Check Point
SOFTWARE TECHNOLOGIES LTD.

15 May 2017

QUALITY OF SERVICE

R80.10

Administration Guide

Check Point
INFINITY

Important Information

Latest Software

We recommend that you install the most recent software release to stay up-to-date with the latest functional improvements, stability fixes, security enhancements and protection against new and evolving attacks.

Check Point R80.10

For more about this release, see the R80.10 home page http://supportcontent.checkpoint.com/solutions?id=sk111841.

Latest Version of this Document

Download the latest version of this document http://supportcontent.checkpoint.com/documentation_download?ID=54837.

To learn more, visit the Check Point Support Center http://supportcenter.checkpoint.com.

Feedback

Check Point is engaged in a continuous effort to improve its documentation.

Please help us by sending your comments mailto:cp_techpub_feedback@checkpoint.com?subject=Feedback on Quality of Service R80.10 Administration Guide.

Searching in Multiple PDFs

To search for text in all the R80.10 PDF documents, download and extract the complete R80.10 documentation package http://downloads.checkpoint.com/dc/download.htm?ID=54846.
Use **Shift-Control-F** in Adobe Reader or Foxit reader.

Revision History

Date	Description
15 May 2017	First release of this document

SmartConsole Toolbars

For a guided tour of SmartConsole, click **What's New** in the left bottom corner of SmartConsole.

Global Toolbar (top left of SmartConsole)

	Description and Keyboard Shortcut
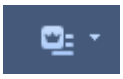	The main SmartConsole Menu
	The **Objects** menu. Also leads to the Object Explorer **Ctrl+E**
	Install policy on managed gateways **Ctrl+Shift+Enter**

Navigation Toolbar (left side of SmartConsole)

	Description and Keyboard Shortcut
	Gateways & Servers configuration view **Ctrl+1**
	Security Policies Access Control view Security Policies Threat Prevention view **Ctrl+2**
	Logs & Monitor view **Ctrl+3**
	Manage & Settings view - review and configure the Security Management Server settings **Ctrl+4**

Command Line Interface Button (left bottom corner of SmartConsole)

	Description and Keyboard Shortcut
	Open a command line interface for management scripting and API **F9**

What's New Button (left bottom corner of SmartConsole)

	Description and Keyboard Shortcut
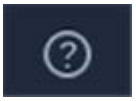	Open a tour of the SmartConsole

Objects and Validations Tabs (right side of SmartConsole)

	Description
Objects	Manage security and network objects
Validations	Validation warnings and errors

System Information Area (bottom of SmartConsole)

	Description
Task List	Management activities, such as policy installation tasks
Server Details	The IP address of the Security Management Server
Connected Users	The administrators that are connected to the Security Management Server

Contents

Terms

Burstiness

Data that is transferred or transmitted in short, uneven spurts. LAN traffic is typically *bursty*. Opposite of streaming data.

CA

Certificate Authority. Issues certificates to gateways, users, or computers, to identify itself to connecting entities with Distinguished Name, public key, and sometimes IP address. After certificate validation, entities can send encrypted data using the public keys in the certificates.

Certificate

An electronic document that uses a digital signature to bind a cryptographic public key to a specific identity. The identity can be an individual, organization, or software entity. The certificate is used to authenticate one identity to another.

Citrix MetaFrame

A client-server software application that enables a client to run a published application on a Citrix server farm from the client's desktop.

Intelligent Queuing Engine

A bandwidth allocation algorithm that guarantees high priority traffic takes precedence over low priority traffic.

Interface

A boundary across which two systems communicate independently with each other.

Jitter

Variation in the delay of received packets. On the sending side, packets are spaced evenly apart and sent in a continuous stream. On the receiving side, the delay between each packet can vary according to network congestion, improper queuing or configuration errors.

Policy

A collection of rules that control network traffic and enforce organization guidelines for data protection and access to resources with packet inspection.

QoS

A policy-based bandwidth management solution.

QoS Action Properties

Properties that define bandwidth allocation, limits, and guarantees for a rule.

RDED

Retransmit Detect Early Drop. The bottleneck that results from the connection of a LAN to the WAN causes TCP to retransmit packets. RDED prevents inefficiencies by detecting retransmits in TCP streams and preventing the transmission of redundant packets when multiple copies of a packet are concurrently queued on the same flow.

Rule

A set of traffic parameters and other conditions that cause specified actions to be taken for a communication session.

Rule Base

The database that contains the rules in a security policy and defines the sequence in which they are enforced.

WFQ

Weighted Fair Queuing. An algorithm to precisely control bandwidth allocation in QoS.

WFRED

Weighted Flow Random Early Drop. A mechanism for managing the packet buffers of QoS. Adjusting automatically and dynamically to the network traffic situation, WFRED remains transparent to the user.

Introduction to QoS

In This Section:

Important

This guide is only for R80.10 and higher.

The Check Point QoS Solution

QoS is a policy based bandwidth management solution that lets you:

- Prioritize business-critical traffic, such as ERP, database and Web services traffic, over lower priority traffic.
- Guarantee bandwidth and control latency for streaming applications, such as Voice over IP (VoIP) and video conferencing.
- Give guaranteed or priority access to specified employees, even if they are remotely accessing network resources.

You deploy QoS with the Security Gateway. QoS is enabled for both encrypted and unencrypted traffic.

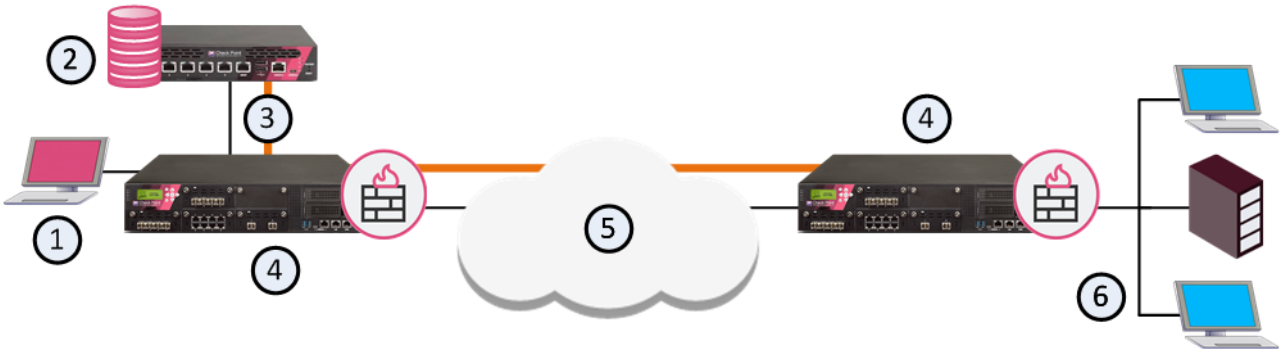

Item	Description
1	SmartConsole
2	Security Management Server
3	QoS Policy
4	Security Gateway with QoS Software Blade
5	Internet
6	Internal network

QoS leverages the industry's most advanced traffic inspection and bandwidth control technologies. Check Point patented Stateful Inspection technology captures and dynamically updates detailed state information on all network traffic. This state information is used to classify traffic by service or application. After traffic has been classified, QoS applies an innovative, hierarchical, Weighted Fair Queuing (WFQ) algorithm to accurately control bandwidth allocation.

Features and Benefits

QoS gives these features and benefits:

- **Flexible QoS policies with weights, limits and guarantees**

  QoS lets you create basic policies that can be modified to include the Advanced QoS features described in this section.

- **Integration with the Security Gateway**

  The integration of an organization's security and bandwidth management policies enables easier policy definition and system configuration. This lets you optimize network performance for VPN and unencrypted traffic

- **Performance analysis**
- Monitor system performance with the Logs & Monitor features in SmartConsole.
- **Integrated DiffServ support**

  Add one or more Diffserv Classes of Service to the QoS Policy Rule Base.

- **Integrated Low Latency Queuing**

  Define special classes of service for "delay sensitive" applications like voice and video to the QoS Policy Rule Base.

- **Integrated Citrix MetaFrame support**

  QoS solution for the Citrix ICA protocol.

- **No need to deploy separate VPN, Firewall and QoS devices**

  QoS and Firewall share a common architecture and many core technology components. User-defined network objects can be used in both solutions.

- **Proactive management of network costs**

  QoS monitoring systems let you to be proactive in managing your network and controlling network costs.

- **Support for end-to-end QoS for IP networks**

  QoS offers full support for end-to-end QoS for IP networks by distributing enforcement throughout network hardware and software.

- **CoreXL and SecureXL support**

  Packet acceleration.

QoS Policy Types

R80.10 includes two QoS Policy types:

- **Express** - Quickly create basic QoS Policies
- **Recommended** - Create advanced Policies with the full set of QoS features

This table shows the difference between the **Recommended** and **Express** policy types.

Features	Recommended	Express	To learn more
Weights	✓	✓	Weight (on page 26)
Limits (whole rule)	✓	✓	Limits (on page 27)
Authenticated QoS	✓*		Authenticated QoS (on page 52)
Logging	✓	✓	Overview of Logging (on page 84)
Accounting	✓*	✓	
Support for UTM-1 Edge Gateways		✓	
Support for hardware acceleration	✓		
High Availability and Load Sharing	✓	✓	
Guarantees (Per connection)	✓		Guarantees (see "Per Connections Guarantees" on page 44)
Limits (Per connection)	✓		Limits (on page 27)
LLQ (controlling packet delay in QoS)	✓		Low Latency Queuing (on page 46)
DiffServ	✓		Differentiated Services (DiffServ) (on page 45)
Sub-rules	✓		
Matching by URI resources	✓		
Matching by DNS string	✓		
Matching Citrix ICA Applications	✓*		
SecureXL support	✓		

Features	Recommended	Express	To learn more
CoreXL support	✓		
SmartLSM clusters	✓		

* You must disable SecureXL and CoreXL before you can use this feature.

To select a QoS Policy type:

1. In SmartConsole menu, click **Manage policies and layers**.
2. In the Manage Policies window, click **New** or select an existing Policy and then click **Edit**.
3. Select **QoS**, and then select **Recommended** or **Express**.

Acceleration Support for R77 Policies

After a clean install or upgrade to R80.10, QoS supports SecureXL and CoreXL acceleration technologies.

Important: After a clean install or upgrade, SecureXL and CoreXL are enabled by default. If you have a QoS policy created for R77 and earlier, these features are not supported when acceleration is enabled:

- User Authority Server
- IPSO
- Citrix printing rules
- Security Gateways below R77.10
- SmartView Monitor - QoS views do not correctly show traffic accelerated by SecureXL

To use these features you must disable QoS ("Disabling QoS Acceleration Support" on page 104).

Workflow

This topic shows a high-level workflow for creating an effective QoS Policy.

Do these steps in SmartConsole:

1. Enable QoS for each applicable Security Gateway.
2. Configure QoS Global Properties.
3. Create or change a QoS Policy ("Working with QoS Policies" on page 62).
4. Configure log collection and system monitoring for QoS.
5. Publish the changes.

Do these steps in SmartDashboard:

1. Define the gateway networks, services and other related objects.
2. Define QoS rules ("Working with Rules" on page 65) (basic and advanced).
3. Configure specialized QoS features:
   a) Differentiated Services (DiffServ) (see "Working with Differentiated Services (DiffServ)" on page 74)
   b) Low Latency Queuing (see "Working with Low Latency Queuing" on page 77)

c) Authenticated QoS (see "Working with Authenticated QoS" on page 78)

d) Citrix ICA Applications (see "Managing QoS for Citrix ICA Applications" on page 79)

Go back to SmartConsole to do these steps:

1. Publish the changes.
2. Install Policy. See Implementing the Rule Base (see "Using Policies" on page 29).

QoS Deployment

In This Section:

Deploying QoS

This section covers topology restrictions.

QoS Topology Restrictions

QoS can manage up to the maximum number of external interfaces supported by Firewall, subject to the following restrictions. (Refer to the Firewall documentation for further information):

1. All of the traffic on a managed line must go through the gateway.
2. Each managed line must be connected (directly or indirectly via a router) to a separate physical interface on the QoS machine. Two managed lines cannot share a physical interface to the QoS gateway, and two network segments cannot connect to the same router.

For example, in the configuration depicted in the following diagram, the routers can pass traffic to each other through the hub without the QoS gateway being aware of the traffic.

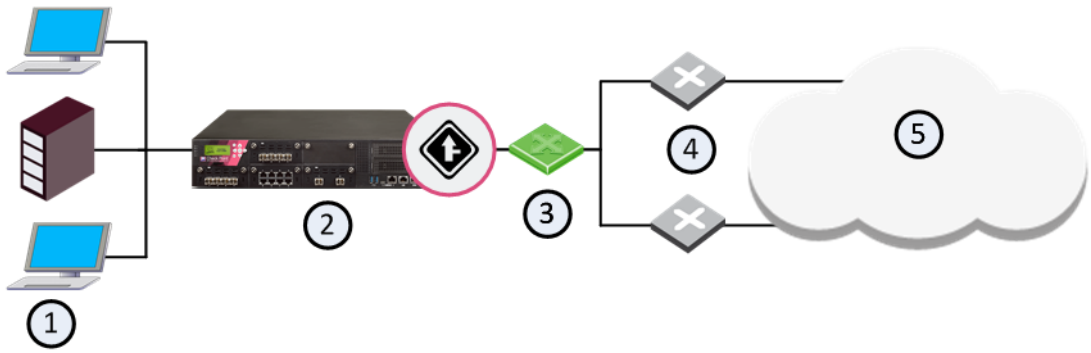

Item	Description
1	Internal network
2	Security Gateway with QoS enabled
3	Hub
4	Routers
5	Internet

You cannot manage two networks connected to a single router since traffic may pass from one line to the other directly through the router, without the QoS gateway being aware of the traffic:

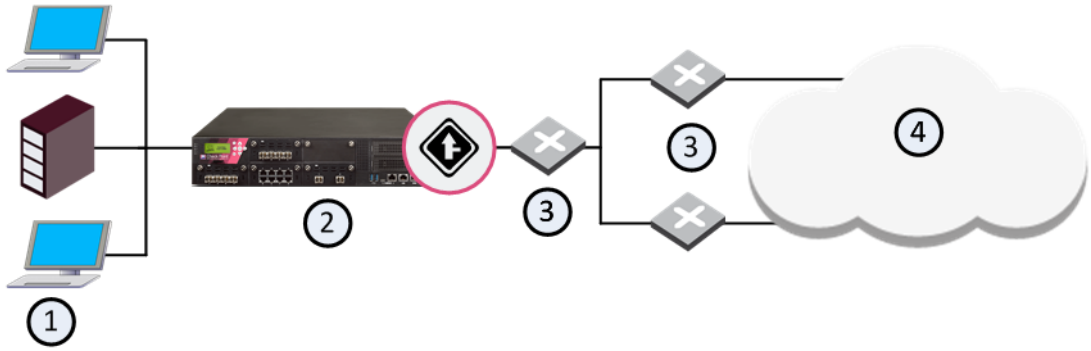

Item	Description
1	Internal network
2	Security Gateway with QoS enabled
3	Routers
4	Internet

In a correct configuration, the routers connect directly to the QoS gateway.

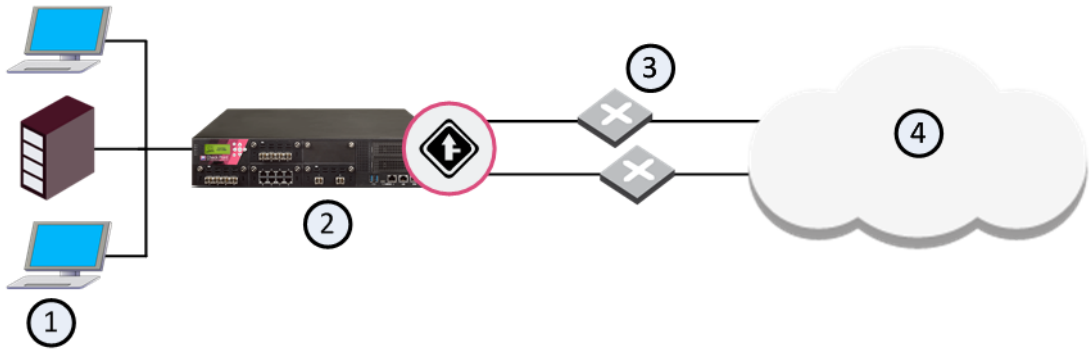

Sample Bandwidth Allocations

Frame Relay Network

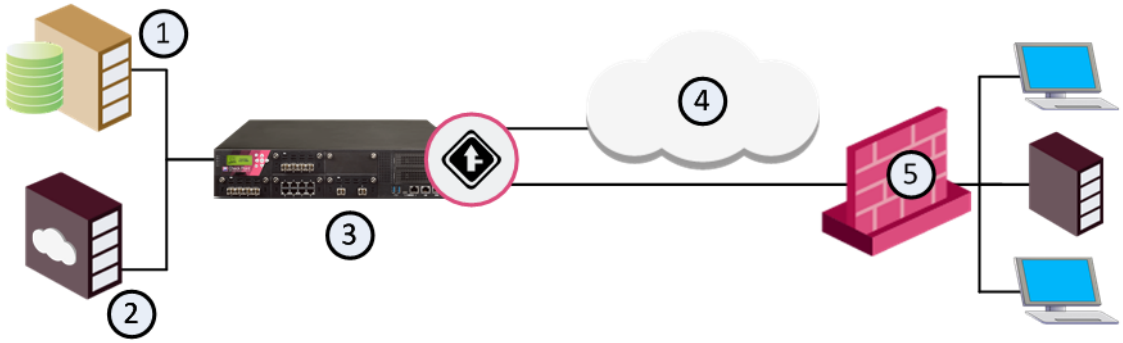

Item	Description
1	Database server
2	Web server
3	Security Gateway with QoS enabled
4	Internet
5	Branch office

The previous diagram shows that the branch office communicates with the central site and the opposite. It only communicates directly with the Internet through the central site. The Web server makes important company documents available to the branch office, and the database server supports the company's mission-critical applications.

The problem is that most of the branch office traffic is internal and external Web traffic, and the mission-critical database traffic suffers as a result. The network administrator has considered upgrading the 56K lines, but is reluctant to do so, not only because of the cost but also because upgrading would probably not solve the problem. The upgraded lines would still be filled mostly with Web traffic.

The goals are as follows:

1. Allocate the existing bandwidth so that access to the database server gets the largest share.
2. Take into account that the branch offices are connected to the network by 56K lines.

These goals are accomplished with the following Rule Base:

Main Rules

Rule Name	Source	Destination	Service	Action
Office 1	Office 1	Any	Any	Weight 10 Limit 56KBps
Office n	Office n	Any	Any	Weight 10 Limit 56KBps
Default	Any	Any	Any	Weight 10

Each office has sub-rules, as follows:

Office Sub-Rules

Rule Name	Source	Destination	Service	Action
Start of Sub-Rule				
Database Rule	Any	Database server	Database service	Weight 50
Web Rule	Any	Web Server	http	Weight 10

Rule Name	Source	Destination	Service	Action
Branch Offices	Any	Any	Any	Weight 10
End of Sub Rule				

The sub-rules give database traffic priority over Web traffic and other traffic.

Assumptions

The following assumptions are made in this example:

- The problem (and its solution) apply to traffic outbound from the central site.

  Note that QoS shapes the branch office lines in the outbound direction only. QoS shapes inbound traffic only on directly controlled interfaces (that is, interfaces of the QoS machine).
- The central site has the capacity to handle the network's peak traffic load.
- There is no traffic between the offices.

QoS Architecture

Basic Architecture

The architecture and flow control of QoS is similar to firewall.
QoS has three components:

- SmartConsole
- Security Management Server
- Gateway

The components can be installed on one machine or in a distributed configuration on a number of machines.

Bandwidth policy is configured using SmartConsole. On the Security Management Server, the policy is verified and installed on the QoS gateways. The QoS Security Gateway uses:

- The firewall chaining mechanism to receive, process and send packets.
- A proprietary classifying and rule-matching infrastructure to examine a packet.

Logging information is created using the firewall kernel API.

The QoS Blade

The primary role of the QoS blade is to:

- Implement a QoS policy at network access points
- Control the flow of inbound and outbound traffic.

QoS has two components:

- QoS kernel driver
- QoS daemon

QoS Kernel Driver

The kernel driver is the heart of QoS operations. It is in the kernel driver that IP packets are examined, queued, scheduled and released, enabling QoS traffic control abilities.

QoS Daemon (fgd50)

The QoS daemon is a user mode process that:

- Resolves DNS for the kernel (used for Rule Base matching).
- Resolves Authenticated Data for an IP (using UserAuthority - again for Rule Base matching).
- In a Cluster Load Sharing configuration, updates the kernel of changes in the cluster status. For example, if a cluster member goes down. The daemon recalculates the relative loads of the gateways and updates the kernel.

QoS SmartConsole

You use the R80.10 SmartConsole and SmartDashboard to create "bandwidth rules" for the QoS policy. Use the Logs & Monitor features in SmartConsole to see information about the active QoS Security Gateways and their Policies.

