

CHAPTER 5: CHECK POINT QoS

.....

Quality of Service (QoS) is a set of intelligent network protocols and services used to efficiently manage the movement of information through local and wide-area networks. QoS allows Security Administrators to prioritize traffic flows and provide better service to certain flows. Check Point QoS is included in every VPN-1 NGX product installation.

Objectives

Given a variety of Check Point QoS configurations, determine how to allocate bandwidth.

1. Determine if Check Point QoS is an appropriate solution, given a variety of business scenarios.
2. Configure Check Point QoS to meet the requirements, given a variety of business requirements.
3. Determine how bandwidth will be allocated, given a variety of Check Point QoS configurations.
4. Identify situations where Low Latency Queuing and Differentiated Services are an appropriate part of a QoS solution.

▪
▪
▪
▪
▪

Key Terms

- Check Point QoS
- IQ Engine
- WFRED
- RDED
- Differentiated Services
- Weighted Fair Queuing
- Low Latency Queuing



CHECK POINT QOS OVERVIEW

Check Point QoS is a bandwidth-management solution used to alleviate bandwidth congestion at network access points. Check Point QoS controls both inbound and outbound traffic flows. Check Point QoS uses four technologies to control traffic:

- Stateful Inspection
- Intelligent Queuing Engine (IQ Engine)
- Weighted Flow Random Early Drop (WFRED)
- Retransmission Detection Early Drop (RDED)

STATEFUL INSPECTION

Check Point QoS uses Check Point's patented Stateful Inspection technology to derive complete state and context information for all network traffic.

INTELLIGENT QUEUING ENGINE (IQ ENGINE)

Check Point QoS's **IQ Engine** uses state-derived information to classify traffic and place it in the proper transmission queue. Check Point QoS uses a packet scheduler to move packets through a dynamically changing scheduling tree at different rates, in accordance with the QoS Policy.

WEIGHTED FLOW RANDOM EARLY DROP (WFRED)

WFRED is a mechanism for managing packet buffers, by selectively dropping packets during periods of network congestion. WFRED is transparent to users and requires no configuration.

RETRANSMISSION DETECTION EARLY DROP (RDED)

RDED is a mechanism for reducing the number of retransmits and retransmit storms during periods of network congestion. 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.

Check Point QoS Architecture

The architecture and flow of Check Point QoS is similar to VPN-1 NGX. Check Point QoS has three components:

SmartConsole — configures and monitor Check Point QoS

SmartCenter Server — stores and distribute the Check Point QoS Policy

QoS Module — enforces Check Point QoS Policy

Like other NGX products, the Check Point QoS components can be installed in stand-alone or distributed deployments.

SMARTCONSOLE AND CHECK POINT QOS

The SmartConsole suite of tools allows a Security Administrator to configure, verify, install, and monitor the Check Point QoS Policy. SmartDashboard is used to configure the Policy. SmartView Tracker, SmartView Monitor, and Eventia Reporter provide useful information that can be used to monitor and tune the Policy.



The QoS information in SmartView Monitor is provided by interface. When reviewing Traffic > Top QoS Rules, and interface must be selected. If an interface is not selected, SmartView Monitor will display an error message.

Check Point's integrated architecture allows Security Administrators to reuse objects created for Security Policies when defining QoS Policies in SmartDashboard. A single SmartCenter Server can control and monitor multiple QoS modules.

SMARTCENTER SERVER AND CHECK POINT QOS

After the QoS Policy is defined in SmartDashboard, it is downloaded to the SmartCenter Server. The SmartCenter Server verifies the Policy and distributes it to QoS Modules. Policy distribution is handled by the Check Point Daemon (CPD), which runs on both the SmartCenter Server and the QoS Module. The SmartCenter Server also manages the Check Point Log Repository, and acts as a log server for SmartView Tracker.



QOS MODULE

The QoS Module enforces the QoS Policy and provides log information to the SmartCenter Server. The QoS Module's functions are divided between the QoS kernel driver and the QoS daemon. The QoS kernel driver examines, queues, schedules, and releases packets. The QoS daemon is a user-mode process that performs tasks that are difficult for the kernel. The QoS daemon performs the following tasks for the kernel:

- DNS name resolution
- Resolving authenticated data for an IP (UserAuthority integration)
- Updating the kernel regarding cluster status (in a Load Sharing configuration)



The Check Point QoS module must be installed with VPN-1 NGX. Both products share similar architecture, and Check Point QoS relies on VPN-1 NGX technology components to perform its bandwidth-management functions.

