20-8](#)
- zone [20-12](#)
- query
 - display format [20-5](#)
 - reporting device ranking [2-27](#)
- Query page [20-1](#)

R

- rank by [20-7](#)
 - bytes transmitted [20-8](#)
 - incident count [20-8](#)
 - session count [20-7](#)
 - time [20-8](#)
- remediation policy
 - suggested content [1-1](#)
- removing
 - user [23-12](#)
- report
 - adding [20-24](#)
 - delete [20-25](#)
 - edit [20-26](#)
 - new [20-24](#)
- reported user ranking [20-7](#)
- reporting device ranking [20-7](#)
- reporting device type ranking [20-7](#)
- reports
 - viewing [20-19, 20-25](#)
- reports, view type, CSV [20-24](#)
- reports, view type, recent [20-24](#)
- reports, view type, total [20-24](#)
- report views, CSV [20-24](#)
- report views, peak, reports, view type, peak [20-24](#)
- report views, recent [20-24](#)
- report views, total [20-24](#)
- rules
 - destination IP
 - ANY [21-8](#)
 - devices [21-8](#)

- DISTINCT [21-8](#)
- IP addresses [21-8](#)
- IP ranges [21-8](#)
- Network Groups [21-8](#)
- networks [21-8](#)
- SAME [21-8](#)
- variables [21-8](#)
- device [21-11](#)
 - ANY [21-11](#)
 - Unknown Reporting Device [21-11](#)
 - variables [21-11](#)
- event type grouping [21-10](#)
- event types [21-10](#)
 - ANY [21-10](#)
 - variables [21-10](#)
- reported user
 - ANY [21-11](#)
 - Invalid User Name [21-11](#)
 - NONE [21-11](#)
 - variables [21-11](#)
- service
 - ANY [21-9](#)
 - defined groups [21-10](#)
 - defined services [21-10](#)
 - service variables [21-9](#)
- severity
 - ANY [21-12](#)
 - green [21-12](#)
 - red [21-12](#)
 - yellow [21-12](#)
- source IP
 - devices [21-7](#)
 - IP addresses [21-7](#)
 - IP ranges [21-7](#)
 - Network Groups [21-7](#)
 - networks [21-7](#)
 - variables [21-7](#)
- runtime logging [24-1](#)

S

scheduling
 discovery [2-39](#)

security contexts
 add discovered [4-10](#)
 define reporting options [4-11](#)
 make MARS aware of [4-8](#)

security policies
 objectives of [1-1](#)

security policy
 suggested content [1-1](#)

see CVE [23-2](#)

seed file
 CSV file [2-20](#)
 loading [2-24](#)

service
 adding [23-8](#)
 deleting [23-8](#)
 editing [23-8](#)
 editing groups [23-7](#)

service group
 adding [23-7](#)

service management [23-7](#)

service provider
 adding [22-11, 23-11](#)

services
 adding group [23-7](#)

session count [20-7](#)

setting
 runtime logging levels [24-1](#)

Severity icons [19-3](#)

Short Message Service
 See SMS. [21-15](#)

Simple Network Management Protocol
 See SNMP. [21-15](#)

SNMP RO, unsupported characters [2-9, 2-22, 2-29](#)

Solaris host, bootstrap [10-2](#)

source IP address ranking [20-6](#)

source network group ranking [20-6](#)

source network ranking [20-6](#)

source port ranking [20-6](#)

stacked charts [17-13](#)

static information [19-10](#)

system determined false positive type [19-8](#)

T

table
 incidents [19-5](#)

Time [19-3](#)

time ranges
 incidents [19-4](#)

Topology
 toggle device display [17-12](#)

traffic flows
 identify and enable [1-4, 16-8](#)

troubleshoot,cannot add device [2-19](#)

troubleshoot,cannot re-add device [2-19](#)

tuning
 false positives [19-5, 19-9](#)

U

unconfirmed false positive type [19-8](#)

unknown event report [20-7](#)

use only firing events [20-8](#)

user
 adding [22-10, 23-9](#)
 editing [23-12](#)
 removing [23-12](#)

user confirmed false positive type [19-8](#)

user confirmed positive type [19-8](#)

user group
 adding [23-12](#)

user management [23-8](#)
 roles defined [23-8](#)

V

valid networks [2-38](#)

variables [20-10](#), [20-11](#), [21-7](#), [21-8](#)