QoS R80.10 SmartConsole

Use R80.10 SmartConsole to create and change QoS Policies. Use SmartDashboard to work with rules, together with their related network objects and services.

The QoS Policy rules are shown the QoS Rule Base.

QoS Configuration

The Security Management Server and the QoS Security Gateway can be installed on the same machine or on two different machines. When they are installed on different machines, the configuration is known as distributed.

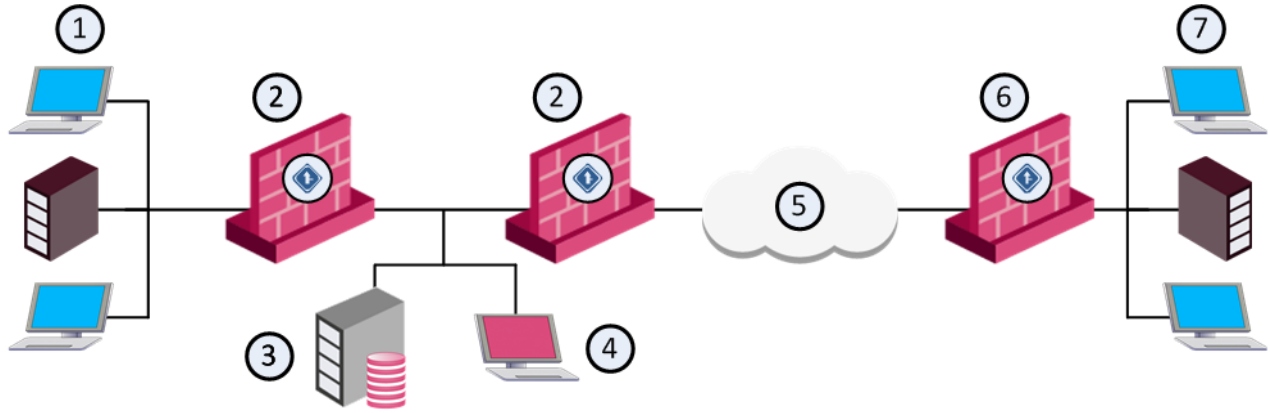

Item	Description
1	Internal network (main office)
2	Security Gateway with QoS enabled
3	Security Management Server
4	SmartConsole
5	Internet

Item	Description
6	Security Gateway with QoS enabled (branch office)
7	Internal network (branch office)

The example shows a distributed configuration, in which one Security Management Server (consisting of a Security Management Server and a SmartConsole controls four QoS gateways. The four QoS gateways manage bandwidth allocation on three QoS enabled lines.

One Security Management Server can control and monitor multiple QoS gateways. The QoS Security Gateway operates independently of the Security Management Server. QoS gateways can operate on more Internet gateways and interdepartmental gateways.

Client-Server Interaction

SmartConsole and the Security Management Server can be installed on the same machine or on two different machines. When they are installed on two different machines, QoS implements the Client/Server model, in which a SmartConsole controls a Security Management Server.

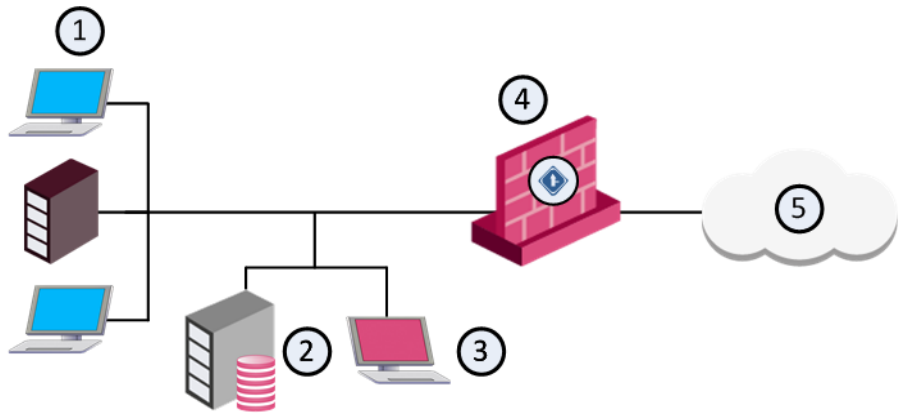

Item	Description
1	Internal network (main office)
2	Security Management Server
3	SmartConsole
4	Security Gateway with QoS enabled
5	Internet

In the configuration depicted in the above figure, the functionality of the Security Management Server is divided between two workstations (Tower and Bridge). The Security Management Server with the database is on Tower. The SmartConsole is on Bridge.

The user, working on Bridge, maintains the QoS Policy and database, which reside on Tower. The QoS Security Gateway on London enforces the QoS Policy on the QoS enabled line.

The Security Management Server is started with the `cpstart` command, and must be running if you wish to use the SmartConsole on one of the client machines.

A SmartConsole can manage the Server only if both the administrator logged into SmartConsole and the computer on which the SmartConsole is running have been authorized to access the

Security Management Server. Use *cpconfig* to:

- Add SmartConsole as GUI client authorized to access the Security Management Server
- Define administrators for the Security Management Server.

Concurrent Sessions

More than one administrator can work with QoS Policies at the same time, each in a different session. A locking mechanism prevents administrators from working on the same object at one time. After you complete you work in a session, click Publish to make your changes available to other sessions and administrators.

Interaction with VPN

Interoperability

QoS and firewall share many core technology components. The same user-defined network objects can be used in both solutions. The integration of an organization's security and bandwidth management policies gives easy policy definition and system configuration. For efficient traffic inspection and enhanced performance, the blades share state table information. The QoS blade and firewall blade let users define bandwidth allocation rules for encrypted and NATed traffic.

Security Management Server

QoS uses the Security Management Server and shares the objects database (network objects, services and resources) with the firewall. Some objects have properties that are product specific. For example, the Firewall has encryption properties which are not related to QoS. A QoS network interface has speed properties that are not related to the firewall.

Basic Policy Management

In This Section:

This section covers basic policy management.

Overview

This chapter describes the basic QoS Policy management that is required to enable you to define and implement a working QoS Rule Base. More advanced QoS Policy management features are discussed in Advanced QoS Policy Management (on page 42).

Rule Base Management

Opening the GUI Clients

To open SmartConsole, click SmartConsole in the Windows Start menu.

SmartDashboard opens automatically when you open an existing QoS Policy, or after you create a new QoS Policy. It is generally not necessary to open SmartDashboard manually.

To open SmartDashboard manually:

1. In SmartConsole, open a QoS Policy:
2. Click **Security Policies** > **Access Control** > **QoS**.
3. In the **QoS** view, click **Open QoS Policy in SmartDashboard**.

SmartDashboard opens and the QoS view shows.

Overview

QoS policy is implemented by defining a set of rules in the Rule Base. The Rule Base specifies what actions are to be taken with the data packets. The Rule Base specifies:

- Source and destination of the traffic
- Services that can be used
- Times
- Logging and logging level

The Rule Base comprises the rules you create and a default rule (see Default Rule (on page 27)). The default rule is automatically created with the Rule Base. It can be modified but cannot be deleted. Unless other rules apply, the default rule is applied to all data packets. The default rule is therefore always the last rule in the Rule Base.

It is a best practice to create your QoS rules based on actual traffic patterns. Use the Logs & Monitor features in SmartConsole to analyze traffic logs.

QoS inspects packets in a sequential manner. When QoS receives a packet for a connection, it compares it against the first rule in the Rule Base. Then against the second, then the third. When QoS finds a rule that matches, it stops checking and applies that rule.

If the matching rule has sub-rules the packets are then compared against the first sub-rule. Then the second, third, and other sub-rules until it finds a match.

If the packet fails to match a rule or sub-rule, the default rule or default sub-rule is applied. The first rule that matches is applied to the packet, not the rule that best matches.

After you have defined your network objects, services and resources, you can use them in building a Rule Base. For instructions on building a Rule Base, see Editing QoS Rules (see "Working with QoS Policies" on page 62).

Note - It is best to organize lists of objects (network objects and services) into groups. Using groups gives you a better overview of your QoS Policy and leads to a more readable Rule Base. New objects added to groups are automatically included in the rules.

Connection Classification

A connection is classified according to four criteria:

- **Source**

  A set of network objects such as specified computers, networks, user groups or domains.

- **Destination**

  A set of network objects such as specified computers, networks, user groups or domains.

- **Service**

  A set of IP services, TCP, UDP, ICMP or URLs.

- **Time**

  Specified days or time periods.

Network Objects

The network objects that can be used in QoS rules include workstations, networks, domains, and groups.

User Groups

QoS lets you define Groups of predefined users. For example, all the users in the marketing department can be grouped together in a User Group called Marketing. When defining a rule, you can use this group as the **Source** instead of adding individual users to the **Source** column of the rule.

Services and Resources

QoS allows you to define QoS rules, not only based on the source and destination of each communication, but also according to the service requested. The services that can be used in QoS rules include TCP, Compound TCP, UDP, ICMP and Citrix TCP services, IP services

Resources can also be used in a QoS Rule Base. They must be of type **URI for QoS**.

Time Objects

QoS allows you to define Time objects. Time objects are used to specify when a rule is enforced. Time objects can be defined for specified times or days. Days can be divided into days of the month or days of the week.

Bandwidth Allocation and Rules

A rule can specify three factors to be applied to bandwidth allocation for classified connections:

Weight

Weight is the percentage of the available bandwidth allocated to a rule. This is not the same as the *weight* in the QoS Rule Base, which is a manually assigned priority.

To calculate what percentage of the bandwidth the connections matched to a rule receives:

```
                        Priority in SmartDashboard
The weight = ------------------------------------------------------
             Total priority of all the rules with open connections
```

For example:

- if this rule's weight (priority in SmartDashboard) is 12
- the total weight (priority in SmartDashboard) of all the rules for which connections are currently open is 120

Then all the connections open under this rule are allocated 12/120, or 10%. The weight of this rule is 10%. The rule gets 10% of the available bandwidth if the rule is active. In practice, if other rules are not using their maximum allocated bandwidth, a rule can get more than the bandwidth allocated by this formula. Unless a per connection limit or guarantee is defined for a rule, all connections under a rule receive equal weight.

Allocating bandwidth according to weights ensures full use of the line even if a specified class is not using all of its bandwidth. In such a case, the left over bandwidth is divided between the remaining classes in accordance with their relative weights. Units are configurable, see Defining QoS Global Properties (on page 59).

Guarantees

A guarantee allocates a minimum bandwidth to the connections matched with a rule.

Guarantees can be defined for:

- The sum of all connections in a rule

  A total rule guarantee reserves a minimum bandwidth for all the connections below a rule. The actual bandwidth allocated to each connection depends on the number of open connections that match the rule. The total bandwidth allocated to the rule cannot be less than the guarantee. The more connections that are open, the less bandwidth each connection receives.

- Individual connections in a rule

A per-connection guarantee means that each connection that matches the specified rule is guaranteed a minimum bandwidth.

Note: Although weights guarantee the bandwidth share for specified connections, only a guarantee lets you to specify an absolute bandwidth value.

Limits

A limit specifies the maximum bandwidth that is assigned to all the connections together. A limit defines a point after which connections below a rule are not allocated more bandwidth, even if there is surplus bandwidth available.

Limits can also be defined for the sum of all connections in a rule or for individual connections within a rule.

For more information on weights, guarantees and limits, see Action Type (on page 27).

Note - Bandwidth allocation is not fixed. As connections are opened and closed, QoS continuously changes the bandwidth allocation to accommodate competing connections, in accordance with the QoS Policy.

Default Rule

A default rule is automatically added to each QoS Policy Rule Base, and assigned the weight specified in the **QoS** page of the **Global Properties** window. You can change the weight, but you cannot delete the default rule (see Weight (on page 26)).

The default rule applies to all connections not matched by the other rules or sub-rules in the Rule Base.

A default rule is automatically added to each group of sub-rules, and applies to connections not classified by the other sub-rules in the group. For more, see: To Verify and View the QoS Policy.

QoS Action Properties

In the **QoS Action Properties** window you can define bandwidth allocation properties, limits and guarantees for a rule.

Action Type

These are the two types of QoS actions:

Action Type	Recommended	Express
Simple	Yes	Yes
Advanced	Yes	No

Simple

The Simple action type has these action properties:

- Apply rule only to encrypted traffic
- Rule weight
- Rule limit
- Rule guarantee

Advanced

The Advanced rule type has these properties:

- Per rule
- Per connection
- Per rule guarantee
- Per connection guarantee
- Number of permanent connections
- Accept additional connections

Example of a Rule Matching VPN Traffic

VPN traffic is traffic that is encrypted by the Security Gateway. VPN traffic does not refer to traffic that was encrypted by a non-Check Point product prior to arriving at this Security Gateway. This type of traffic can be matched using the IPSec service.

When **Apply rule only to encrypted traffic** is selected in the **QoS Action Properties** window, only VPN traffic is matched to the rule. If this field is not checked, all types of traffic (both VPN and non-VPN) are matched to the rule.

Use the **Apply rule only to encrypted traffic** option to create a Rule Base that applies only to VPN traffic. These actions are different from actions applied to non-VPN traffic. Since QoS uses the First Rule Match concept, the VPN traffic rules must be defined as the top rules in the Rule Base. Below them define rules that apply to all other types of traffic. Other types of traffic skip the top rules and match to one of the non-VPN rules. To separate VPN traffic from non-VPN traffic, define this rule at the top of the QoS Rule Base:

Name	Source	Destination	Service	Action
VPN rule	Any	Any	Any	VPN Encrypt, and other configured actions

All the VPN traffic is matched to this rule. The rules below this VPN Traffic Rule are then checked only against non-VPN traffic. You can define sub-rules below the VPN Traffic rule that classify the VPN traffic more granularly.

Bandwidth Allocation and Sub-Rules

When a connection is matched to a rule with sub-rules, the sub-rules are checked for match. If none of the sub-rules apply, the default rule for the sub-rules is applied (see Default Rule (on page 27)).

Sub-rules can be nested, meaning that sub-rules themselves can have sub-rules. The same rules then apply to the nested sub-rules. If the connection matches a sub-rule that has sub-rules, the nested sub-rules are checked for a match. If none of the nested sub-rules apply, the default rule for the nested sub-rules is applied.

Bandwidth is allocated on a top/down basis. This means that:

- Sub-rules cannot give more bandwidth to a matching rule, than the rule in which the sub-rule is located.
- A nested sub-rule cannot give more bandwidth than the sub-rule in which it is located.

A Rule Guarantee must always be greater than or equal to the Rule Guarantee of a sub-rule in that rule. The same applies to Rule Guarantees in sub-rules and their nested sub-rules.

Example:

Bandwidth Allocation in Nested Sub-Rules:

Rule Name	Source	Destination	Service	Action
Rule A	Any	Any	ftp	Rule Guarantee - 100KBps Weight 10
Start of Sub-Rule A				
Rule A 1	Client-1	Any	ftp	Rule Guarantee - 100KBps Weight 10
Start of Sub-Rule A1				
Rule A1.1	Any	Any	ftp	Rule Guarantee - 80KBps Weight 10
Rule A1.2	Any	Any	ftp	Weight 10
End of sub-rule A1				
RuleA2	Client-1	Any	ftp	Weight 10
End of sub-rule A				
Rule B	Any	Any	http	Weight30

In this example, surplus bandwidth from the application of Rule A1.1 is applied to Rule A2 before it is applied to Rule A1.2.

Using Policies

After you define your QoS rules in the Rule Base, you must publish your session in SmartConsole, and then install the Policies on your Security Gateways. The policy installation procedure automatically validates the rules and objects. If there verification errors, a message shows in the in the **Install Policy Details** tab.

After policy installs successfully, the Security Gateways enforce the policy rules.

Note - Make sure the QoS blade is enabled on the Security Gateway before you install the policy.

Installing a QoS Policy

To install a QoS Policy:

1. In SmartDashboard, make changes to Policy rules and then click **Update**.
2. In SmartConsole, click **Install Policy**.
3. From the **Policy list**, select the policy to install.

4. Click **Policy Targets** and select the Security Gateways that will get this Policy.

   Note - By default, no gateways are selected for QoS. You must select them manually.

5. Click **Install**.

If the installation is successful, the new Policy is enforced by the Security Gateways on which it is installed. If installation fails, do these steps to see the error messages:

1. Click the Task Information area, in the lower, left hand corner of SmartConsole.
2. In the **Recent Tasks** area, click **Details** on the applicable error.

In the **Install Policy Details** window, click the **^** icon in the **Status** column to see the error messages. You must resolve all errors before you can successfully install the Policy.

QoS Tutorial

In This Section:

This chapter includes a step by step guide for creating a sample deployment with a QoS Policy. We recommend that you have a working knowledge of these Check Point products and concepts to use this tutorial effectively:

- Security Gateways and management servers
- Security Policies and the Rule Base
- SmartConsole and SmartDashboard
- Firewall and related Software Blades

Deployment Scenario for this Tutorial

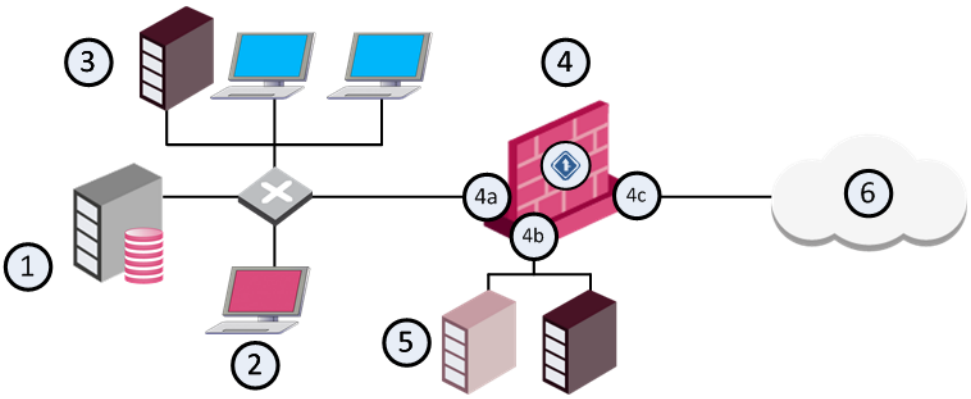

Item	Description
1	Oxford - Security Management Server
2	Cambridge - SmartConsole client
3	Local area network - Engineering and Marketing
4	London - Security Gateway with QoS
4a	Interface eth2 - 199.199.199.32
4b	Interface eth1 - 199.199.199.32
4c	Interface eth0 - 199.199.199.32
5	DMZ with Web and FTP servers
6	Internet

This scenario is an organization with offices located in London, Oxford and Cambridge. The QoS Security Gateway is in London and has three interfaces, one of which is connected to the Internet. The Security Management Server is in Oxford and the SmartConsole is in Cambridge. The local network includes the Marketing and Engineering departments.

Tutorial Workflow

This tutorial is a simplified exercise that shows you how to do these QoS activities:

1. Install and configure the system components.
2. Create a new QoS Policy with SmartConsole.
3. Select one of these QoS Policies types:
   - **Express -** Quickly create basic QoS Policies.
   - **Recommended** - Create advanced Policies with the full set of QoS features.
4. Configure the network objects used by QoS rules.
5. Configure specialized services for use in QoS rules.
6. Create QoS Policy rules.
7. Install the Policy on the Security Gateway.

Installing the System Components

To install and configure system components for this tutorial:

1. Enable QoS, Firewall, and other Software Blades on the London Security Gateway.
2. Install R80.10 Security Management Server on the Oxford server platform.
3. Install SmartConsole on the Cambridge PC.
4. In SmartConsole, define Cambridge as a trusted client.
5. In SmartConsole, define the administrators who can manage the QoS Policy.
6. Make sure that there is SIC trust between the Oxford Security Management Server and the London QoS Security Gateway.

Starting R80.10 SmartConsole

This section describes how to open SmartDashboard and access the QoS tab.

Creating a New QoS Policy

To create a New Policy:

1. On the gateway, make sure that the QoS blade is enabled.
2. In SmartConsole, from the **File** menu, select **Manage Policies and Layers**.
3. Click **New** .
4. In the **Policy** window, enter a Policy name.

This name cannot:

- Contain any reserved words or spaces.
- Start with a number.

- Contain any of the following characters: %, #, ', &, *, !, @, ?, <, >, /, \, :.
- End with any of the following suffixes: .w, .pf, .W.

5. Select **QoS** and then select a QoS Policy type:
   - **Express** - Quickly create basic QoS Policies
   - **Recommended** (default) - Create advanced Policies with the full set of QoS features

   Note: There are some limitations that can prevent you from enabling SecureXL or CoreXL with QoS Policies.

   For more, see: QoS Policy limitations ("Acceleration Support for R77 Policies" on page 15).
6. Click **OK**.

   The system saves the new Policy and SmartDashboard opens automatically. You can start to define your rules here.

Planning the QoS Policy

To implement a good QoS Policy, find out how the network is used. Identify and prioritize the types of traffic. Identify users and their needs. For example:

- HTTP traffic must be allocated more bandwidth than RealAudio.
- Marketing must be allocated more bandwidth than Engineering.