Check Point QoS Deployment Considerations

A distributed deployment scenario is strongly recommended for Check Point QoS. Distributed deployments scale gracefully, and dedicated servers reduce the likelihood that competition for server resources will create a bottleneck. CPU performance is the main factor affecting the performance of Check Point QoS modules. A server-class system is strongly suggested. Memory requirements are based on the number of connections the Check Point QoS Module is expected to manage. The table below provides guidelines for memory. On average, each connection requires 1,300 bytes of memory.

Number of Connections	Suggested Memory
10,000	39 MB
25,000	57 MB
50,000	91 MB
100,000	156 MB

Check Point QoS Rule Base

The Check Point QoS Rule Base can be configured in either Express or Traditional mode. Express mode allows a Security Administrator to define basic Policies quickly and easily. Traditional mode incorporates the more advanced features of Check Point QoS.

EXPRESS MODE FEATURES

The following Check Point QoS features are available in Express mode:

- Weights
- Per Rule Guarantees
- Per Rule Limits
- Logging & Accounting
- Hardware Accelerator Support
- High Availability & Load Sharing

TRADITIONAL MODE FEATURES

In addition to all the features available in Express mode, Traditional mode also has the following features available:

- Authenticated QoS (UserAuthority integration)
- Per Connection Guarantees and Limits
- Low Latency Queuing
- Differentiated Services Support
- Sub-Rules
- Matching by URI Resources
- Matching by DNS String
- TCP Retransmission Detection (RDED) Mechanism
- Matching Citrix ICA Applications

QoS Action Properties

The fields in QoS rules can be thought of defining:

- Who or What
 - Source
 - Destination
 - Service
- When
 - Time
- How
 - Action

The Action fields in the QoS rules determine how bandwidth will be allocated. Most actions allow for dynamic allocation of bandwidth based on current connections. Low Latency Queuing is an exception, and it is discussed later in this chapter.

SIMPLE AND ADVANCED ACTION TYPES

When a Security Administrator edits the Action field of a QoS rule, either Simple or Advanced Action Types may be selected. The configuration options for Simple Action type allow a Security Administrator to:

- Apply the rule only to encrypted traffic.
- Set a weight for the rule.
- Set a limit for the rule.
- Set a guarantee for the rule.

The Advanced action type selection adds configuration options for setting per-connection limits and guarantees.



The Advanced action type option is only available for Traditional mode, not Express mode. If per-connection limits and guarantees are needed, convert the Policy to Traditional mode, using the Policy > Convert To option from the main menu.

WEIGHTS

Weights are used to set the relative priority for different traffic flows. Weight is the relative portion of the available bandwidth allocated to a rule or connection. Bandwidth is allocated dynamically, and only open connections are included in the calculation. The formula for determining a rule's portion of the available bandwidth is:

this rule's portion = this rule's weight / total weight of all rules with open connections

Weights can only be set per rule, not per connection.

GUARANTEES

A guarantee allocates a minimum amount of bandwidth to the connections matched with a rule. Guarantees can be defined per rule and per connection. When Guarantees are specified per connection, the Security Administrator should also determine how many connections will be permitted. Weights ensure proportional shares, but only guarantees allow a Security Administrator to specify an absolute bandwidth value. Guaranteed bandwidth is allocated before bandwidth is distributed according to weights.

LIMITS

Limits define the maximum bandwidth that can be assigned to connections matching a rule. A limit defines a point beyond which connections under a rule are not allocated bandwidth. A limit can be defined for the entire rule, thus setting an absolute maximum quantity of bandwidth the rule may consume. Limits can also be set for each connection matched with a rule.

DEFAULT RULE

A default rule is automatically added to each QoS Policy Rule Base. The weight of the default rule is defined on the QoS screen in Global Properties. The default rule applies to all connections not matched by other rules or sub-rules. The default rule can be modified, but it cannot be deleted.

Bandwidth Allocation and Rules

Figuring out the relationship among weights, guarantees, limits, and open connections is similar to solving a logic problem. The following sections provide examples that start very simply and progress in complexity. Each example includes QoS Rule Base information, open connections, available bandwidth, and an explanation of how to determine how much bandwidth a particular connection will be allocated.



Check Point QoS ensures optimal bandwidth use by allocating bandwidth dynamically. Each example provided is a “snapshot”. In a live network, connections open and close constantly. By the time a Security Administrator figured out per-connection allocation, the allocations would have changed. However, the exercises are still valuable, because they allow you to see how bandwidth is allocated.

SIMPLE QOS RULE BASE WITH WEIGHTS ONLY

Adam is the Security Administrator at ABC Company. ABC Company has only two networks, Execu_Net and Other_Net. Adam has been told by upper management that traffic from Execu_Net needs to be given priority over traffic from Other_Net.