Configuring the Security Gateway

Define these Network Objects:

- London, the Security Gateway on which the QoS is enabled
- Sub-networks for the Marketing and Engineering departments

To define the London Security Gateway:

1. In SmartConsole, click Gateways & Servers.
2. Click **New** > **Gateway** > **Classic Mode**.
3. Configure these parameters in the **General Properties** window.

Field	Value	Notes
Name	London	This is the name by which the object is known on the network; the response to the **hostname** command.
Platform	Select an appliance type or **Open Server**	The platform must be supported for R80.10.
SIC	Click **Communication**	Establishes a secure communication channel between the Security Gateway and the management server.
Version	R80	
OS	Gaia	

Field	Value	Notes
IP Address	192.32.32.32	This is the interface associated with the host name in the DNS — get this by clicking **Get Address**. For gateways, this should always be the IP address of the external interface.
Network Security Tab	**Firewall** and **QoS**	

Defining Interfaces on the Gateway

In this step you configure each interface and its QoS properties.

To configure interface properties:

1. Click **Network Management** in the navigation tree.
2. Click **Get Interfaces** on the toolbar.

   The interfaces show in the **Network Management** window.
3. Double-click each interface and configure parameters in the **Interface > General** window.

eth0

Field	Value	Notes
Net Address	192.32.32.32	
Net Mask	255.255.255.0	
Topology Settings (Click **Modify**)	**Internet External**	This interface connects to the Internet.
Anti-Spoofing	**Perform Anti-Spoofing based on interface topology**	Each incoming packet is examined to make sure that the source IP address is valid.
Spoof Tracking	**Log**	Log Anti-Spoofing events.

eth1

Field	Value	Notes
Net Address	192.32.42.32	
Net Mask	255.255.255.0	
Topology Settings (Click **Modify**)	**Internet External**	This interface connects to the Internet.
Anti-Spoofing	**Perform Anti-Spoofing based on interface topology**	Each incoming packet is examined to make sure that the source IP address is valid.
Spoof Tracking	**Log**	Log Anti-Spoofing events.

eth2

Field	Value	Notes
Net Address	192.199.199.32	
Net Mask	255.255.255.0	
Topology Settings (Click **Modify**)	**Internet External**	This interface connects to the Internet.
Anti-Spoofing	**Perform Anti-Spoofing based on interface topology**	Each incoming packet is examined to make sure that the source IP address is valid.
Spoof Tracking	**Log**	Log Anti-Spoofing events.

To configure interface QoS properties:

1. In the **Interface** window, click the **QoS** tab.
2. Select **Inbound Active** and **Outbound Active**.
3. Set **Inbound Active** and **Outbound Active** to **192000 - T1 (1.5 Mbps)**.

Configuring QoS Properties for Interfaces

Defining the Services

The QoS Policy required for this tutorial does not require the definition of new proprietary services. The commonly used services HTTP and RealAudio are already defined in QoS.

Creating and Configuring Rules

After you define your network objects and services, the next step is to create your QoS policy rules. This tutorial shows you how to create two simple QoS rules. A new QoS Policy always includes a Default Rule (see Default Rule (on page 27)).

To Create a New Policy

1. In **SmartConsole** select **New** from the **File** menu.

   The **New Policy** window opens.
2. Enter the name in the **New policy Package Name** field.
3. Select **QoS**.
4. Select **QoS policy (recommended)**.
5. Click **OK**.

   The new **Policy** is created together with a **Default Rule** and is displayed in the **QoS** tab.

Creating New Rules

When you create a new QoS Policy, the system automatically adds a *default rule*, which must always be the last rule in the Policy. Make sure that you add your new rules above the default rule.

Create these two rules: *Web Rule* and *RealAudio Rule*.

1. In **SmartDashboard** > **QoS** tab, select the default rule.
2. Click the **Before current rule** icon.

3. Enter *Web Rule* in the **Rule Name** window, and then click **OK**.

Do this procedure again for *RealAudio Rule*.

Rule Properties

A new rule has the default values assigned by the administrator. The next procedure describes how to change these rules to the values shown in the table below.

Changing Rules Default Values

Rule Name	Source	Destination	Service	Action
Web Rule	Any	Any	HTTP	Weight 35
RealAudio Rule	Any	Any	RealAudio	Weight 5
Default	Any	Any	Any	Weight 10

Changing New Rule Properties

The system automatically assigns the default parameters as defined in the **Global Properties** > **QoS** to new rules. Use this procedure to change these rules to the values shown in the table below.

Rule Name	Source	Destination	Service	Action
Web Rule	Any	Any	HTTP	Weight 35
RealAudio Rule	Any	Any	RealAudio	Weight 5
Default	Any	Any	Any	Weight 10

To change the properties in a rule:

1. In the **QoS** tab, right-click in the **Service** field of the Web Rule.
   Select **Add Objects**, and then select **HTTP** from the list.
2. Double-click the **Action** field, and then change the **Rule Weight** (see "Defining QoS Global Properties" on page 59) property to 35.

Do this procedure again for the *RealAudio* and *Default* rules.

Classifying Traffic by Service

Usually, a full Rule Base will not explicitly define rules for all the "background" services (such as DNS and ARP). Background services are handled by the **Default** rule.

The structure of the Rule Base is shown at the left of the window as a tree, with the **Default Rule** at the bottom. (For a description of the Rule Base window, see Basic Policy Management (on page 24)).

Connections receive bandwidth according to the weights (priority) assigned to the rules that apply to them. The table below describes what occurs when there are four active connections. Note that bandwidth allocation is constantly changing.

Service Rules - Four Active Connections

Connections	Relevant rule	Bandwidth	Comments
HTTP	**Web Rule**	70%	35 / 50 (the total weights)
RealAudio	**RealAudio Rule**	10%	5 / 50
FTP	**Default**	sharing 20%	10 /50; a rule applies to all the connections together
TELNET	**Default**	sharing 20%	10 /50; a rule applies to all the connections together

Bandwidth is allocated between connections according to relative weight. As connections are opened and closed, QoS changes the bandwidth allocation according to the QoS Policy.

For example:

- If the HTTP, FTP and TELNET connections are all closed. The only remaining connection is the RealAudio connection. RealAudio receives 100% of the bandwidth.
- If the TELNET and FTP connections are closed, both HTTP and RealAudio benefit from the released bandwidth.

Service Rules - Two Active Connections

Connections	Relevant rule	Bandwidth	Comments
HTTP	**Web Rule**	87/5%	35 / 40 (the total weights)
RealAudio	**RealAudio Rule**	12.5%	5 / 40

Although RealAudio is assigned a very small weight compared to HTTP, it will not be starved of bandwidth no matter how heavy the HTTP traffic.

In practice, you will probably want to give a high relative weight to interactive services such as TELNET, which transfers small amounts of data but involves users issuing commands.

Classifying Traffic by Source

The second part of the QoS Policy (Marketing must be allocated more bandwidth than Engineering (see "Planning the QoS Policy" on page 33)) is implemented by these rules:

Marketing is Allocated More Bandwidth Than Engineering

Rule Name	Source	Destination	Service	Action
Marketing Rule	Marketing	Any	Any	Weight 30
Engineering Rule	Engineering	Any	Any	Weight 20
Default	Any	Any	Any	Weight 10

Using the same principles described in To Create a New Rules (see "Creating New Rules" on page 35) and To Modify New Rules (see "Changing New Rule Properties" on page 36), create new rules in SmartConsole and change them to match the values shown in the table above. The effect of these rules is equivalent to the rules shown here:

Connections	Relevant rule	Bandwidth	Comments
HTTP	**Web Rule**	70%	35 / 50 (the total weights)
RealAudio	**RealAudio Rule**	10%	5 / 50
FTP	**Default**	sharing 20%	10 /50 A rule applies to all the connections together
TELNET	**Default**	sharing 20%	10 /50 A rule applies to all the connections together

Except for:

- the different weights
- the fact that allocation is based on source rather than on services

Classifying Traffic by Service and Source

The table below shows all the rules in one Rule Base.

All the Rules Together

Rule Name	Source	Destination	Service	Action
Web Rule	Any	Any	HTTP	Weight 35
RealAudio Rule	Any	Any	RealAudio	Weight 5
Marketing Rule	Marketing	Any	Any	Weight 30
Engineering Rule	Engineering	Any	Any	Weight 20
Default	Any	Any	Any	Weight 10

In this Rule Base, bandwidth allocation is based both on sub-networks and on services.

First Rule Match Principle

In the Rule Base shown below:

Rule Name	Source	Destination	Service	Action
Web Rule	Any	Any	HTTP	Weight 35
RealAudio Rule	Any	Any	RealAudio	Weight 5
Marketing Rule	Marketing	Any	Any	Weight 30

Rule Name	Source	Destination	Service	Action
Engineering Rule	Engineering	Any	Any	Weight 20
Default	Any	Any	Any	Weight 10

In a production environment, a connection can match more than one rule. QoS works according to a first rule match principle. Each connection is examined against the QoS Policy and receives bandwidth according to the *Action* defined in the first rule that is matched.

If a user in Marketing initiates an HTTP connection, the connection matches the Web Rule and the Marketing Rule. The Web Rule comes before the Marketing Rule in the Rule Base, so the connection is matched to the Web Rule and given a weight of 35.

To differentiate HTTP traffic by source, create sub-rules for the Web Rule. See Sub-Rules (on page 40).

Guarantees and Limits

Bandwidth allocation can also be defined using guarantees and limits. You can define guarantees and limits for rules or for individual connections in a rule.

Rule Name	Source	Destination	Service	Action
Web Rule	Any	Any	HTTP	Weight 35
RealAudio Rule	Any	Any	RealAudio	Weight 5
Marketing Rule	Marketing	Any	Any	Weight 30
Engineering Rule	Engineering	Any	Any	Weight 20
Default	Any	Any	Any	Weight 10

The Web Rule shown in the Rule Base allocates 35% of available bandwidth to all the HTTP connections combined. The actual bandwidth allocated to connections that match this rule depends on:

- Total available bandwidth
- Open connections that match other rules

Note: 35% of available bandwidth (specified in the example above) is assured to Web Rule. Web Rule will get more bandwidth if there are fewer connections matched to other rules, but never less than 35%.

As an alternative to relative weights, a guarantee can be used to specify bandwidth as an absolute value (in Bytes per second). In this table, Web Rule is guaranteed 20 KBps:

Guarantee Example

Rule Name	Source	Destination	Service	Action
Web Rule	Any	Any	HTTP	Guarantee 20 KBps Weight 35
RealAudio Rule	Any	Any	RealAudio	Weight 5

Rule Name	Source	Destination	Service	Action
Marketing Rule	Marketing	Any	Any	Weight 30
Engineering Rule	Engineering	Any	Any	Weight 20
Default	Any	Any	Any	Weight 10

Connections matched to Web Rule will receive a total bandwidth of 20 KBps. Remaining bandwidth will be allocated to all the rules, Web Rule included, according to their weights.

For more on guarantees and limits, see Examples: Guarantees and Limits (on page 42) and Bandwidth Allocation and Rules (on page 26).

Sub-Rules

Sub-rules are rules nested in a rule. For example, you can create a sub-rule that allocates more bandwidth to HTTP connections that originate in Marketing. Connections whose Source is marketing receive more bandwidth than other HTTP traffic. In this example, the marketing sub-rule and default sub-rule is below the **Web Rule**:

Defining Sub-Rules

Rule Name	Source	Destination	Service	Action
Web Rule	Any	Any		Weight 20
Start of Sub-Rule				
Marketing HTTP	**Marketing**	Any	Any	Weight 10
Default	Any	Any	Any	Weight 1
End of Sub-Rule				

Bandwidth is allocated to **Web Rule** according to its weight (20). This weight is divided between its sub-rules in a 10:1 ratio. Connections below **Web Rule** are allocated bandwidth according to the weights specified:

- 10 for HTTP traffic from the Marketing department
- 1 for everything else.

Note:

- There are two **Default** rules: one for the Rule Base and one for the **Web Rule** sub-rule.
- The Source, Destination and Service fields of the sub-rule must always be a "sub-set" of the parent rule.

To create a sub-rule:

1. Right-click in the **Name** field of the rule in which you want to create the sub-rule.
2. Select **Add Sub-Rule**.

Installing a QoS Policy

To install a QoS Policy:

1. In SmartDashboard, make changes to Policy rules and then click **Update**.
2. In SmartConsole, click **Install Policy**.
3. From the **Policy list**, select the policy to install.
4. Click **Policy Targets** and select the Security Gateways that will get this Policy.

   Note - By default, no gateways are selected for QoS. You must select them manually.
5. Click **Install**.

If the installation is successful, the new Policy is enforced by the Security Gateways on which it is installed. If installation fails, do these steps to see the error messages:

1. Click the Task Information area, in the lower, left hand corner of SmartConsole.
2. In the **Recent Tasks** area, click **Details** on the applicable error.

In the **Install Policy Details** window, click the **^** icon in the **Status** column to see the error messages. You must resolve all errors before you can successfully install the Policy.

Advanced QoS Policy Management

In This Section:

Overview

This chapter covers more advanced QoS Policy management procedures that let you to refine the basic QoS Policies described in Basic Policy Management (on page 24).

Examples: Guarantees and Limits

The QoS Action properties defined in the rules and sub-rules of a QoS Policy Rule Base decide bandwidth allocation.

The guidelines and examples in the sections that follow show how to use effectively guarantees and limits.

Per Rule Guarantees

- The bandwidth allocated to the rule equals the guaranteed bandwidth plus the bandwidth allocated to the rule because of its weight. To uphold the guarantee, the guaranteed bandwidth is subtracted from the total bandwidth and set aside. The remaining bandwidth is then distributed according to the weights specified by all the rules.

The bandwidth guaranteed to a rule is the guaranteed bandwidth plus the rule's share of bandwidth according to weight.

Total Rule Guarantees

Rule Name	Source	Destination	Service	Action
Rule A	Any	Any	ftp	Rule Guarantee - 100KBps Weight 10
Rule B	Any	Any	http	Weight 20

- The link capacity is 190KBps.
- In this example, Rule A receives 130KBps, 100KBps from the guarantee, plus (10/30) * (190-100).
- Rule B receives 60KBps, which is (20/30) x (190-100).

- If a guarantee is defined in a sub-rule, then a guarantee must be defined for the rule above it. The guarantee of the sub-rule can also not be greater than the guarantee of the rule above it.

Guarantee is Defined in Sub-rule A1, But Not in Rule A Making the Rule Incorrect

Rule	Source	Destination	Service	Action
Rule A	Any	Any	ftp	Weight 10
Start of Sub-Rule				
Rule A1	Client-1	Any	ftp	Rule Guarantee - 100KBps Weight 10
Rule A2	Client-2	Any	ftp	Weight 10
End of Sub-Rule				
Rule B	Any	Any	http	Weight 30

This Rule Base is not correct. The guarantee is defined in sub-rule A1, but not in Rule A. To correct this, add a guarantee of 100KBps or more to Rule A.

- A rule guarantee must not be smaller than the sum of guarantees defined in its sub-rules.

Example of an Incorrect Rule Base

Rule	Source	Destination	Service	Action
Rule A	Any	Any	ftp	Rule Guarantee - 100KBps Weight 10
Start of Sub-Rule				
Rule A1	Client-1	Any	ftp	Rule Guarantee - 80KBps Weight 10
Rule A2	Client-2	Any	ftp	Rule Guarantee - 80KBps Weight 10
Rule A3	Client-3	Any	ftp	Weight 10
End of Sub-Rule				
Rule B	Any	Any	http	Weight 30

This Rule Base is incorrect. The sum of guarantees in Sub-Rules A1 and A2 is (80 + 80) = 160, which is greater that the guarantee defined in Rule A (100KBps). To correct this, define a guarantee not smaller than 160KBps in Rule A, or decrease the guarantees defined in A1 and A2.

- If a rule's weight is low, connections that match the rule might receive little bandwidth.

If a Rule's Weight is Low, Some Connections Might Receive Little Bandwidth

Rule	Source	Destination	Service	Action
Rule A	Any	Any	ftp	Rule Guarantee - 100KBps Weight 1
Start of Sub-Rule				
Rule A1	Client-1	Any	ftp	Rule Guarantee - 100KBps Weight 10
Rule A2	Client-2	Any	ftp	Weight 10
End of Sub-Rule				
Rule B	Any	Any	http	Weight 30

The link capacity is 190KBps.

Rule A is entitled to 103KBps, which are the 100KBps guaranteed, plus (190-100) x (1/31). FTP traffic classified to Sub-Rule A1 receives the guaranteed 100KBps which is almost all the bandwidth to which Rule A is entitled. All connections classified to Sub-Rule A2 together receive only 1.5KBps, which is half of the remaining 3KBps.

- The sum of guarantees in rules in the top level must not be more than 90% of the capacity of the link.
- The guarantee rule reserves the bandwidth only if a connection matches the guarantee rule. If no connection matches the guarantee rule, the bandwidth is not reserved.
- When the connection speed is less than the bandwidth guarantee, the guarantee rule makes unused bandwidth available to other connections.

For example, if the guarantee is 5MB and the connection speed is 3MB. The unused 2MB reserved by the rule is made available for other connections.

Per Connections Guarantees

1. If the **Accept additional connections** option is selected, connections exceeding the number defined in the **Number of guaranteed connections** are opened. If the field adjacent to **Accept additional connections** is empty, additional connections receive bandwidth allocated according to the defined **Rule Weight**.
2. You can define **Per connection guarantees** for a rule and for its sub-rule. The **Per connection guarantee** of the sub-rule must not be greater than the **Per connection guarantee** of the rule.

When such a Rule Base is defined, a connection classified to the sub-rule receives the **Per connection guarantee** that is defined in the sub-rule. If a sub-rule does not have a **Per connection guarantee**, it still receives the **Per connection guarantee** defined in the parent rule.

Limits

A rule can have both a **Rule limit** and a **Per connection limit**. But the **Per connection Limit** must not be greater than the **Rule Limit**.

If a limit is defined in a rule with sub-rules, and limits are defined for all the sub-rules, the rule limit has a restriction. *The rule limit must not be greater than the sum of limits defined in the sub-rules*. It is not possible to give more bandwidth to a rule than the bandwidth determined by the sum of the limits of its sub-rules.

Guarantee - Limit Interaction

- If a **Rule Limit** and a **Guarantee per rule** are defined in a rule, the limit must not be less than the guarantee.
- If both a Limit and a Guarantee are defined in a rule, and the Limit is equal to the Guarantee, connections might not receive bandwidth.

Example:

No Bandwidth Received:

Rule	Source	Destination	Service	Action
Rule A	Any	Any	ftp	Rule Guarantee — 100KBps Rule Limit 100KBps Weight 10
Start of Sub-Rule				
Rule A 1	Client-1	Any	ftp	Rule Guarantee - 100KBps Weight 10
Rule A2	Client-2	Any	ftp	Rule Guarantee - 80KBps Weight 10
End of Sub-Rule				
Rule B	Any	Any	http	Weight 30

The Guarantee in sub-rule A1 equals the Guarantee in rule A (100KBps). When there is sufficient traffic on A1 to use the full Guarantee, traffic on A2 does not receive bandwidth from A. (There is a limit on A of 100KBps).

In this example:

- A rule has both a guarantee and a limit, such that the limit equals the guarantee.
- The rule has sub-rules with Total Rule Guarantees that add up to the Total Rule Guarantee for the rule.
- The rule also has sub-rule(s) with no guarantee.

In such a case, the traffic from the sub-rule(s) with no guarantee might receive little or no bandwidth.

Differentiated Services (DiffServ)

Overview

DiffServ is an architecture for giving different types or levels of service for network traffic.

When on the enterprise network, packets are marked in the IP header TOS byte as belonging to some Class of Service (QoS Class). When outside on the public network, these packets are granted priority according to their class.

DiffServ markings have meaning on the public network, not on the enterprise network. Good implementation of DiffServ requires that packet markings be recognized on all public network segments.

DiffServ Markings for IPSec Packets

When DiffServ markings are used for IPSec packets, the DiffServ mark can be copied between headers by setting these properties in: `$FWDIR/conf/objects_5_0.c.`

- `:ipsec.copy_TOS_to_inner`

  The DiffServ mark is copied from the IPSec header to the IP header of the packet after decapsulation/decryption.

- `:ipsec.copy_TOS_to_outer`

  The DiffServ mark is copied from the packet's IP header to the IPSec header of the encrypted packet after encapsulation.

  The default setting are:

```
:ipsec.copy_TOS_to_inner (false)
:ipsec.copy_TOS_to_outer (true)
```

Interaction Between DiffServ Rules and Other Rules

Just like QoS Policy Rules, a DiffServ rule specifies not only a QoS Class, but also a weight. These weights are enforced only on the interfaces on which the rules of this class are installed.

For example, if a DiffServ rule specifies a weight of 50 for FTP connections. That rule is installed only on the interfaces for which the QoS Class is defined. On other interfaces, the rule is not installed. FTP connections routed through the other interfaces do not get the weight specified by the rule. To specify a weight for all FTP connections, add a rule below "Best Effort."

DiffServ rules can be installed only on interfaces for which the related QoS Class has been defined. QoS class is defined on the **QoS** tab of the **Interface Properties** window. For more, see: Define the QoS Properties for the Interfaces ("Configuring QoS Properties for Interfaces" on page 35).

"Best Effort" rules (that is, non-DiffServ rules) can be installed on all interfaces of gateways with QoS gateways installed. Only rules installed on the same interface interact with each other.

Note:

- QoS supports adding DiffServ markings to packets that match a rule
- QoS does not support matching packets based on DiffServ tagging.

Low Latency Queuing

Overview

For most traffic on the Web (most TCP protocols), the WFQ (Weighted Fair Queuing, see Intelligent Queuing Engine) paradigm is sufficient. Packets reaching QoS are put in queues and forwarded according to the interface bandwidth and the priority of the matching rule.

Using this standard Policy, QoS avoids dropping packets. Dropped packets adversely affect TCP. Avoiding drops means holding (possibly) long queues, which can lead to non-negligible delays.

For some types of traffic, such as voice and video, bounding this delay is important. Long queues are inadequate for these types of traffic. Long queues can result in substantial delay. For most "delay sensitive" applications, it is not necessary to drop packets from queues to keep the queues short. The fact that the streams of these applications have a known, bounded bit rate can be utilized. If QoS is configured to forward as much traffic as the stream delivers, only a small number of packets are queued and delay is negligible.

QoS Low Latency Queuing makes it possible to define special Classes of Service for "delay sensitive" applications like voice and video. Rules below these classes can be used together with other rules in the QoS Policy Rule Base. Low Latency classes require you to specify the maximal delay that is tolerated and a Constant Bit Rate. QoS then guarantees that traffic matching rules of this type is forwarded within the limits of the bounded delay.

Low Latency Classes

For each Low Latency class defined on an interface, a constant bit rate and maximal delay must be specified for active directions. QoS checks packets matched to Low Latency class rules to make sure they have not been delayed for longer than their maximal delay permits. If the maximal delay of a packet has been exceeded, it is dropped. Otherwise, it is transmitted at the defined constant bit rate for the Low Latency class to which it belongs.

If the Constant Bit Rate of the class is not smaller than the expected arrival rate of the matched traffic, packets are not dropped. The maximal delay must also exceed some minimum. For more, see Computing Maximal Delay (see "Calculating Maximal Delay" on page 48)).

When the arrival rate is higher than the specified Constant Bit Rate, packets exceeding this constant rate are dropped. This is to make sure that transmitted packets comply with the maximal delay limitations.

Note - The maximal delay set for a Low Latency class is an upper limit. Packets matching the class are always forwarded with a delay not greater, but often smaller, than specified.

Low Latency Class Priorities

In most cases, one Low Latency class is sufficient for all bounded delay traffic. In some cases, it might be necessary to define more than one Low Latency class. For this reason, Low Latency classes are assigned one out of five priority levels (not including the Expedited Forwarding class, see Low Latency versus DiffServ (on page 51)). These priority levels are relative to other Low Latency classes.

As a best practice, define more than one Low Latency class if different types of traffic require different maximal delays.

The class with the lower maximal delay must get a higher priority than the class with the higher delay. When two packets are ready to be forwarded, one for each Low Latency class, the packet from the higher priority class is forwarded first. The remaining packet (from the lower class) then encounters greater delay. The maximal delay that can be set for a Low Latency class depends on the Low Latency classes of higher priority.

Other Low Latency classes can affect the delay incurred by a class. Other Low Latency classes must be taken into consideration when determining the minimal delay that is possible for the class. This is best done by:

- Initially setting the priorities for all Low Latency classes according to maximal delay

- Defining the classes according to descending priority

When you define class two, for example, class one must already be defined.

For more on the effects of class priority on calculating maximal delay, see: Computing Maximal Delay (see "Calculating Maximal Delay" on page 48).

Logging LLQ Information

The system logs data for all aspects of LLQ.

Calculating the Correct Constant Bit Rate and Maximal Delay

Limits on Constant Bit Rate

For the inbound or outbound interface direction, the sum of the constant bit rates of all the Low Latency classes has a limit. This sum cannot exceed 20% of the total designated bandwidth rate. This 20% limit makes sure that "Best Effort" traffic does not suffer substantial jitter because of the existing Low Latency class(es).

Calculating Constant Bit Rate

To calculate the Constant Bit Rate of a Low Latency class, you must know the bit rate of one application stream in traffic that matches the:

- Class
- Number of expected streams that are simultaneously opened.

The Constant Bit Rate of the class equals the bit rate of one application multiplied by the expected number of streams opened at the same time.

If the number of streams is greater than the number you expected, the total incoming bit rate will exceed the Constant Bit Rate. Many drops will occur. To prevent drops, limit the number of concurrent streams. For more, see Ensuring that Constant Bit Rate is Not Exceeded (Preventing Unwanted Drops) (see "Making sure that Constant Bit Rate is not Exceeded" on page 50).

Note - Unlike bandwidth allocated by a Guarantee, the constant bit rate allocated to a Low Latency class on an interface in a given direction is not increased when more bandwidth is available.

Calculating Maximal Delay

To calculate the maximal delay of a Low Latency class, take into account the:

- Maximal delay that streams matching the class can tolerate in QoS
- Minimal delay that QoS can guarantee this stream

It is important not to define a maximal delay that is too small, which can result in unwanted drops. The delay value defined for a class determines the number of packets that can be queued in the Low Latency queue before drops occur. The smaller the delay, the shorter the queue. A maximal delay that is not sufficient can cause packets to be dropped before they are forwarded. Allow for some packets to be queued, as explained in the steps below.

Best Practice - Use the default Class Maximal Delay defined in the LLQ log. To obtain this default number:

- First configure the correct Constant Bit Rate for the Class
- Give an estimation for the Class Maximal Delay

You can also set the Class Maximal Delay by obtaining estimates for the upper and lower bounds. Set the delay to a value between the bounds.

1. Estimate the greatest delay that you can set for the class:

   a) Refer to the technical details of the streaming application and find the delay that it can tolerate.

      For voice applications, the user generally starts to experience irregularities when the overall delay exceeds 150 ms.

   b) Find or estimate the bound on the delay that your external network (commonly the WAN) imposes. Many Internet Service Providers publish Service Level Agreements (SLAs) that guarantee some bounds on delay.

   c) The maximal delay must be set at no more than:

      (i) The delay that the streaming application can tolerate minus

      (ii) The delay that the external network introduces

   This makes sure that the delay introduced by QoS plus the delay introduced by the external network is no more than the delay tolerated by the streaming application.

2. Estimate the smallest delay that you can set for the class:

   - Find the bit rate of the streaming application in the application properties, or use SmartView Monitor.

     Note: Even if you set the Constant Bit Rate of the class to accommodate multiple simultaneous streams, do the next calculations with the rate of a single stream:

   - Estimate the typical packet size in the stream.
     - Find it in the application properties, or monitor the traffic.
     - If you do not know the packet size, use the size of the MTU of the LAN behind QoS. For Ethernet, this number is 1500 Bytes.
   - Many LAN devices, switches and NICs, introduce some burstiness to flows of constant bit rate by changing the delay between packets. For constant bit rate traffic generated in the LAN and going out to the WAN, monitor the stream packets on the QoS Security Gateway. To get an estimate of burst size, monitor the internal interface that precedes the QoS Security Gateway.
   - If no burstiness is detected, the minimal delay of the class must be no smaller than:

     ```
     3 x packet size
     ---------------
        bit rate
     ```

     This enables three packets to be held in the queue before drops can occur.

     The bit rate must be the bit rate of one application, even if the Constant Bit Rate of the class is for multiple streams.

   - If burstiness is detected, set the minimal delay of the class to at least:

     ```
     (burst size + 1) x packet size
     ------------------------------
                bit rate
     ```

The maximal delay that you select for the class must be between the smallest delay (step 2) and the greatest delay (step 1). Setting the maximal delay near to one of these values is not recommended. If you expect the application to burst occasionally, or if you don't know whether the application generates bursts at all, set the maximal delay close to the value of the greatest delay.

This error message can show after you enter the maximal delay: "The inbound/outbound maximal delay of class... must be greater than... milliseconds." The message shows if Class of Service that you define is not of the first priority (see Low Latency Class Priorities (on page 47)). The delay value displayed in the error message depends on the Low Latency classes of higher priority, and on interface speed.

Set the maximal delay to a value no smaller than the one printed in the error message.

Making sure that Constant Bit Rate is not Exceeded

If the total bit rate going through the Low Latency class exceeds the Constant Bit Rate of the class, then drops occur. (See: Logging LLQ Information (on page 48).)

This occurs when the number of streams opened exceeds the number you expected when you set the Constant Bit Rate.

To limit the number of streams opened through a Low Latency Class:

1. Define one rule under the class, with a per connection guarantee as its **Action**.
2. In the **Per Connection Guarantee** field of the **QoS Action Properties** window, define the per connection bit rate that you expect
3. In the **Number of guaranteed connections** field, define the maximal number of connections that you allow in this class.

   Do not select the **Accept additional non-guaranteed connections** option.

The number of connections is limited to the number you used to calculate the Constant Bit Rate of the class.

Interaction between Low Latency and Other Rule Properties

To activate a Low Latency class, define at least one rule below it in the QoS Policy Rule Base. Traffic matching a Low Latency class rule receives the delay and Constant Bit Rate properties defined for the specified class. The traffic is handled according to the rule properties (weight, guarantee and limit).

You can use all types of properties in the rules below the Low Latency class:

- Weight
- Guarantee
- Limit
- Per Connection Guarantee
- Per Connection Limit.

Think of the Low Latency class with its rules as a separate network interface:

- Forwarding packets at a rate defined by the Constant Bit Rate with delay bounded by the class delay
- With the rules defining the relative priority of the packets before they reach the interface

If a rule has a relatively low priority, then packets matching it are entitled to a small part of the Constant Bit Rate. More packets will be dropped if the incoming rate is not sufficiently small.

Note:

- Using sub-rules under the low latency class is not recommended. Sub-rules make it difficult to calculate streams that suffer drops and the drop pattern.
- Guarantees and limits are not recommended for the same reason. Except for **Per Connection Guarantees**, as described in Ensuring that Constant Bit Rate is Not Exceeded (Preventing Unwanted Drops (see "Making sure that Constant Bit Rate is not Exceeded" on page 50)).

When to Use Low Latency Queuing

Use Low Latency Queuing when:

- Low delay is important, and the bit rate of the incoming stream is known. For example video and voice applications. In such cases, specify both the maximal delay and the Constant Bit Rate of the class.
- Controlling delay is important, but the bit rate is unknown. For example, Telnet requires quick responses, but the bit rate is not known. If the stream occasionally exceeds the Constant Bit Rate, you do not want to experience drops. A longer delay is recommended.
  - Set the Constant Bit Rate of the class to a high estimate of the stream rate.
  - Set a large maximal delay (such as 99999 ms).

  The large delay makes sure that packets are not dropped if a burst exceeds the Constant Bit Rate. The packets are queued and forwarded according to the Constant Bit Rate.

  Note - When the incoming stream is smaller than the Constant Bit Rate, the actual delay is much smaller than 99999 ms. (As in the example above). Packets are forwarded almost as soon as they arrive. The 99999 ms bound is effective only for large bursts.

Do not use a Low Latency Class when controlling delay is not of primarily importance. For most TCP protocols (such as HTTP, FTP and SMTP) the other type of QoS rule is more applicable. Use Weights, Limits and Guarantees. The correct priority is imposed on traffic without having to adjust bit rate and delay.

QoS enforces the policy with minimal drops. Weights and guarantees dynamically fill the pipe when expected traffic is not present. Low Latency Queuing limits traffic according to the Constant Bit Rate.

Low Latency versus DiffServ

Low Latency classes are different from DiffServ classes in that they do not receive type of service (TOS) markings. Not all packets are marked as Low Latency. Preferential treatment is guaranteed only while the packets are passing through the QoS Security Gateway.

The exception to this rule is the Expedited Forwarding DiffServ class. A DiffServ class defined as an Expedited Forwarding class automatically becomes a Low Latency class of highest priority. Such a class receives the conditions afforded it by its DiffServ marking both in QoS and on the network.

Note: To use the Expedited Forwarding class as DiffServ only, without delay being enforced, specify a **Maximal Delay** value of **99999** in the **Interface Properties** tab (see Low Latency Classes (on page 47)).

When to Use DiffServ and When to Use LLQ

Do not use Low Latency Queuing to delay traffic when your ISP:

- Supports DiffServ

  Despite the DiffServ marking that you apply, the IP packets might get a different QoS level from the ISP.

- Offers you a number of Classes of Service using MPLS

  DiffServ marking communicate to your ISP the Class of Service that you expect all packets to receive.

For these two cases, mark your traffic using a DiffServ class (see When to Use Low Latency Queuing (on page 51)):

Authenticated QoS

Check Point Authenticated QoS gives Quality of Service (QoS) for end-users in dynamic IP environments, such as remote access and DHCP environments. This lets priority users, such as corporate CEOs, to receive priority service when remotely connecting to corporate resources.

Authenticated QoS dynamically prioritizes end-users, based on information gathered during network or VPN authentication. The feature leverages Check Point UserAuthority technology to classify both inbound and outbound user connections. The User Authority Server (UAS) maintains a list of authenticated users. When you query the UAS, QoS retrieves the data and allocates bandwidth accordingly.

QoS supports Client Authentication, Encrypted Client Authentication, and SecuRemote/SecureClient Authentication. User and Session Authentication are not supported.

Note - Authenticated QoS is available for backward compatibility, but only works in QoS policy mode and does not support CoreXL or SecureXL acceleration technologies.

Citrix MetaFrame Support

This section covers support for Citrix Metaframe.

Overview

Citrix MetaFrame is a client/server software application that enables a client to run a published application on a Citrix server farm from the client's desktop. Citrix MetaFrame:

- Load balances by:
  - Automatically directing a client to the server with the lightest load
  - Allows publishing and application management from only one server in that farm.
- Supplies a secure encryption option using the ICA (Independent Computing Architecture) protocol developed by Citrix.

Note:

- Uncontrolled printing traffic on Citrix ICA networks with a slow internet connection can remove bandwidth from mission critical applications. On slow networks, it is necessary to differentiate between Citrix traffic and other types of:
  - Network traffic
  - Traffic in the same layer (layer 7)

- *The Citrix Print Manager service can only be used in QoS policy mode when SecureXL or CoreXL acceleration technologies are not enabled.*

QoS solves the problem of uncontrolled printing traffic on Citrix ICA networks by:

- Identifying all ICA applications running over Citrix through layer 7.
- Differentiating between the Citrix traffic based on ICA published applications, ICA printing traffic (Priority Tagging) and NFuse.

For more, see: Managing QoS for Citrix ICA Applications (on page 79).

To manage printing over Citrix, QoS uses the `Citrix_ICA_printing` service. The Citrix ICA printing traffic service is supported in *NG with Application Intelligence (R55)* and higher, but is not supported:

- In Express policy mode
- When SecureXL or CoreXL acceleration technologies are enabled.

In QoS policy mode, when SecureXL or CoreXL is enabled on the Security Gateway, you cannot that uses the `Citrix_ICA_printing` service in a rule. Policy installation will fail.

For more, see: Managing QoS for Citrix Printing (on page 82).

Limitations

- Citrix services are supported in QoS policy mode only, but not when SecureXL or CoreXL are enabled on the Security Gateway.
- Session Sharing must be disabled.
- The inspection infrastructure detects a maximum of 2048 applications. After the 2048 limit, console errors are sent. These errors do not affect your system. To stop the errors, restart the machine.
- Versions of MetaFrame prior to 1.8 are not supported because there is no packet tagging in these versions.
- Only one Citrix TCP service can be allocated per rule.

Load Sharing

Overview

Load Sharing is a mechanism that distributes traffic in a cluster of gateways to increase the total throughput. QoS architecture guarantees that Load Sharing uses either:

- Two-way Stickiness - all packets of a connection use the same Security Gateway in both directions.
- Conversation Stickiness - all packets of control/data connections in a conversation use the same Security Gateway in both directions.

In Load Sharing configurations, all functioning gateways in the cluster are active, and handle network traffic (Active/Active operation). If one of the member gateways fails, its connections are redistributed to other members of the cluster.

If one Security Gateway in the cluster becomes unreachable, transparent failover occurs to the other members. Connections are shared between the remaining gateways without interruption.

Note - The new Check Point High Availability is a special type of load sharing that automatically works with QoS Load Sharing. These modes can be safely switched. To enforce the change though, The QoS policy has to be reinstalled.

All cluster servers share the same set of virtual interfaces. Each virtual interface corresponds to an outgoing link. The next example shows a typical cluster.

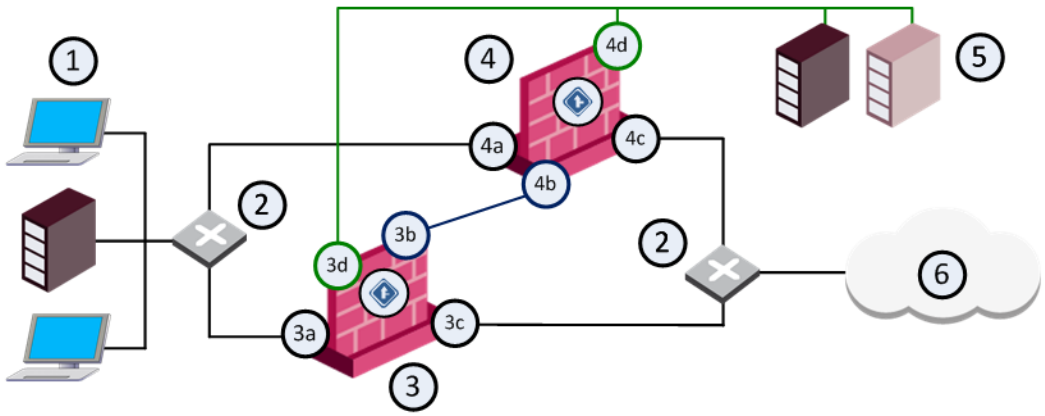

Item	Description
1	Internal network
2	Router
3	Security Gateway cluster member 1 with QoS
3a	Virtual interface to the internal network
3b	Interface to the Cluster Sync network
3c	Virtual interface to the external network (Internet)
3d	Virtual interface to the DMZ
4	Security Gateway cluster member 2 with QoS
4a	Virtual interface to the internal network
4b	Interface to the Cluster Sync network
4c	Virtual interface to the external network (Internet)
4d	Virtual interface to the DMZ
5	DMZ with Web and FTP servers
6	Internet

QoS gives a fault-tolerant QoS solution for cluster Load Sharing that deploys a unique, distributed WFQ bandwidth management technology. The user specifies a unified QoS policy for each virtual interface of the cluster. The resulting bandwidth allocation is the same as that obtained by installing the policy on one server.

Note - In a situation of heavy load, a few connections are backlogged active for short periods of time. The load is not spread evenly. In such cases the load is not spread evenly, but in this case there is no congestion and therefore no need for QoS.

QoS Cluster Infrastructure

This section describes the cluster infrastructure needed for QoS load sharing.

Cluster State

ClusterXL introduces a member's load value. A member's load, calculated in percentages, is assigned to each member by the cluster. The load is different for ClusterXL multicast and unicast modes.

Usually, the load for N members in the cluster equals (100 / N)%. If the number of cluster members changes dynamically (due to failover or recovery) the load is dynamically adjusted by the cluster to the applicable value.

Changes in Cluster State

All cluster members recalculate their rates when the load on one of the members changes. If a member fails, on the next rate calculation the members bandwidth is divided between the active cluster members.

Rates Calculation Algorithm

QoS Load Sharing uses a member's load value to get the correct rates allocation for QoS rules. QoS calculates the actual rate according to these criteria:

For a centralized policy rule:

- The rate, limit and guarantee values in the rule are proportionally divided between each cluster member according to their load value
- The rate is equally divided between each connection that matches the rule

For a physical network interface, the limit is:

- A fraction of the cluster-interface limit value
- Proportional to the cluster-member's load value

The rates, limits and guarantees are recalculated each time the cluster's state changes.

Note - If the QoS daemon cannot retrieve a load, it calculates the load statically according to the (100 / N)% formula. Where N is the number of members configured in the cluster topology that are not active.

Per-connection guarantees are processed separately (per-connection limit implementation remains unchanged by the Load Sharing mechanism).

Per-Connection Guarantee Allocation

Each rule with a per connection guarantee manages its rate budget. A rule's budget is the sum of all per connection guarantee rates over the number of per-connection guarantee connections allowed by this rule.

To determine if a new connection receives its per-connection guarantee, the overall rate (already granted to the matched rule's per-connection guarantee) is checked. If this rate is below the rule's budget then the new connection is granted its per-connection guarantee.

This budget is also divided between cluster members proportionally to their load. Generally, each member will get only half the allowed per-connection guarantee matched to the rule. The cluster grants per-connection guarantee service according to the QoS policy.

Example of Rates Calculation

This example shows a cluster consisting of two members with one virtual interface configured to the rate of 125KBps. The two members of the cluster equally share the load. The centralized scheduling policy and the corresponding local scheduling policies look like this:

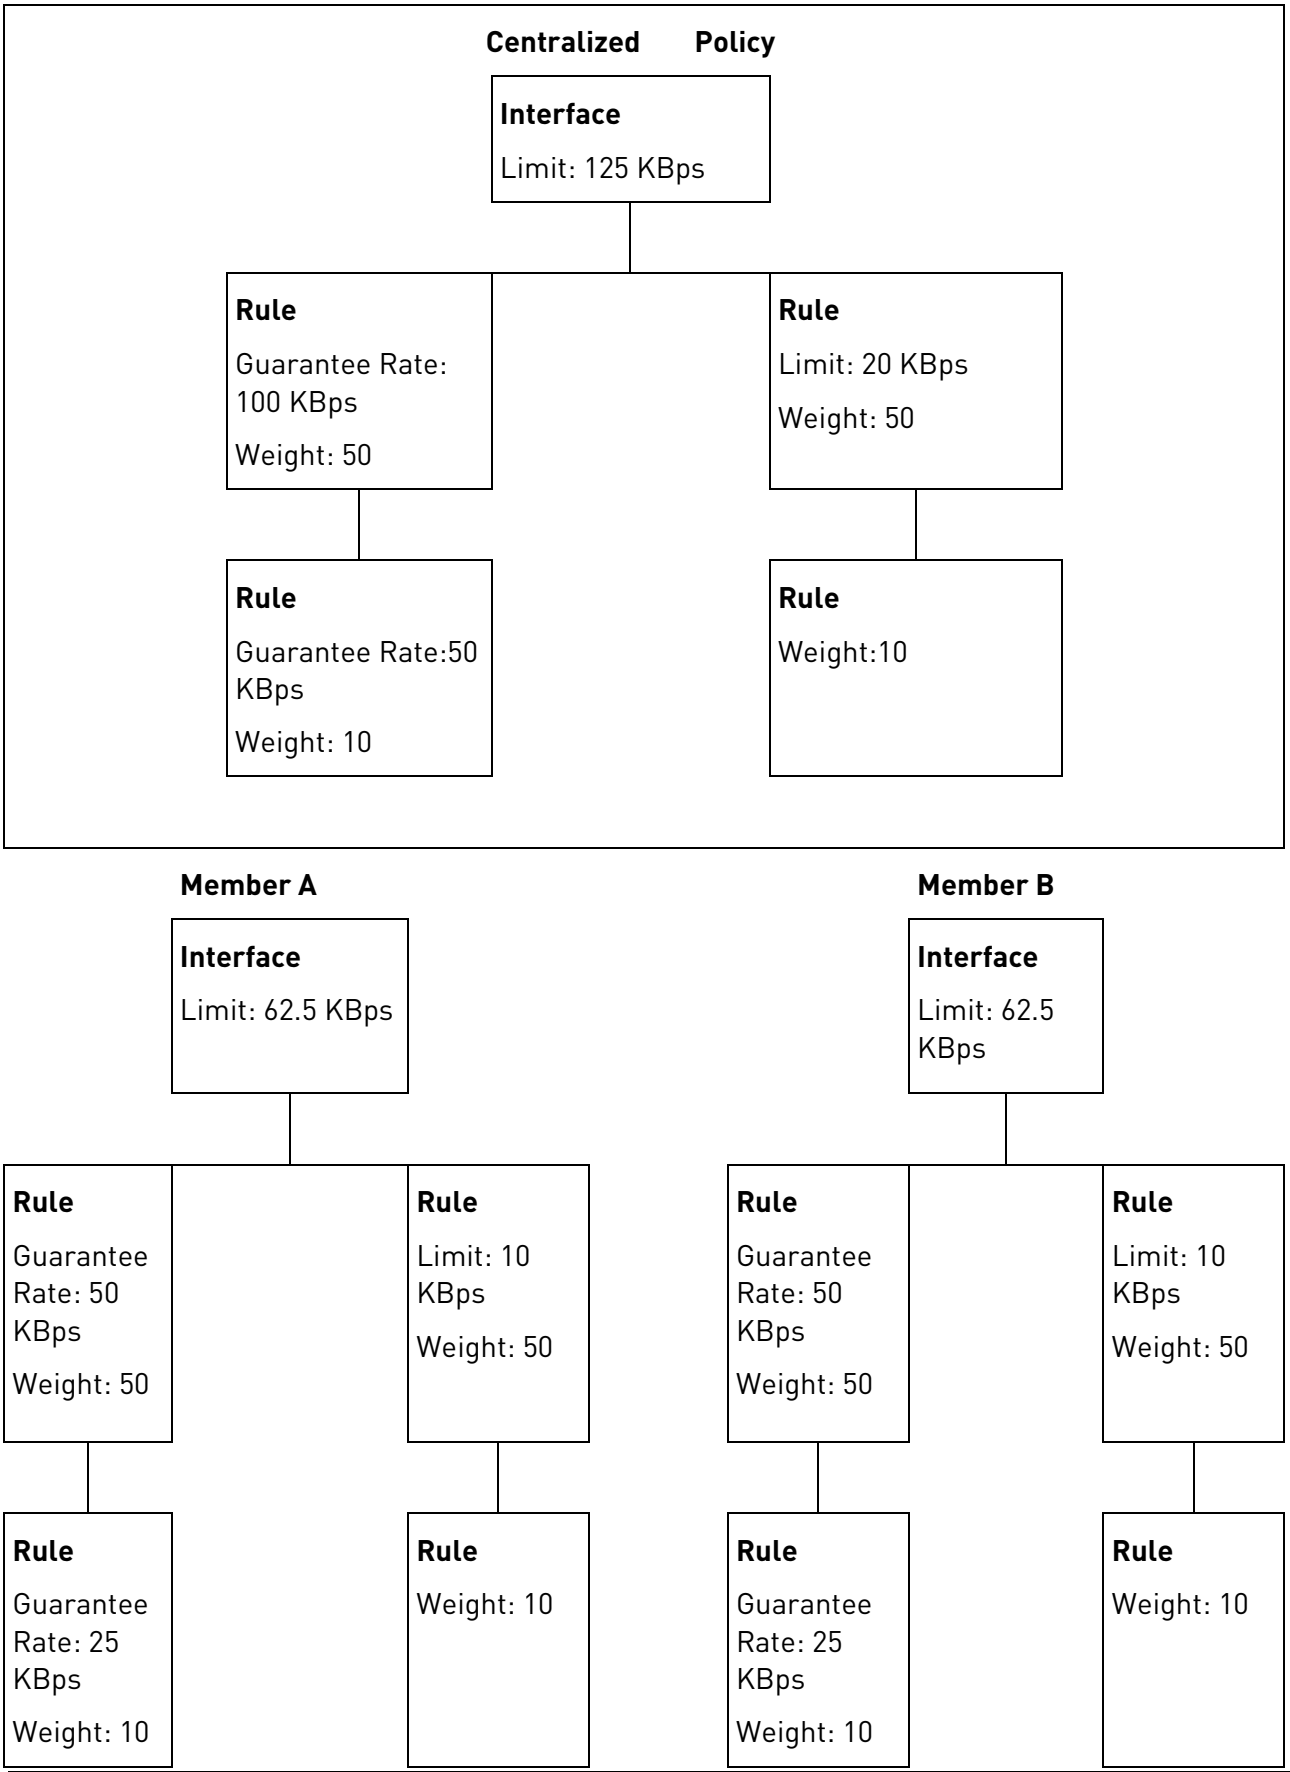

Conclusion

The decision function distributes traffic between all cluster members. The resultant load sharing allocates the same rates to the rules/connections as would be done by a centralized policy.

Managing QoS

In This Section:

This chapter shows you how to configure and manage QoS. These procedures assume that you have opened SmartConsole, as described in Opening the GUI Clients (on page 24).

Defining QoS Global Properties

The QoS global properties include default values for QoS rule parameters, unit of measure, and QoS authentication timeouts. Configure QoS global properties in SmartConsole.

Note: You must close SmartDashboard before you can work with global properties.

To configure QoS Global Properties:

1. In SmartConsole click **Application Menu** > **Global properties** > **QoS**.
2. In the Global Properties window, configure these parameters:

Weight:

- **Maximum weight of rule:** The maximum weight that can be assigned to rules. The default value is 1000, but can be changed to any number.
- **Default weight of rule:** The weight to be assigned in the **Action** column by default to new rules, including new **Default** rules.

Rate:

- **Unit of measure:** The unit specified in QoS windows by default for transmission rates (for example, Bps - Bytes per second).

Authenticated timeout for QoS :

- **Authenticated IP expires after:** If a user has been authenticated, all connections that are opened within the specified time receive the guaranteed bandwidth connection. Any connection opened after the specified time will be queried with the User Authority Server (UAS) again.
- **Non authenticated IP expires after:** If a user has previously tried and failed to be authenticated by the QoS Policy, then all connections that are opened within the specified

time will not receive the guaranteed bandwidth connection. This means that they will not match that specific rule during that time.

- **Unanswered queried IP expires after:** The User Authority Server (UAS) database is queried to see if a user's IP has been previously authenticated using Client Authentication or SSL. Until an answer is received, connections from this user will be classified to the next matching rule. If an answer is not received within the specified time, there will be another query.

Note - Click **Set Default** to restore the default settings for the **Authentication timeout for QoS** parameters.

3. Click **Set Default** to save the default values.

Changing QoS Global Properties

To configure QoS Global Properties:

1. From the **Policy** menu, choose **Global Properties** or click the **Edit Global Properties** icon in the toolbar.

The **Global Properties** window opens showing these fields:

In the **Weight** area:

- **Maximum weight of rule:** The maximum weight that can be assigned to rules. The default value is 1000, but can be changed to any number.
- **Default weight of rule:** The weight to be assigned in the **Action** column by default to new rules, including new **Default** rules.
- In the **Rate** area:
- **Unit of measure:** The unit specified in QoS windows by default for transmission rates (for example, Bps - Bytes per second).

In the **Authenticated timeout for QoS** area:

- **Authenticated IP expires after:** If a user has been authenticated, all connections that are opened within the specified time receive the guaranteed bandwidth connection. Any connection opened after the specified time will be queried with the User Authority Server (UAS) again.
- **Non authenticated IP expires after:** If a user has previously tried and failed to be authenticated by the QoS Policy, then all connections that are opened within the specified time will not receive the guaranteed bandwidth connection. This means that they will not match that specific rule during that time.
- **Unresponded queried IP expires after:** The User Authority Server (UAS) database is queried to see if a user's IP has been previously authenticated using Client Authentication or SSL. Until an answer is received, connections from this user will be classified to the next matching rule. If an answer is not received within the specified time, there will be another query.

Note - Click **Set Default** to restore the default settings for the **Authentication timeout for QoS** parameters.

2. Click **OK** to save the changes to the QoS Global Properties.

Interface QoS Properties

You must first define the network objects, that is, the Security Gateway and its interfaces on which QoS controls traffic flow.

After defining the interfaces you can specify the QoS properties for those interfaces. This is done in the **QoS** tab of the **Interface Properties** window. Defining the interface QoS properties involves setting the Inbound and Outbound active transmission rates and specifying the Differentiated Services (DiffServ) and Low Latency classes. You can change these definitions at any time.

Note - The **QoS** tab is only enabled for the interfaces of gateways that have **QoS** selected on the **General Properties** page of the Security Gateway.

Configuring Interface QoS Properties

To configure Security Gateway interfaces

1. Open SmartConsole.
2. Click **Gateways & Servers** and double-click the applicable **Security Gateway**.
3. In the **General Properties**, click **Network Management**.

   The **Check Point Gateway - Topology** window opens.
4. If a list of interfaces does not show, click **Get Interface.**

   If you choose this method of configuring the Security Gateway, the topology fetched suggests the external interface of the Security Gateway based on the QoS Security Gateway routing table. You must make sure that this information is correct.
5. Double-click the appropriate interface.
6. In the **Interface Properties** window, click the **QoS** tab.
7. In the **DiffServ and Low Latency classes** area, you can specify the Differentiated Services (DiffServ) and Low Latency Queuing classes to be used on the interface.

   You can **Add**, **Edit** or **Remove** a class. Refer to Working with Differentiated Services (DiffServ) (on page 74) and Working with Low Latency Classes **(see "**Working with Low Latency Queuing**" on page** 77**)** for more details on adding or editing DiffServ and Low Latency Classes.

   For information about DiffServ and Low Latency classes, see Differentiated Services (DiffServ) (on page 45) and Low Latency Queuing (on page 46).
8. Click **OK**

   Changes to the interface QoS properties are saved.

Do steps 4 - 7 for each applicable interface.

Notes:

- Interfaces on the WAN side (or interfaces connected to a slower network) are typically defined as active. On a gateway with only two interfaces, enable QoS only on the interface connected to the WAN. If the gateway controls DMZ traffic, you can install QoS on the interface connected to the DMZ.
  - Select **Inbound Active** to control traffic on this interface in the inbound direction.
  - From the **Rate** list, select or enter the available bandwidth in the inbound direction.
  - Check **Outbound Active** to control traffic on this interface in the outbound direction.
  - From the **Rate** list select the or enter the available bandwidth in the outbound direction.

- *Make sure that the rates correspond to the actual physical capacity of the interfaces.*

  QoS cannot not make sure the defined rates are compatible with the interface hardware.

  If the defined rate is less than the physical capacity, QoS uses only specified capacity. Excess capacity is not used. If the defined rate greater than the physical capacity, QoS cannot control traffic correctly.

Working with QoS Policies

QoS policy is an ordered set of QoS rules in the Rule Base. The Rule Base contains rules that you create, and a default rule. The default rule is automatically created with the Rule Base. It can be modified but cannot be deleted. The fundamental concept is that unless other rules apply, the default rule is applied to all data packets. The default rule is therefore always the last rule in the Rule Base.

The Rule Base specifies what actions are to be taken with the data packets. It specifies the source and destination of the communication, what services can be used, at what times, whether to log the connection and the logging level.

A QoS Rule Base is applied to specific gateways and interfaces. After you have created the Policy and defined its QoS rules you must install it on the relevant QoS gateways.

Creating a New QoS Policy

To create a New Policy:

1. On the gateway, make sure that the QoS blade is enabled.
2. In SmartConsole, from the **File** menu, select **Manage Policies and Layers**.
3. Click **New** .
4. In the **Policy** window, enter a Policy name.

   This name cannot:
   - Contain any reserved words or spaces.
   - Start with a number.
   - Contain any of the following characters: %, #, ', &, *, !, @, ?, <, >, /, \, :.
   - End with any of the following suffixes: .w, .pf, .W.
5. Select **QoS** and then select a QoS Policy type:
   - **Express** - Quickly create basic QoS Policies
   - **Recommended** (default) - Create advanced Policies with the full set of QoS features

   Note: There are some limitations that can prevent you from enabling SecureXL or CoreXL with QoS Policies.

   For more, see: QoS Policy limitations ("Acceleration Support for R77 Policies" on page 15).
6. Click **OK**.

   The system saves the new Policy and SmartDashboard opens automatically. You can start to define your rules here.

Opening an Existing QoS Policy

To Open an Existing Policy:

1. In SmartConsole, click **Security Policies** > Manage **Policies**.
2. In the Manage Policies window, double-click a QoS Policy.

SmartDashboard opens.

Creating New Rules

You work with rules in SmartDashboard. When you add rules, you can put the new rule anywhere in the Rule Base except after the last rule. The Default Rule must always be at the bottom of the Rule Base.

To create a new rule:

1. In the **QoS** tab, at the position where you want to add a new rule.
2. Add a new rule from the **Rule** menu, the toolbar, or right-click a name in the **Name** column of a rule to display the **Rule** menu.

The **Rule Name** window opens.

3. Enter the name of the rule in the **Rule Name** field.
4. Click **OK**.

The rule is added to the Rule Base at the selected position, with the values defined in the **QoS** page of the **Global Properties** window.

To add a rule	Select from Menu	Toolbar button
After the last rule	**Rules > Add Rule > Bottom**	
Before the first rule	**Rules > Add Rule > Top**	
After the current rule	**Rules > Add Rule > Below**	
Before the current rule	**Rules > Add Rule > Above**	
To the current rule	**Rules > Add Sub-Rule**	

Right-click a rule to use these menu commands:

Menu Option	Explanation
Add Rule above	Adds a rule before the current rule.
Add Rule below	Adds a rule after the current rule.
Add Sub-Rule	Deletes the current rule.
Delete Rule	Deletes the current rule.

Menu Option	Explanation
Copy Rule	Copies the current rule to the clipboard.
Cut Rule	Deletes the current rule and puts it in the clipboard.
Paste Rule	Pastes the rule in the clipboard (a sub-menu is displayed from which you can select whether to paste the rule above or below the current rule).
Add Class of Service	Specifies a Class of Service (see Differentiated Services (DiffServ) (on page 45) and Low Latency Queuing (on page 46)). A sub-menu is displayed from which you can select whether the Class of Service is to be added above or after the current rule.
Hide Rule	Hides the current rule. The rule is still part of the Rule Base and will be installed when the QoS Policy is installed.
Disable Rule	Disables the current rule. The rule appears in the Rule Base but is not enforced by the QoS Policy.
Rename Rule	Renames the current rule.

Changing the Rule Name

To change the rule name:

1. In the **QoS** tab, double-click the **Name** column in the rule to rename.
2. In the **Rule Name** window, enter the new rule name in the **Rule Name** field.
3. Click **OK**.

To Copy, Cut or Paste a Rule

You can copy, cut or paste a rule using either the **Edit** or **Rules** menus or the right-click menu of the selected rule.

1. In the **QoS** tab, select the rule you want to copy, cut or paste.
2. From the **Edit** or **Rules** menu, select one of the options described in the table below.

Action	From Menu select
Cut	Edit > Cut
Copy	Edit > Copy
Paste	Edit > Paste

If you select **Paste**, then the **Paste** menu will be opened. You must then select **Bottom**, **Top**, **Above**, or **Below** to specify where in the Rule Base to paste the rule.

To Delete a Rule

You can delete a rule using either the right-click menu of the selected rule or clicking the Delete button on the toolbar.

1. In the **QoS** tab, select the rule you want to delete.
2. Click **Delete** on the toolbar.
3. Click **Yes** to delete the selected rule.

Working with Rules

You can change rule fields, as often as you like, until the rule is in the form that you require. Configure the source and destination of each communication, services that can be used (TCP, Compound TCP, UDP, and ICMP), actions to be taken with the data packets, whether to maintain a log of the entries for the selected rule, and interfaces of the QoS Security Gateway that the rule is enforced.

This section describes the procedures for modifying the various fields in a rule. Refer to Overview (on page 24) for more details about rules.

Modifying Sources in a Rule

You can modify the source(s) of the communication in a rule. You can add as many sources as required. In addition, you can restrict the sources of the rule to particular user groups, or to user groups originating from specific locations.

To Add Sources to a Rule

1. From the **Rule Base** select the rule to modify.
2. Right-click the **Source** column of the selected rule and select **Add**.

The **Add Object** window shows listing the network objects defined in the Security Policy and the QoS Policy.

Note - You can also use the **Add Object** window to define new objects and delete or modify objects.

3. Select one or more network objects (using the standard Windows selection keys) to add to the rule's **Source**.
4. Click **OK**.
   - The objects are added to the **Source** field.
   - You can add as many sources as required.

To Add User Access to the Sources of a Rule

1. From the **Rule Base** select the rule you want to modify.
2. Right-click in the **Source** column of the selected rule and select **Add Users Access**. The **User Access** window is displayed.
3. Select one of the user groups to add to the rule's **Source**.

4. Select whether you want to restrict the **Location**, as follows:
   - **No restriction**: There is no restriction on the source of the users. For example, if you select **All Users** and check **No restriction**, then **AllUsers@Any** will be inserted under **Source** in the rule.
   - **Restrict to**: The source is restricted to the network object you select in the list box. For example, the source object in the rule will be **AllUsers@Local_Net**.
5. Click **OK** to add the user access to the rule source.

To Edit, Delete, Cut, Copy or Paste a Source in a Rule

You can edit, delete, cut, copy or paste a source in a rule using the right-click menu of the selected source.

1. From the **Rule Base** select the rule to modify.
2. Right-click on the **Source** of the selected rule
3. Select one of these options:
   - **Edit:** The appropriate window is opened, according to the type of object selected, and you can change the object's properties. (Alternatively, you can double-click on an object in the **Source** column of the selected rule to edit it.)
   - **Delete:** The selected object is deleted. If you delete the last source object in the rule it is replaced by **Any**.
   - **Cut:** The selected object is cut and put it in the clipboard.
   - **Copy:** The selected object is copied to the clipboard.
   - **Paste:** The object is pasted from the clipboard to the rule's **Source**.

To View Where an Object is Used

You can view where the selected object is used (in queries, active policies, and so on).

1. From the **Rule Base** select the rule to modify.
2. Right-click on the **Source** of the selected rule
3. Select **Where Used**.

   The **Object References** window opens showing where the selected object is used (in queries, active policies, and so on).
4. Click **Close** to return to the rule.

Modifying Destinations in a Rule

You can modify the destination(s) of the communication in a rule. You can add as many destinations as required.

To Add Destinations to a Rule

1. From the **Rule Base** select the rule to modify.
2. Right-click in the **Destination** column of the selected rule.
3. Select **Add**.

The **Add Object** window opens), listing the network objects defined in the Security Policy and the QoS Policy.

Note - You can also use the **Add Object** window to define new objects and delete or modify objects.

4. Select one or more network objects (using the standard Windows selection keys) to add to the rule's **Destination**.
5. Click **OK**.

The objects are added to the **Destination** field. Add as many destinations as required.

To Edit, Delete, Cut, Copy or Paste a Destination in a Rule

You can edit, delete, cut, copy or paste a destination in a rule using the right-click menu of the selected source.

1. From the **Rule Base** select the rule you want to modify.
2. Right-click on the **Destination** of the selected rule and select one of the following options:
   - **Edit:** The appropriate window is opened, according to the type of object selected, and you can change the object's properties. (Alternatively, you can double-click on an object in the **Destination** column of the selected rule to edit it.)
   - **Delete:** The selected object is deleted. If you delete the last destination object in the rule it is replaced by **Any**.
   - **Cut:** The selected object is cut and put it in the clipboard.
   - **Copy:** The selected object is copied to the clipboard.
   - **Paste:** The object is pasted from the clipboard to the rule's **Destination**.

To View Where an Object is Used

You can view where the selected object is used (in queries, active policies, and so on).

1. From the **Rule Base** choose the rule you want to modify.
2. Right-click on the **Source** of the selected rule and choose **Where Used**. The **Object References** window is displayed showing you where the selected object is used (in queries, active policies, and so on).
3. Click **Close** to return to the rule.

Modifying Services in a Rule

You can modify the service(s) in a rule. You can add as many services as required, however, you can only add one URI for QoS resource in a single rule.

Note - Previous versions of QoS have not limited the number of URIs for QoS resources allowed per rule. If you are using a QoS Policy originally designed for use with a previous QoS version, be sure to redefine any rule that has more than one resource in its **Service** Field.

To Add Services to a Rule

1. From the **Rule Base** select the rule to modify.
2. Right-click in the **Service** column of the selected rule.
3. Select **Add**.

The **Add Object** window shows listing the network objects defined in the Security Policy and the QoS Policy.

4. Select one or more network objects (using the standard Windows selection keys) to add to the rule's **Service**.
5. Click **OK**.

The objects are added to the **Service** field.

- You can add as many services as required.
- Only one **TCP Citrix** or **URI for QoS** service is allowed.

To Add a Service with a Resource to a Rule

1. From the **Rule Base** choose the rule you want to modify.
2. Right-click in the **Service** column of the selected rule and select **Add with Resources**.

The **Services with Resource** window opens.

You can only add one service with a resource to a rule, so this option will only be available if you have not already added a service with a resource to this rule.

3. Select one of the services in the **Location** area.
4. Select the appropriate resource from the **Resource** list.

Note:

- Only resources of type URI for QoS can be added to the QoS Rule Base. URI for QoS is used for identifying HTTP traffic according to the URL (URI).
- Do not use the protocol prefix (http://) when setting up a URI resource. HTTP services with URI for QoS resources can be defined on all ports.
- The regular expression ("Appendix: Regular Expressions" on page 103) supported by QoS is of form **a*b** where a and b are strings and * is wildcard.
- Both full and relative URI are supported:
  - **Full URI**: Use the full URI but without protocol prefix (for example, do **not** use"http://"). Valid full URI example: "www.my-site.com/pic/qos.gif"
  - **Relative URI**: Use the URI that starts just after the domain name. The relative URI must start with slash. For example: "/pic/qos.gif"

5. Click **OK** to add the service with a URI for QoS resource to the rule.

Note - Only one resource is allowed in a single rule.

To Edit, Delete, Cut, Copy or Paste a Service in a Rule

You can edit, delete, cut, copy or paste a service in a rule using the right-click menu of the selected service.

1. From the **Rule Base** the select the rule to modify.
2. Right-click on the **Service** of the selected rule.
3. Select one of these options:
   - **Edit:** The appropriate window is opened, according to the type of object selected, and you can change the object's properties. (Alternatively, you can double-click on an object in the **Service** column of the selected rule to edit it.)
   - **Delete:** The selected object is deleted. If you delete the last service object in the rule it is replaced by **Any**.
   - **Cut:** The selected object is cut and put it in the clipboard.

- **Copy:** The selected object is copied to the clipboard.
- **Paste:** The object is pasted from the clipboard to the rule's **Service**.

To View Where an Object is Used

You can view where the selected object is used (in queries, active policies, and so on).

1. From the **Rule Base** select the rule to modify.
2. Right-click on the **Service** of the selected rule.
3. Select **Where Used**.

   The **Object References** window opens showing you where the selected object is used (in queries, active policies, and so on).
4. Click **Close** to return to the rule.

Modifying Rule Actions

You can modify the default properties of a rule. The available options depend on whether it is a simple or advanced type of rule. The advanced rule action type enables you to specify limits and guarantee allocation on a per connection basis.

To Edit the Rule Actions

1. From the **Rule Base** choose the rule you want to modify.
2. Right-click in the **Action** column of the selected rule and select **Edit Properties**.

   The **QoS Action Properties** window opens.
   - If the **Action Type** of the rule is defined as **Simple**, the **QoS Action Properties** window opens:
   - If the **Action Type** of the rule is defined as **Advanced**, the **QoS Action Properties** window opens:

     Note - When Express QoS has been installed, **Advanced** Actions are not available.
3. The following properties are displayed for a QoS rule with a simple action type. You can change any of these fields:

   In the **Action Type** area:
   - **Simple:** The full set of actions with the exception of the **Guarantee Allocation** and the **per connection limit** features.
   - **Advanced:** The full set of actions with the **Guarantee Allocation** feature included.
   - In the **VPN Traffic** area:
   - **Allow rule only to encrypted traffic:** Check this box if you want the rule to be matched only by VPN traffic. If you do not check this field, rules will be matched by all traffic types, both VPN and non-VPN traffic. VPN traffic means traffic that is encrypted in this same Security Gateway by IPsec VPN. This field does not apply to traffic that was encrypted prior to arriving to this Security Gateway. This type of traffic can be matched using the "IPSec" service. For further explanation on how to use this check box for prioritizing VPN traffic over non-VPN, see Example of a Rule Matching VPN Traffic (on page 28).

- In the **Action Properties** area you can define the restrictions on bandwidth for connections to which the rule applies in the following fields:
- **Rule Weight:** Enables you to define the weight of the rule. This field is checked by default and has the value defined in the **Global Properties** window in Defining QoS Global Properties (on page 59). **Best Practice -** Leave this value as is to avoid a complete loss of bandwidth. For detailed information see Weight (on page 26).

  Important - 0 rate in conjunction with 0 guarantee can lead to the rule's complete loss of bandwidth. To prevent this from happening, retain some ratio in the **Rule Weight**. The default is 10.

- **Rule Limit:** Enables you to restrict the total bandwidth consumed by the rule. For detailed information see Limits (on page 27).

  Note - When using weights or guarantees, the weighted fair queuing algorithm that QoS makes use of assures that no bandwidth is ever wasted. Spare bandwidth is divided among the backlogged rules. However, if you set a rule limit, it will not use spare bandwidth above this limit.

- **Rule Guarantee:** Enables you to define the absolute bandwidth allocated to the rule. For detailed information see Guarantees (on page 26).

  Note - The number you enter for the **Rule Guarantee** cannot be larger than the **Rule Limit**.

4. (Optional) The following additional properties are displayed for a QoS rule with an advanced action type. You can change any of these fields:

In the **Limit** area:

- **Rule Limit:** Enables you to restrict the total bandwidth consumed by the rule. For detailed information see Limits (on page 27).

  Note - When using weights or guarantees, the weighted fair queuing algorithm that QoS makes use of assures that no bandwidth is ever wasted. Spare bandwidth is divided among the backlogged rules. However, if you set a rule limit, it will not use spare bandwidth above this limit.

- **Per connection limit:** Enables you to set a rule limit per connection.

  Note - The number you enter for the Rule Guarantee cannot be larger than the Rule Limit.

In the **Guarantee Allocation** area:

- **Guarantee:** Enables you to allocate a minimum bandwidth to the connections matched with a rule. For detailed information see Guarantees (on page 26).
- **Per rule:** Enables you to define the absolute bandwidth allocated to the rule.

  Note - The number you enter for the **Per rule** cannot be larger than the **Rule Limit**.

- **Per connection:** Enables you to manage the bandwidth at the connection-level.
- **Per connection guarantee:** Enables you to restrict the absolute bandwidth allocated per connection.
- **Number of guaranteed connections:** Enables you to allocate a minimum number of guaranteed connections.

Note - The **Number of guaranteed connections** multiplied by the **Per connection guarantee** cannot be greater than the rule limit.

- **Accept additional connections:** Check this option to allow connections without per connection guarantees to pass through this rule and receive any leftover bandwidth. Enter the maximum amount of bandwidth that is allowed for this option in the text box. This only occurs if all other conditions have been met.

Note - Select a non-zero rule weight when **Accept additional non-guaranteed connections** is checked.

5. Click **OK** to update the **QoS Action Properties** for the rule.

To Reset the Rule Actions to Default Values

1. From the **Rule Base** select the rule you want to modify.
2. Right-click in the **Action** column of the selected rule and select **Reset to Default**. The action properties for the selected rule are reset to their default values. The default values are defined in the **QoS** page of the **Global Properties** window (see Defining QoS Global Properties (on page 59)).

Modifying Tracking for a Rule

You can choose whether you want to maintain a log of the entries for the selected rule. If you do want to log the entries, you also have the option of logging the entries in account format. For further information on tracking and logging, see Overview of Logging (on page 84). For information on how to turn logging on, see Enabling Log Collection (on page 83).

To Modify Tracking for a Rule

1. From the **Rule Base** select the rule you want to modify.
2. Right-click in the **Track** column of the selected rule. The menu that is displayed has the following options:
   - **None**. No logging is done for this connection
   - **Log**. Logging is done for this connection
   - **Account**. Logging for this connection is done in Accounting format.
3. Select the required option.

Modifying Install On for a Rule

The **Install On** field specifies on which interfaces of the QoS Security Gateway the rule is enforced. You can select any number of **Install On** objects.

Note - To install a QoS Policy on a Security Gateway, make sure that:

- The Security Gateway has the QoS option selected on the **Network Security** tab of the gateways **General Properties** page.
- The interface is defined in the **QoS** tab of the **Interface Properties** window. (See Defining QoS Global Properties (on page 59) and Specifying Interface QoS Properties (see "Interface QoS Properties" on page 61).)

To Modify Install On for a Rule

1. From the **Rule Base** select the rule you want to modify.
2. Right-click in the **Install On** column of the selected rule and select **Add**. The **Add Interface** window is displayed.
3. (Optional) Click **Select Targets** to select additional installable targets. The **Select Installation Targets** window is displayed.
4. To add any target(s) to the list of Installed Targets, select the target(s) in the **Not in Installation Targets** area and click **Add**.

   The selected target(s) are added to the **In Installation Targets** area.
5. To remove a target(s) from the **In Installation Targets** area, select the target(s) and click **Remove**.

   The selected targets are returned to the **Not in Installation Targets** area.
6. Click **OK**. The selected targets now appear in the **Add Interface** window.
7. Select from the list of targets in the **Add Interface** window:
   - A Security Gateway (and all its interfaces on which QoS is defined), *or*
   - An interface (in both directions), *or*
   - One direction of an interface
8. Click **OK**. The selected interface is added to the **Install On** field.

To Delete an Install On for a Rule

You can remove an interface for a rule. The rule will no longer be enforced for the interface.

1. From the **Rule Base** select the rule to modify.
2. Right-click on the **Service** of the selected rule.
3. Select **Delete**.

   The selected object is deleted.

To View Where an Object is Used

You can view where the selected object is used.

1. From the **Rule Base** select the rule to modify.
2. Right-click on the **Install On** of the selected rule.
3. Select **Where Used**.

   The **Object References** window opens showing where the selected object is used.
4. Click **Close** to return to the rule.

Modifying Time in a Rule

You can specify the times that the rule is enforced. You add any number of time objects to a rule.

To Modify Time in Rules

1. From the **Rule Base** select the rule to modify.
2. Right-click in the **Time** column of the selected rule.
3. Select **Add**.

The **Add Object** window opens.

4. (Optional) You can edit a time object:
   a) Select the required time object and click **Edit** to modify a time object.

      The **Time Properties** window opens. (Alternatively, you can double-click on an object in the **Time** column of the selected rule to edit it.)

   b) Edit the fields in the **Time Properties** window, as required.

   c) Click **OK**.
5. Select the required time object in the **Add Object** window.

The time object is added to the rule.

To Edit or Delete a Time Object for a Rule

You can edit or delete a time object in a rule using the right-click menu of the selected service.

1. From the **Rule Base** choose the rule to modify.
2. Right-click on the **Time** column of the selected rule.
3. Select one of these options:
   - **Edit:** The appropriate window is opened, according to the type of object selected, and you can change the object's properties. (Alternatively, you can double-click on an object in the **Time** column of the selected rule to edit it.)
   - **Delete:** The selected object is deleted. If you delete the last time object in the rule it is replaced by **Any**.

To View Where an Object is Used

You can view where the selected object is used (in queries, active policies, and so on).

1. From the **Rule Base** select the rule to modify.
2. Right-click on the **Service** of the selected rule.
3. Select **Where Used**.

   The **Object References** window opens showing you where the selected object is used (in queries, active policies, and so on).
4. Click **Close** to return to the rule.

Adding Comments to a Rule

You can add a comment to a rule.

To Add Comments to Rules

1. From the **Rule Base** select the rule to modify.
2. Right-click in the **Comment** column of the selected rule.
3. Select **Edit**.

   The **Comment** window opens. You can also open this window by double-clicking in the **Comment** column of the selected rule.
4. Type relevant comments in the text box.
5. Click **OK**.

The comment is added to the rule.

Defining Sub-Rules

Sub-rules are rules that allocate bandwidth more specifically within a rule. For example, consider the rule shown in the figure below.

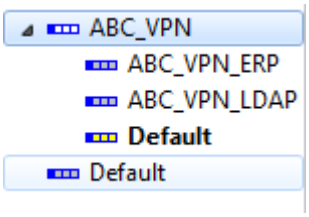

The bandwidth allocated to the **ABC_VPN** rule is further allocated among the sub-rules **ABC_VPN_ERP** through **Default** under **ABC_VPN**.

To Define Sub-Rules

1. Select the rule under which the sub-rule is to be defined.
2. Right-click in the **Rule Name** column.
3. Select **Add Sub-Rule** from the menu. The **Rule Name** window is displayed.
4. Enter the sub-rule name and click **OK**. The new sub-rule together with a default sub-rule is automatically created, under the rule selected in 1 above, using the default values defined.
5. You may modify the sub-rules by following the same procedures for editing rules described on page in Editing QoS Rule Bases (see "Working with QoS Policies" on page 62).
6. Add new sub-rules by following the same procedures for creating rules described in Editing QoS Rule Bases (see "Working with QoS Policies" on page 62).

To View Sub-Rules

The sub-rules under a main rule can be seen by expanding the rule in the QoS Rule Tree. To view sub-rules in the Rule Base, click one of the sub-rules in the relevant main rule. The Rule Base shows all the sub-rules for that rule.

Working with Differentiated Services (DiffServ)

A DiffServ rule specifies not only a QoS Class, but also a weight, in the same way that other QoS Policy Rules do. These weights are enforced only on the interfaces on which the rule is installed.

For more on DiffServ, see: Differentiated Services (DiffServ) (on page 45).

Defining a DiffServ Class of Service

To define a DiffServ class of service:

1. From the SmartDashboard menu, select **Manage** > **QoS** > **QoS Classes**.
2. In the **QoS Classes** window, click **New** > **DiffServ Class of Service**.
3. In the **Class of Service Properties** window, configure these settings:
   - **Name** - The name of the Class of Service.
   - **Comment** - The text to be displayed when this class is selected in the **QoS Classes** window
   - **Color** - Select a color from the list.

- **Type** - Select a type from the list. You may select a predefined or user defined class.
- **DiffServ code** - This is a read-only field that displays the DiffServ marking as a bitmap.

4. Click **OK**.

Defining a DiffServ Class of Service Group

To define a DiffServ class of service group:

1. In SmartDashboard, click **Manage** > **QoS** > **QoS Classes**.
2. In the **QoS Classes** window, click **New** > **DiffServ Class of Service Group**.
3. In the **Group Properties** configure these properties:
   - **Name** - The name of the group.
   - **Comment** -The text to be displayed when this class is selected in the **QoS Classes** window.
   - **Color** - Select a color from the list.
   - To add a DiffServ class to the group, double-click a class in the list in the **Not in Group** list.
   - To delete a class from the group, double-click a class **In Group** list.
4. Click **OK**.

Configuring an Interface for DiffServ

Use these procedures to configure interfaces and to add a DiffServ class to an interface.

To configure interface for DiffServ:

1. In SmartConsole, go to **Gateways & Servers**.
2. Double-click the applicable Security Gateway.
3. In the **Check Point Gateway** window, click **Network Management**.
4. Double-click the applicable interface.
5. In the **Interface** window, click the **QoS** tab.
6. In the **Diffserv and Low Latency classes** section, click **Add** > **DiffServ Classes** > **Others**.
7. Select **Inbound Active** and/or **Outbound Active** and set the **Rate** properties.
8. In the **Object Editor** window, select a **QoS Class** from the list.
9. Select and configure these parameters for **Inbound** and/or **Outbound** traffic:
   - **Guaranteed bandwidth -** The bandwidth guaranteed marked for priority.

     IMPORTANT: Make sure you do not exceed the guaranteed bandwidth.
   - **Bandwidth Limit** - The maximum bandwidth for this class.
     Traffic volume greater than the Bandwidth Limit is marked for QoS priority.

   Note: You must configure these properties for at least one traffic direction.
10. Click **OK**.

To add QoS Classes to the Rule Base:

1. Open SmartDashboard.
2. Do one of these actions:
   - In the **Name** column of a QoS rule, click the rule **Add Class of Service > Above**.
   - In a class header, right click the header and then click **Add Class of Service Above** or **Add Class of Service Below**.

3. Select a class from the list, and then click **OK**.

   The DiffServ class header shows in the Rule Base. If this is the first defined class, the *Best_Effort* header shows directly below the new DiffServ class header.
4. Follow the steps in the next sections to define the class properties.

Defining Expedited Forwarding Class Properties

To define Expedited Forward class properties:

1. In the SmartDashboard **Network Objects** tree, double-click the applicable Security Gateway.
2. In the **Gateway** window, click **Network Management**.
3. In the **Interface** window, click the **QoS** tab.
4. In the **DiffServ and Low Latency classes** section, click **Add** or **Edit.**
5. Click **DiffServ Classes > Expedited Forwarding**.
6. Configure these properties:
   - **Class:** Select a Low Latency class from the list of defined classes.
   - **Inbound:** Define the portion of the interface's inbound capacity to be reserved.
   - **Constant Bit Rate:** The constant bit rate at which packets of this class will be transmitted.
   - **Maximal Delay:** The maximum delay that will be tolerated for packets of this class. Those packets that exceed this delay are dropped.
   - **Outbound:** Define the portion of the interface's outbound capacity to be reserved by defining a **Constant Bit Rate** and a **Maximum Delay** as described above.

   You must configure at least one of the two directional properties (Inbound/Outbound), and you can configure both.
7. Click **OK**.

Defining DiffServ Class Properties

To define DiffServ class properties:

1. In SmartDashboard, locate the relevant Security Gateway.
2. In the **Gateway Properties** window, click **Network Management**.
3. In the **Interface** window, click the **QoS** tab.
4. In the **DiffServ and Low Latency classes** section, click **Add** or **Edit.**
5. Click **DiffServ Classes > Others.**
6. Configure these properties:
   - **Class:** Select a DiffServ class from the list of defined classes.
   - **Inbound:** Define the portion of the interface's inbound capacity to be reserved.
   - **Guaranteed bandwidth:** The bandwidth guaranteed to be marked with the QoS Class.
   - **Bandwidth Limit:** The upper limit of the bandwidth to be marked with the QoS Class. Traffic in excess of the **Bandwidth Limit:** will not be marked. For example, if the interface's capacity is 256MB and **Bandwidth Limit** to 192MB, then traffic beyond 192MB will not be marked.
   - **Outbound:** Define the portion of the interface's outbound capacity to be marked by defining a **Guaranteed Bandwidth** and a **Bandwidth Limit** as described above.
7. Click **OK**.

Working with Low Latency Queuing

QoS Low Latency Queuing makes it possible to define special classes of service for "delay sensitive" applications like voice and video. Rules under these classes can be used together with other rules in the QoS Policy Rule Base. Low Latency classes require you to specify the maximal delay that is tolerated and a Constant Bit Rate. QoS then guarantees that traffic matching rules of this type are forwarded within the limits of the bounded delay.

For more, see: Low Latency Queuing (on page 46).

Defining a Low Latency Class

To define a Low Latency class:

1. In SmartDashboard select **Manage** > **QoS** > **QoS Classes**.
2. In the **QoS Classes** window, click **New** > **Low Latency Class of Service**.
3. In the **Class of Service Properties** window, configure these class properties:
   - **Name -** The name of the Class of Service.
   - **Comment -**The text to be displayed when this class is selected in the **QoS Classes** window
   - **Color -** Select a color from the list.
   - **Type -** Select a type from the list.
4. Click **OK.**

Configuring an Interface for Low Latency

Use these procedures to configure interfaces to use a Low Latency or DiffServ Expedited Forwarding class.

To configure an interface for Low Latency:

1. Make sure that SmartDashboard is closed.
2. In SmartConsole, go to **Gateways & Servers**.
3. Double-click the applicable Security Gateway.
4. In the **Check Point Gateway** window, click **Network Management**.
5. Double-click the applicable interface.
6. In the **Interface** window, click the **QoS** tab.
7. Select **Inbound Active** and/or **Outbound Active** and set the **Rate** properties.
8. In the **Diffserv and Low Latency classes** section, click **Add** > **Low Latency Classes**.
9. In the **Low Latency QoS** window, select a class from the list.
10. Select **Inbound Active** and/or **Outbound Active**.

    Note: You must set at least one traffic direction to Active.
11. Configure these Low Latency properties:
    - **Constant Bit Rate** - The constant bit rate at which packets of this class will be transmitted.
    - **Maximal Delay** - The maximum delay allowed for packets of this class. Packets that exceed this value are dropped.

      Note: To configure an Expedited Forwarding interface to work as a DiffServ interface, set the Maximal Delay property to *99999*.

Do these steps for each applicable interface on a Security Gateway.

Defining Low Latency Class Properties

To define Low Latency class properties:

1. In SmartDashboard, click a Gateways & Servers and double click the applicable Security Gateway.
2. In the **Gateway** window, click **Network Management**.
3. In the **Interface** window, click the **QoS** tab.
4. In the **DiffServ and Low Latency classes** section, click **Add** or **Edit.**
5. Click **Low Latency**.
6. Configure these properties:
   - **Class:** Select a Low Latency class from the list of defined classes.
   - **Inbound:** Define the portion of the interface's inbound capacity to be reserved.
   - **Constant Bit Rate:** The constant bit rate at which packets of this class will be transmitted.
   - **Maximal Delay:** The maximum delay that will be tolerated for packets of this class. Those packets that exceed this delay are dropped.
   - **Outbound:** Define the portion of the interface's outbound capacity to be reserved by defining a **Constant Bit Rate** and a **Maximal Delay** as described above.

   You must configure at least one of the two directional properties (Inbound/Outbound), and you can configure both.
7. Click **OK**.

Working with Authenticated QoS

Authenticated QoS provides Quality of Service (QoS) for end-users in dynamic IP environments, such as remote access and DHCP environments. This enables priority users, such as corporate CEOs, to receive priority service when remotely connecting to corporate resources.

For more detailed information, please see Authenticated QoS (on page 52).

Using Authenticated QoS

To apply Authenticated QoS in a rule:

1. Make sure that the UAS package is installed on the Security Gateway that does Authenticated QoS.
2. Make sure that the **User Authority Server** option under **Check Point Products Installed** is selected on the Security Gateway on which you are installing the policy.
3. Open SmartDashboard.
4. Create a group in **Manage > Users > New > Group**.
5. In the **Group Properties** window, add all the priority users.
6. Create a rule.
7. In the **Source** column, right-click and select **Add object > Add legacy user access**.

   Best Practice - To minimize the resources taken up by Authenticated QoS, it is recommended that Authenticated QoS rules refer to specific services, and unless absolutely necessary, you should not include **Any** in the **Service** field.
8. Install the policy.

For example, if the CEO of your company is in a remote location and wants to access his email and without waiting too long, create a rule like this:

Rule Name	Source	Destination	Service	Action
CEO	CEO@localnet	Any	Pop-3	Weight 10 Guarantee 50,000 Bps

Notes -

- The user must be authenticated in the UAS, for the QoS policy to be enforced.
- Policy-wide properties for Authenticated QoS can be defined in the **QoS** page of the **Global Properties** window.

  For more information, see Defining QoS Global Properties (on page 59).

Managing QoS for Citrix ICA Applications

In order to deliver a QoS solution for the Citrix ICA protocol:

1. Disable session sharing in the **Citrix Program Neighborhood**.
2. Modify your Security Policy to allow the **Citrix_ICA and Citrix_ICA_Browsing services**.

   

   Note - The **Any** service does not include the Citrix ICA service.
3. Discover the Citrix application names, as defined by the Citrix Administrator, and retrieve your Citrix ICA application names from the log.

   This includes turning on the application detection check box and installing Security and QoS Policies.
4. Define new Citrix TCP services with the application names you have detected.
5. Add the appropriate Citrix TCP services to rules in your QoS Policy.
6. Install the QoS Policy.

Disabling Session Sharing

Citrix enables session sharing by default. In this mode, traffic from all the applications used by a specific client share the same TCP connection. In order for QoS to prioritize different Citrix ICA applications from the same client, you *must* disable session sharing. This means that every application uses a separate TCP connection (all going to the same server port, 1494, from different source ports).

You should contact the Citrix Administrator to configure the correct mode.

To Disable Session Sharing:

1. Double-click the **Citrix Program Neighborhood** icon placed on the desktop by the Citrix install program.
2. Click the **Settings** icon or, from the **File** menu, select **Application Set Settings**.
3. Select the **Default Options** tab.
4. From the **Window Size** list, select something other than **Seamless Window**.

Modifying your Security Policy

You must modify your Security Policy to enable the new Citrix_ICA TCP and Citrix_ICA_Browsing UDP services. The Citrix_ICA service initializes the stateful inspection of the Citrix ICA protocol. The Citrix_ICA service is not included in the Any service of the Security Policy and must therefore be enabled in one of these ways:

In the **Security** tab add a rule to your Security Policy with the Citrix_ICA TCP service. Similarly, add a rule for the Citrix_ICA_Browsing service. Alternatively, you can add simply add the Citrix_metaFrame group, which incorporates both the Citrix_ICA TCP and Citrix_ICA_Browsing UDP services.

Or:

1. Expand the TCP branch of the Services Tree.
2. Double-click on the **Citrix_ICA** service.

The **TCP Service Properties - Citrix _ICA** window opens.

3. Click **Advanced**.

The **Advanced TCP Service Properties** window opens.

4. Select **Match for Any** to turn on the Citrix ICA protocol inspection without having to add a specific rule for the Citrix_ICA service (if the **Any** service is allowed).

Discovering Citrix ICA Application Names

To discover the Citrix ICA application name, as defined by the Citrix Administrator, use QoS to snoop the wire and send logs (of type alert) to the log server, recording the Citrix ICA application name. The Citrix ICA application detection is turned off by default.

Note - The frequency of recording an application name log (alert) is 24 hours.

Advanced: If you want to reset the application detection cache in order to re log a Citrix ICA application on the wire even if it was logged in the past 24 hours use this command line instruction:

```
fw tab -t fg_new_citrix_app -x
```

To Enable Citrix ICA Application Name Logging:

1. Double-click on the Security Gateway in the Network Objects Tree. The **General Properties** window shows.
2. Select **Logs and Masters > Additional Logging** in the tree on the left side of the **General Properties** window. The **Additional Logging Configuration** window shows.
3. Select **Detect new Citrix ICA application names** to enable QoS to log the Citrix application names.
4. Click **OK**.

Citrix ICA application name detection is enabled.

5. Create a Security Policy with a valid rule that uses the Citrix_ICA service.
6. Install the QoS and Security policies on the QoS Security Gateway and let it run for a period of time. See Installing a QoS Policy.

Note - The QoS policy content is irrelevant to the application detection feature.

7. View the QoS log entries in SmartConsole (the entries are of Type Alert and contain the Citrix ICA application names). Once you have the application names you can turn off the application detection, as well as define new Citrix TCP services to use in a QoS policy.

Note - It is a pre-requisite that the Citrix_ICA TCP be enabled in the Security Policy.

To Disable Citrix ICA Application Name Logging:

1. Double-click on the Security Gateway in the Network Objects Tree.

The **Gateway - General Properties** window opens.

2. Select **Logs and Masters > Additional Logging**.

The **Additional Logging Configuration** page opens.

3. Clear the **Detect new Citrix ICA application names** option so that QoS will not log the Citrix application names.
4. Click **OK**.

Citrix ICA application name detection is disabled.

5. Install the QoS Policy.

Defining a New Citrix TCP Service

A new service type was introduced in the SmartConsole, Citrix TCP.

To Define a New Citrix TCP Service

1. Right-click on the **Citrix TCP** branch of the Services Tree, and select **New Citrix TCP**. The **Citrix Service Properties** window is displayed.
2. Enter the following details in the **Citrix Service Properties** window, as shown in the example below:
   - **Name**: The name of the new service.
   - **Comment**: A comment describing the new service.
   - **Color**: Select a color from the list.
   - **Application**: The exact name (case insensitive) of the Citrix Application.

Note - The application name is case insensitive.

3. Click **OK** to create the new Citrix Class of Service.

Adding a Citrix TCP Service to a Rule

In QoS policy mode, you can add a Citrix TCP service to a rule in the QoS Policy (see "Working with QoS Policies" on page 62).

Installing the Security and QoS Policies

Security and QoS Policies must be installed on the gateways before the policies are enforced.

Managing QoS for Citrix Printing

Printing generates relatively large quantities of data, causing a TCP connection to consume excessive quantities of bandwidth. This type of connection should be identified and the bandwidth made available to these connections limited.

There are three primary methods of printing in the MetaFrame environment:

- IP Network printing
- MetaFrame Auto-Creation of printers
- Local MetaFrame printing.

The QoS policy solution for printing traffic using the MetaFrame Auto-creation method is to classify each ICA connection as a printing or a non-printing connection.

A connection that is classified as printing is assigned to a Citrix printing rule. This rule can be configured to limit printing traffic and avoid excessive consumption of bandwidth. A connection that is classified as non-printing is assigned to a rule according to the regular matching method.

Classification of the connection is dynamic and is based on examining the ICA priority bits of each packet. An ICA connection is therefore matched dynamically to one of two different rules depending on the type of data passing through the connection at any point in time.

Best Practice - Limit the bandwidth per printing connection to a low value, depending on your network speed. This saves bandwidth for other traffic.

Citrix Printing rules are designed for slow networks that might have all their bandwidth consumed by printing connections.

Citrix printing rules are not supported when CoreXL or SecureXL acceleration technologies are enabled on the QoS gateway.

Configuring a Citrix Printing Rule

Define a printing rule to which all ICA connections that are in a printing state are assigned.

To Configure a Citrix Printing Rule

1. Position your cursor in the **Name** field of the **QoS** tab, at the position where you want to add a new rule.
2. Right-click and select one of the **Add Rule** options.

   The **Rule Name** window opens.
3. Enter the rule name in the **Rule Name** field.
4. Click **OK**.
5. Right-click in the **Service** column
6. Select **Add**.

   The **Add Object** window opens listing the network objects defined in the Security Policy and the QoS Policy.
7. Select the predefined **Citrix_ICA_printing** service.
8. Click **OK**.
9. Right-click in the **Action** column.
10. Select **Edit Properties**.

The **QoS Action Properties** window opens.

11. Select **Advanced** in the **Action Type** area.
12. Select **Per connection limit** in the **Limit** area.
13. In the **Per connection limit** field, enter a low value.
14. Click **OK**.

Viewing QoS Security Gateway Status

To see the QoS Security Gateway status, click **Security Gateway** in the **Gateways & Servers** view in SmartConsole. The status information shows on the **Summary** tab at the bottom of the view.

Display QoS Gateways Configured by SmartConsole

See SmartView Monitor. For more, see the *R80.10 Logging and Monitoring Administration Guide.*

Configuring QoS Topology

When the MetaFrame Auto-Creation of printers printing method is used, the Citrix printing traffic passes from the Citrix Server to the Citrix Client. To enforce QoS on this traffic, QoS must be installed

- On the Security Gateway external interface on the inbound direction or
- On the Security Gateway internal interface on the outbound direction.

Enabling Log Collection

In order for a connection to be logged, the QoS logging flag must be turned on and the connection's matching rule must be marked with either **Log** or **Account** in the **Track** field of the rule. For further information on how logging features work, see Overview of Logging (on page 84).

To Turn on QoS Logging

A QoS Security Gateway logs to the log if **Turn on QoS Logging** is checked in the **Additional Logging** page (under **Logs and Masters**) of the **Properties** window. By default, QoS Logging is turned on.

Confirming a Rule is logged

1. In SmartDashboard, select the rule whose connection will be logged.
2. Confirm that either **Log** or **Account** appear in the **Track** field (see "To Modify Tracking for a Rule" on page 71).

Logs & Monitor

In This Section:

This chapter shows you how configure rules to create logs for specified conditions. You can use the powerful Logs & Monitor features in SmartConsole to see logs and to monitor the effectiveness of QoS Policies.

Overview of Logging

These events are logged. The table below describes features unique to event logs.

Non-Accounting Log Events

Log Event	Data Returned	Presentation	Policy Mode
Connection Reject			
QoS rejects a connection when the number of guaranteed connections is exceeded and/or when you have configured the system not to accept additional connections.	The name of the matching rule on account of which the connection was rejected.	Generated as a reject log. Unified with the initial connection log.	Recommended policy only.
Running Out of Packet Buffers			
One of the interface-direction's packet buffers is exhausted. A report is generated a maximum of once per 12 hours.	A string explaining the nature of the problem and the size of the relevant pool.	New log record created each time a global problem is reported.	Recommended policy only.
LLQ Packet Drop			
When a packet is dropped from an LLQ connection. A report is generated a maximum of once per 5 minutes.	Logged data: • Number of bytes dropped due to delay expiration • Average packet delay • Jitter (maximum delay difference between two consecutive packets)	Unified with the initial connection log.	Recommended policy only.

The next table describes the features unique to accounting logs.

Explaining the Accounting Log

Logged	Data Returned	Policy Mode
General Statistics		
The total bytes transmitted through QoS for each relevant interface and direction.	Inbound and outbound bytes transmitted by QoS.	Recommended and Express policies.
Drop Policy Statistics		
• Total bytes dropped from the connection as a result of the QoS policy. • Count of the bytes dropped from the connection because the maximum used memory fragments for a single connection was exceeded.		Recommended policy mode only.
LLQ Statistics		
Statistics about the LLQ connection.	Logged data: • Number of bytes dropped due to delay expiration • Average packet delay • Jitter (maximum delay difference between two consecutive packets)	Recommended policy mode only.

These conditions must be met for a connection to be logged:

- The QoS logging checkbox must be selected in the Gateway Properties - Additional Logging Configuration window (see "To Turn on QoS Logging" on page 83). (By default this is automatically selected.)
- The connection's matching rule must be marked with either **Log** or **Account** in the **Track** field of the rule. See To Confirm that the Rule is Marked for Logging **(see "**Confirming a Rule is logged**" on page** 83**)** and To Modify Tracking for a Rule (on page 71).

Examples of Log Events

This section describes the log events.

Connection Reject Log

The connection is rejected because the rule exceeds the number of guaranteed connections, where **Accept additional non-guaranteed connections** is unchecked in the **QoS Action Properties** window (see QoS Action Properties (on page 27)). The log will include the name as well as the class of the rule in the following format: **rule_name:<***class***>-><***name***>**.

In the following example, the rule belongs to the class **Best_Effort**. The name of the rule (**rule_name**) is **udp2**.

Connection Reject Log — Example

Time	Product	Interface	Type	Action	Information
15:17:09	QoS	daemon	log	reject	rule_name:Best_Effort->udp2

LLQ Drop Log

When a packet from the LLQ connection is dropped, LLQ information is computed and logged from the *last* time a log was generated. This information includes *significant* data logged from the relevant interface-direction. In the following example, the information logged includes:

- **s_in_llq_drops:** The number of bytes dropped from the connection on the Server-In interface direction.
- **s_in_llq_avg_xmit_delay:** The average delay computed for all the connection's packets that were not dropped on the Server-In interface direction.
- **s_in_llq_max_delay:** The maximum delay of a connection packet that was not dropped on the Server-In interface direction.
- **s_in_llq_xmit_jitter:** The maximum delay difference between two consecutive *successfully* transmitted packets of the connection on the Server-In interface direction. Any packets which are dropped in between the two successfully transmitted packets are ignored.
- **s_in_llq_recommended_delay:** The default delay that can be entered into the **Add Low Latency QoS Class Properties** window in order to achieve a minimal number of dropped bytes.

LLQ Drop Log — Example

Product	Type	Information
QoS	log	s_in_llq_drops:3000 s_in_llq_avg_xmit_delay: 900 s_in_llq_max_delay: 1351 s_in_llq_xmit_jitter: 1351 s_in_llq_recommended_delay:2000

In the above example relevant data was observed only on the Server-In interface direction, therefore only Server-In counters are available.

Note - There are several reasons why logging might not occur on a specified interface direction:

- QoS might not be installed on all the interfaces directions.
- No packets were seen on other interface directions.
- Data on other interface directions might not be significant, for instance, the values logged might be zero.

Pool Exceeded Log

A log for when the designated size of the **ifdir** pool is exceeded. In this example, the log shows:

- An interface direction (**ifdir**) has a pool size of 8 fragments.
- The interface name is **E100B1**, and the direction is outbound (*outbound* shown by the cube with an outward pointing arrow).

Pool Exceeded Log — Example

Product	Interface	Type	Information
QoS	E100B1	control	info:Ifdir Memory Pool Exceeded Pool_size:8

Examples of Account Statistics Logs

Logs always include the **segment_time** information (the time from which the information about the log was gathered) in the **Information** column.

The Mandatory Fields in Account Logs

Product	Type	Information
QoS	Account	segment_time 8May2002 12:24:57

Account Logs may include any or all of the above information:

Note - Only significant data is logged and presented in the same log record.

General Statistics Data

These statistics include the number of bytes transmitted through QoS in any relevant interface direction. In the following example:

- **s_in_bytes:** 5768 bytes were transmitted through QoS on the Server-In interface direction.
- **s_out_bytes:** 154294 bytes were transmitted through QoS on the Server-Out interface direction.

General Statistics Data — Example

...	Information	...
	s_in_bytes:5768 s_out_bytes: 154294	

Drop Policy Statistics Data

The number of bytes dropped from the connection in any relevant interface direction as a result of drop policy are logged. The drop policy is aimed at managing QoS packet buffers, see WFRED (Weighted Flow Random Early Drop). This includes the total number of bytes dropped from the connection since it exceeded its allocation. In the following example:

- **s_out_total_drops:** 3914274 bytes were dropped from the connection as a result of drop policy, on the Server-Out interface direction.
- **s_out_exceed_drops:** Out of total number of drops **(s_out_total_drops)** 3914274 bytes were dropped from the connection because it exceeded its allowed number of fragments, on the Server-Out interface direction.

Drop Policy Statistics Data — Example

...	Information	...
	s_out_total_drops:3914274 s_out_exceed_drops: 3914274	

LLQ Statistics Data

Data items are the same as in LLQ Drop Log (on page 86), but are generated from the beginning of the connection, *not* from the last time a log was created.

Command Line Interface

In This Section:

QoS Commands

QoS Command	Description
etmstart	Starts QoS
etmstop	Stops QoS
fgd50	QoS daemon

Note: On Windows gateways, running `etmstop` can result in this error message: *The Check Point FloodGate-1 service could not be stopped*. This is caused by a too-short Windows service check timeout, not `etmstop` failure. To resolve, run `etmstop` again.

Setup

cpstart and cpstop

Generally, to stop and start the QoS Security Gateway you are required to stop the Firewall using the **cpstop** and **cpstart** commands. In the event that you would like to stop the QoS Security Gateway only, you can use the QoS specific **etmstart** and **etmstop** commands.

etmstart

etmstart loads the QoS Security Gateway, starts the QoS daemon (**fgd50**), and retrieves the last policy that was installed on the QoS Security Gateway.

etmstop

etmstop kills the QoS daemon (`fgd50`) and then unloads the QoS Policy and Security Gateway.

fgate Menu

Typing **fgate** on the command line shows this menu:

```
# fgate
Usage:
fgate load <rules-file.F> [targets]  # install targets
```

```
fgate unload [targets]               # uninstall targets
fgate fetch [-f | servers]           # fetch last policy installation
fgate stat [targets]                 # display status
fgate ver [-k]                       # display version
fgate log [args]                     # control logging
fgate debug <on | off>               # control daemon debug
fgate kill [-t sig_no] procname      # send signal to FloodGate-1 daemon
fgate fetch_robo [servers]           # fetch the robo-cluster policy


    [targets] and [servers] are lists of host names or IP addresses.Specifying no
target performs the operation locally.
```

Control

fgate

The **fgate** program is used to manage QoS. Its specific action is determined by the first command line argument, as described in the following sections:

fgate load

fgate load runs a verifier on the policy file. If the policy file is valid, **fgate** compiles and installs a QoS Policy to the specified QoS gateways. It can only be run from the Security Management Server.

1. **Syntax**

```
fgate load <rule-file.F> [targets]
```

If **targets** is not specified, the QoS Policy is installed on the local host.

fgate unload

fgate unload uninstalls a QoS Policy from the specified QoS gateways. It can only be run from both the Security Management Server and localhost.

1. **Syntax**

```
fgate unload [targets]
```

If **targets** is not specified, the QoS Policy is uninstalled from the local host.

fgate fetch

fgate fetch retrieves the QoS Policy that was last installed on the local host. You must specify the machine where the QoS Policy is found. Use "localhost" in case there is no Security Management Server or if the Security Management Server is down. You may specify a list of Security Management Servers, which will be searched in the order listed.

fgate fetch -f attempts to retrieve policies from all management stations, one after the other until it succeeds. If the Security Gateway fails to retrieve a policy from a Security Management Server, it tries to retrieve one from itself.

Syntax

```
fgate fetch [-f | servers]
```

Examples

```
fgate fetch localhost
fgate fetch -f
fgate fetch mgmt_server_name
```

Monitor

fgate stat

fgate stat displays the status of target hosts in various formats. If this command is launched from a Security Management Server, it can be run on more than one Security Gateway. If this command is launched from a Security Gateway, the status of the Security Gateway is returned.

Usage

```
fgate stat [targets]
```

The default format displays the following information for each host: product, version, build number, policy name (QoS policy mode, or Express mode), install time and interfaces number.

If no target is specified, the status of **localhost** is shown. For example:

```
# fgate stat

Blade:          QoS
Version:        R77.10
Kernel Build:   11
Policy:    Standard
Install time:   Wed Oct 23 12:30:33 2013
Interfaces Num: 1

Interface table
-----------------------------------------------------------------
|Name|Dir|Limit (Bps)|Avg Rate (Bps)|Conns|Pend pkts|Pend bytes|
-----------------------------------------------------------------
|eth0|in |    5625000|             0|    3|        0|         0|
|eth0|out|    5625000|            58|    2|        0|         0|
-----------------------------------------------------------------
```

Examples

```
fgate stat
fgate stat gateway1 gateway2
```

fgate ver

fgate ver displays the QoS version number. If the -k option is included, both the kernel version build number and QoS executable version build number are returned. Without the -k, only the QoS executable version is specified.

Syntax

```
fgate ver [-k]
```

Utilities

fgate log

fgate log turns logging on or off in the kernel. It can be used in order to save resources without reinstalling your QoS policy. The **stat** option returns the current state of logging.

Syntax

```
fgate log < on | off | stat >
```

By default, **fgate log** is turned on.

fgate debug

fgate debug turns on a debug flag which sends additional debugging information to the fgd log file: **$FGDIR/log/fgd.elg**. The default is off.

Syntax

```
fgate debug < on | off >
```

fgate kill

fgate kill sends a signal to a QoS daemon. The Security Management Server does not run the QoS daemon therefore this command is valid only on gateways.

Syntax

```
fgate kill [-t sig_no] proc-name
```

Parameter	Meaning
`[-t sig_no] proc-name`	If the file `$FWDIR/tmp/<proc-name>.pid` exists, send `sig_no` to the PID in the file. If no signal is specified, signal 15 (`sigterm`) is sent.

The QoS daemon writes the PIDs to files in the log directory upon startup. These files are named `$FWDIR/tmp/<daemon_name>.pid`. For example, the file containing the PID of the QoS SNMP daemon is `$FWDIR/log/snmpd.pid`.

Examples

Commands	Example and Description
`fgate kill`	• `fgate kill fgd` Sends signal 15 to the QoS fgd daemon. • fgate kill -t 1 fgd Sends signal 1 to the QoS fgd daemon.

Commands	Example and Description
`fgate fetch_robo`	• `fgate fetch_robo` Fetches the local robo-cluster policy • fgate fetch_robo [server] Fetches the robo-cluster policy from the given server

FAQ

In This Section:

QoS Basics

When should I use Recommended Policy type **and** when should I use **Express Policy** type**?** — Use the Recommended Policy type when you need fine-tuned functionality and advanced QoS features. Use Express if your system requires only basic QoS.

What are the benefits of using each mode? — Recommended gives you advanced QoS functionality. Express mode gives you better performance and requires less CPU and memory.

Can I change the Policy types? — You can change a policy type from Express to Recommended, but you cannot change Recommend to Express. We recommend that you start with Express if you are not certain. This way, you can change to Recommended if you require advanced QoS functionality.

What is the highest weight I can use in a rule? — Weights are relative. The only limitation is the **Maximum weight of rule** parameter, which is defined in the **Global Properties** window under QoS. The default parameter is 1000, but can be changed to any number.

Note - This parameter is only used to assist in input validation.

In the example shown here:

Example of Highest Weight Differentiation

Policy 1	HTTP gets	...and equals	Comment
HTTP weight = 500, FTP weight =500	500/(500+500)	= ½	Equal weight is given to each rule.
Policy 2			
HTTP weight = 2, FTP weight =2;	2/(2+2)	= ½	Equal weight is given to each

Policy 1	HTTP gets	...and equals	Comment
Policy 1 + third rule			
HTTP weight = 500, FTP weight =500, SMTP weight = 100	500/(500+500+100)	= 500/1100	Due to the initial high value of the weights in Policy 1, the amount of bandwidth available to the HTTP connection is only marginally less than in Policy 1 even after the introduction of the third rule.
Policy 2 + third rule			
HTTP weight = 2, FTP weight =2; SMTP weight = 100	2/(2+2+100)	= 2/104	Due to the low value of the weights in Policy 2, the amount of bandwidth available to the HTTP connection is now significantly less as a result of the introduction of the third rule.

You can see the significance of the value of the weight allocated in two different policies. In the example both the HTTP and FTP connections initially enjoy an equal share of the available bandwidth, although they each had a weight of 500 in Policy 1 and a weight of 2 in Policy 2.

By adding a third rule to both policies you can significantly change the result. For example, an SMTP connection with a weight of 100 can be added to each policy. Due to the high initial weights used in Policy 1, there is an insignificant change to the amount of bandwidth available for the HTTP connection in Policy 1 + third rule. However, due to the low initial weights used in Policy 2, the amount of bandwidth that is available to the HTTP connection in Policy 2 + third rule is significantly reduced.

Should I install QoS on the external or the internal interface? — While QoS can run on both interfaces, it is highly recommended to position QoS on the external interface only.

What is the difference between guarantees and weights? — Guarantees and weights are similar in their behavior. Despite the difference in their dictionary meaning, they both guarantee the allocated bandwidth to the matched traffic. The differences between them are:

- Guarantees are stated in absolute numbers (for example, 20000bps) and weights are stated in relative numbers (for example, 100).
- Guarantees are allocated their share of bandwidth before weights. For example if you have a link of 1.5 MB:

Your Rule Base is:

- HTTP Guarantee 1Mb
- FTP Weight 40
- SMTP Weight 10

The result is:

- first 1 MB for HTTP is allocated, *then*
- 0.4 MB for FTP is allocated *and* 0.1MB for SMTP is allocated.

Use guarantees to define bandwidth in absolute terms or for per connection guarantees.

How does QoS handle TCP retransmitted packets? — When a retransmission is detected, QoS checks to see if the retransmitted data is already contained in the QoS queue. If so, the packet is dropped. This unique QoS capability eliminates retransmissions that consume up to 40% of a WAN link, and saves memory required to store duplicated packets.

Which Firewall resources does QoS support in the Rule Base? — QoS can use its resources to inspect HTTP traffic. Resources are defined using the **URI for QoS** option and can contain specific URLs or files. For example, you can limit Web surfing to the site **http://www.restrict-access-to-this-site.com.** You need to add a QOS URI resource that looks for the string "**www.restrict-access-to-this-site.com**" (without http://). Then use the resource in a QoS rule and add a limit.

Do guarantees waste bandwidth? — No. QoS uses a sophisticated queuing mechanism. An application only takes as much bandwidth as it needs. Any unused bandwidth is then available for use by other applications.

How do I know if loaned bandwidth is available for applications that may need it back? — There is no loaned bandwidth in QoS. Bandwidth that is not utilized by a guarantee/weighted rule is immediately (on a per-packet basis) distributed to the other connections, according to their relative priorities. The important thing to remember is Resolution *(referring to level of granularity)*. QoS allocates bandwidth on a per packet basis. Therefore, only one packet is allocated at a time, resulting in the most accurate scheduling policy.

Other Check Point Products - Support and Management

Where is QoS placed in the Multi-Domain Security Management Inspection chain? — QoS is composed of two components:

- QoS Policy, which is in charge of rule matching
- QoS Scheduling, which is in charge of packet scheduling

Does QoS work With Multi-Domain Security Management? — Yes. One of the most important QoS features is its unique and sophisticated integration with Multi-Domain Security Management. Its integration features include:

- accurate classification of VPN traffic (inside the VPN tunnel)
- classification of NATed traffic
- shared network objects and topology (that save you time and effort in administration)
- common SmartDashboard with an advanced GUI but a familiar look and feel
- authenticated Quality of Service allows you to assign bandwidth to VPN remote users
- DiffServ Support and QoS bring Better than Frame Relay QoS to the VPN world
- log verification

Is SmartView Monitor a part of QoS? — No. As of NG with Application Intelligence (R55), SmartView Monitor is a separate product that is bundled with QoS.

Does QoS support Load Sharing configurations? — Yes, QoS supports all ClusterXL configurations. QoS supports the SYNC mechanism and therefore can be used with CPLS/CPHA or third-party solutions. For OPSEC partner solutions, see the OPSEC Website.

Does QoS support NATed traffic? — QoS has full support for NATed traffic, including matching, scheduling, limiting and all other QoS features.

What is the maximum number of QoS gateways I can manage? — QoS Security Gateway management is identical to that for any Security Gateway. Thus, the maximum number of gateways is identical to the maximum number of gateways that are managed.

Do I need to run QoS on the Security Management Server? — Yes, in order to manage a QoS Security Gateway you need to install QoS on the Security Management Server.

Policy Creation

When should I use LLQ (Low Latency Queuing)? — LLQ is best suited for VoIP applications, Video conferencing and other multimedia applications. LLQ is targeted for applications where:

- a minimal guaranteed bandwidth is required for adequate performance
- low delay and Jitter are required

Is QoS Rule Base "first match"? — From QoS NG forward, all QoS rules are matched on the "first match" principle. Meaning that only the first rule that applies to a connection is activated.

For example, if you have a rule for CEO traffic and a rule for HTTP traffic, the rule that appears first within the Rule Base will be matched to all CEO surfing.

Correct Rule Base (CEO is the first match)

1. SRC=CEO => Guarantee = 128Kbps
2. Service=HTTP => Limit = 64Kbps

Incorrect Rule Base (CEO traffic will be limited)

1. Service=HTTP => Limit = 64Kbps
2. SRC=CEO => Guarantee = 128Kbps

I am using QoS on multiple gateways. What is the best way to organize my Rule Base?

- If you are managing gateways with identical bandwidth and you want an identical policy for all gateways, define as **All** in the **Install On** field.
- If you are managing gateways with varied bandwidths and want an identical policy for all gateways, you can have one policy installed on all gateways. It is best to use weights since they assign relative bandwidth and not a fixed one. Remember that weights also guarantee bandwidth allocation.
- If you are managing gateways with varied bandwidths and want a different policy for all gateways, you can use different sub-rules for each Security Gateway. You can also use common rules that are matched for gateways.

When should I use Sub-rules? — Sub-rules should be used when there is hierarchy between objects. For example, when you want to manage bandwidth according to organizational structure, such as within an organization that has R&D, Marketing and operation divisions.

How can I see the top bandwidth-hogging applications? — From the command line run the command **rtmtopsvc**.

Capacity Planning

What are the QoS memory requirements? — To run QoS, the following amount of free memory is needed (in addition to the memory needed for Multi-Domain Security Management):

QoS memory requirements

Number of connections	Management	Gateway (or Management and gateway)
5,000	0 MB	32.5 MB
10,000	0 MB	39 MB
25,000	0 MB	57 MB
50,000	0 MB	91 MB
100,000	0 MB	156 MB

- These numbers include SmartView Monitor and UserAuthority.
- Connections are counted in the Firewall connection table.
- Note that the default size for the connection table is 25,000.
- On an average, each connection requires 1300 bytes.

How do I know which machine I need to run QoS? — Deciding on a hardware platform and vendors involves many aspects and each buyer has their own specific considerations such as support, price, appliances, knowledge, and so on.

As far as performance is concerned, CPU performance is the main factor in QoS performance. The reduced memory footprint and low memory prices, memory should not usually be the cause of a bottleneck.

How do I tune QoS performance? — Here are some tips on fine-tuning QoS performance:

1. Upgrade to the newest QoS version available. Major improvements in performance have been introduced in
   QoS NG FP1 and NG FP2.
2. In most cases you need to install QoS only on the external interfaces of the gateway.
3. Unless you are using limits for inbound traffic, installing QoS only in the outbound direction will provide you with most of the functionality and improvements.
4. Put more frequent rules at the top of your Rule Base. You can use SmartView Monitor to analyze how much a rule is used.
5. Turn "per connection limits" into "per rule limits".
6. Turn "per connection guarantees" into "per rule guarantees".

What is the maximum bandwidth supported by QoS? — QoS NG FP1 can support up to 1.13MBps and 890MBps (in Traditional Mode) of traffic of long UDP packets. In real-world traffic and in the Rule Base, QoS supports 330MBps (In Express Mode) and 255 MBps (in Traditional Mode) of traffic.

.

Installation/Backward Compatibility/Licensing/Versions

When will QoS next feature pack be available? — QoS feature packs/releases are usually shipped at the same time Multi-Domain Security Management feature packs are released.

How do I?

How do I guarantee performance for my mail server? — You need to add a rule matching your email traffic. You can do this by either matching the source/destination of your mail server, or matching mail protocols (SMTP, POP3, Exchange). For this rule, define a weight or guarantee that meets the needs of the priorities you want to set.

How do I ensure Quality of Service for Voice Over IP? — QoS FP1 introduced the VoIP-tuned mechanism Low Latency Queuing (LLQ). This mechanism is tuned to achieve best latency for constant bit rate applications, like VoIP.

To limit the number of connections admitted, use LLQ with a per connection guarantee. For voice, you want to give each conversation a guaranteed bandwidth. Usually you would want an admission policy that does not accept additional calls if bandwidth is not adequate.

Note - This is equivalent to the busy tone in old voice system.

How can I prioritize traffic for remote users? — Using the Authenticated QoS feature of QoS, you can prioritize bandwidth allocation for remote VPN users and Windows domain user groups.

How do I guarantee performance for my ERP applications? — You need to add a rule matching your ERP traffic. You can do this by either matching the source/destination of your ERP server, or matching application protocols (SAP, BAAN, ORACLE). For this rule, define a weight or guarantee that meets the needs of the priorities you want to set. If your ERP application is not a predefined service, you can either add it manually or use the first method.

If you are using ERP over HTTP, check "**How can I provide bandwidth for my intranet applications**"?

Can I use QoS to prevent Denial of Service Attacks? — QoS is not an Anti-Denial of Service tool. However, there are many situations in which QoS can be used to detect, monitor and prevent such attacks. Using SmartView Monitor and QoS you can perform detection and monitoring.

Prevention can be achieved in the following ways:

- by limiting applications that are known to be a part of DOS attacks (for example, ICMP, suspicious URLs).
- by providing guarantees for important traffic (for example, ERP, MAIL, VoIP).
- by providing guaranteed bandwidth for authenticated users using Authenticated QoS. Authenticated users can be identified with digital signatures and can rely on VPN authentication and encryption. QoS guarantees that these users will get their bandwidth. The attacker cannot authenticate to the VPN and will not get bandwidth for the attack.

Why is limiting bandwidth for Napster better than blocking it? — Blocking "non-work related" applications might cause users to find a way to bypass blocking. Prioritizing bandwidth lets users continue with their activities without damaging critical business processes. Consider a university where the Internet connection is being used for peer-to-peer file downloads like Napster and Kazaa. Blocking these services completely may encourage the students find a smart way to

bypass the block, which in turn might cause legal problems. QoS offers smarter solutions:

- Limiting the allocated bandwidth for such applications – this can be done with or without the students' knowledge.
- Limiting the allocated bandwidth during daytime, and providing more bandwidth at night.
- Providing guarantees to important users (Professors, MIS) while allowing students to use the reminder of the bandwidth.

General Issues

My machine is experiencing certain technical failures. What should I do? — Check the Web for updated release notes on known issues and limitations. Contact your vendor for further support.

I set up a guarantee/limit but in SmartView Monitor it seems to be broken? — If you are looking at very low traffic limit (for example, 1000 Bytes per second) at a high frequency (update every 2 seconds) it might look, as if the limit is broken since QoS does not fragment packets. If you lower the sampling frequency of SmartView Monitor (update every 8 seconds) you will see that limits are kept.

Can QoS prompt a user for authentication in order to use the Authenticated QoS feature? — No. In order to use Authenticated QoS, Multi-Domain Security Management must perform an authentication session prior to the classification of the connection by QoS.

Can I deploy QoS on LAN environments? — Yes. You will need to position the hardware to support the network traffic you want to prioritize. QoS is best deployed in congestion points for network traffic.

What happens if a line's bandwidth (as defined in the QoS tab of the Interface Properties window) is less than its physical ("real") bandwidth? — QoS will only allocate as much bandwidth as is defined in the **Interface Properties** window. Additional bandwidth will not be allocated regardless of the physical bandwidth of the interface.

What happens if a link bandwidth (of the link defined in QoS) is more than its physical ("real") bandwidth? — QoS will attempt to transmit more than the physical bandwidth allows. This can cause random traffic drops in the next hop that result in the loss of critical packets.

Debug Flags

In This Section:

Error and Debug Codes for QoS

Note:

- Error is turned on by default
- All commands begin with: `fw ctl debug -m fg +`

Command Line	Usage
driver	Driver Attachment
error	General Error Flag
chain	Main Steps Of QoS Packet Processing
install	For Future Use
pkt	Packet recording mechanism
citrix	Citrix processing
ls	Load sharing
tcp	TCP Retransmission Detection
sched	Packet Scheduling
policy	QOS Policy Rules Matching
url	QOS URL Matching
dns	DNS Related Messages
rtm	SmartView Monitor Interaction
auth	Authenticated QOS
log	Logging
conn	Connections Processing
drops	Drop Policy
rates	Reporting Rule/Connection Rates

Command Line	Usage
dropsv	Verbose Version Of Drop Policy
timers	Timer Events
chainq	Internal Chain Q Mechanism
llq	Low Latency Queuing
verbose	Used With Other Flags - Adds More Information
automatch	Report Matching Process (Debug Version Only)
autosched	Report Scheduling Process (Debug Version Only), a good way of reporting of rates on rules
qosaccel	Report SecureXL related data to QoS
multik	Report CoreXL related data to QoS

Appendix: Regular Expressions

In This Section:

Regular Expression Syntax

This table shows the Check Point implementation of standard regular expression metacharacters.

Metacharacter	Name	Description
\	Backslash	escape metacharacters non-printable characters character types
[ ]	Square Brackets	character class definition
( )	Parenthesis	sub-pattern, to use metacharacters on the enclosed string
{min[,max]}	Curly Brackets	min/max quantifier {n} - exactly n occurrences {n,m} - from n to m occurrences {n,} - at least n occurrences
.	Dot	match any character
?	Question Mark	zero or one occurrences (equals {0,1})
*	Asterisk	zero or more occurrences of preceding character
+	Plus Sign	one or more occurrences (equals {1,})
|	Vertical Bar	alternative
^	Circumflex	anchor pattern to beginning of buffer (usually a word)
$	Dollar	anchor pattern to end of buffer (usually a word)
-	hyphen	range in character class

Using Non-Printable Characters

To use non-printable characters in patterns, escape the reserved character set.

Character	Description
\a	alarm; the BEL character (hex 07)
\cx	"control-x", where x is any character
\e	escape (hex 1B)
\f	formfeed (hex 0C)
\n	newline (hex 0A)
\r	carriage return (hex 0D)
\t	tab (hex 09)
\ddd	character with octal code ddd
\xhh	character with hex code hh

Using Character Types

To specify types of characters in patterns, escape the reserved character.

Character	Description
\d	any decimal digit [0-9]
\D	any character that is not a decimal digit
\s	any whitespace character
\S	any character that is not whitespace
\w	any word character (underscore or alphanumeric character)
\W	any non-word character (not underscore or alphanumeric)

Disabling QoS Acceleration Support

If you have a QoS policy created for R77 and earlier, you will have to disable QoS acceleration to use other features ("Acceleration Support for R77 Policies" on page 15).

To manually disable QoS acceleration:

1. On the Security Gateway, run: `cpconfig` to turn off SecureXL and CoreXL.
2. Reboot the Security Gateway.
3. After reboot, run:

```
cpprod_util CPPROD_SetValue FG1 FgWithAcceleration 1 0 1
```

To manually enable QoS acceleration:

1. On the Security Gateway, run:
```
cpprod_util CPPROD_SetValue FG1 FgWithAcceleration 1 1 1
```
2. Use `cpconfig` to turn on SecureXL/CoreXL.
3. Reboot the Security Gateway.