Adam configures a simple QoS Policy with the following rules:

Name	Executive Net Rule	Default
Source	Execu_Net	Any
Destination	Any	Any
Service	Any	Any
Action	Weight 30	Weight 10
Track	Log	Log
Install On	Any	Any
Time	Any	Any
Comment	Give Execu_Net Priority	Blank

The total bandwidth available to the QoS Module is 190 KBps.

The current open connections on the QoS Module are:

- 8 HTTP connections from Execu_Net.
- 3 SMTP connections from Other_Net.
- 2 RealAudio connections from Execu_Net.

Bandwidth is distributed among the open connections according to weight, using the formula:

this rule's portion = this rule's weight / total weight of all rules with open connections

In this example, connections match both rules, so the total weight of all rules with open connections is 40 (Executive Net Rule = 30, Default Rule = 10).

The bandwidth allocated to the Executive Net Rule is:

$$(30/40) * 190 = 142.5 \text{ KBps}$$

The 142.5 KBps allocated to the Executive Net Rule is shared among the 8 HTTP sessions and the 2 RealAudio sessions.

The bandwidth allocated to the Default Rule is:

$$(10/40) * 190 = 47.5 \text{ KBps}$$

The 47.5 KBps is shared among the 3 SMTP sessions from Other_Net.

Q&A

Q.) How would bandwidth be allocated if all connections originated from Execu_Net?

A.) All 190 KBps would be shared among the connections from Execu_Net. If no open connections match a QoS rule, that rule is not allocated any bandwidth.

Q.) How would bandwidth be allocated, if no open connections matched the Executive Net Rule?

A.) All bandwidth would be allocated to the Default Rule.



The bandwidth allocated to the IMAP Rule is:

$$(50/70) * 260 \text{ KBps} = 185.7 \text{ KBps}$$

The 185.7 KBps allocated to the IMAP Rule is shared among the 55 connections matching that rule.

The bandwidth allocated to the Marketing rule is:

$$(20/70) * 260 \text{ KBps} = 74.3 \text{ KBps}$$

The 74.3 KBps allocated to the Marketing Rule is shared among the 57 connections matching the rule.

Q&A

Q.) How would the bandwidth allocation change if all other connections stayed open, but an additional connection was opened from the Engineering_Net using the Telnet service?

A.) The Telnet session from the Engineering_Net would match the Default Rule. Check Point QoS would dynamically reallocate the bandwidth to provide the Default Rule its portion of bandwidth. The total weight of all rules with open connections becomes 80. Bandwidth is allocated by weight:

$$(50/80) * 260 \text{ KBps} = 162.5 \text{ KBps for the IMAP Rule}$$

$$(20/80) * 260 \text{ KBps} = 65 \text{ KBps for the Marketing Rule}$$

$$(10/80) * 260 \text{ KBps} = 32.5 \text{ KBps for the Default Rule}$$

Q.) How is the bandwidth divided, by connection, in the question above?

A.) The connections matching each rule share the allocated bandwidth, so:

55 connections share the 162.5 KBps allocated to the IMAP Rule.

57 connections share the 65 KBps allocated to the Marketing Rule.

1 connection gets the entire 32.5 KBps allocated to the Default Rule.

Bandwidth is allocated to rules with guarantees before it is distributed by weights. Allocation by guarantee only happens if there is an open connection matching the rule. In this example, the RealAudio Rule has a guarantee, and there are 10 open connections matching that rule. First the QoS Module allocates 100 MBps to the RealAudio Rule. The remaining 400 MBps is divided among all rules with open connections, including the RealAudio Rule, by weight.

All 4 rules have matching connections, so the total weight of all rules with open connections is 70. Bandwidth distribution by weight occurs next:

$[(10/70) * 400 \text{ MBps}] + 100 \text{ MBps} = 157.1 \text{ MBps}$ for the RealAudio Rule

$(30/70) * 400 \text{ MBps} = 171.4 \text{ MBps}$ for the IMAP Rule

$(20/70) * 400 \text{ MBps} = 114.3 \text{ MBps}$ for the Execu_Net Rule

$(10/70) * 400 \text{ MBps} = 57.1 \text{ MBps}$ for the Default Rule

Q&A

Q.) How many connections are sharing the Execu_Net Rule's 114.3 MBps?

A.) 9 HTTP connections match the Execu_Net Rule. The RealAudio and IMAP connections from the Execu_Net network matched on earlier rules.

Q.) If all the RealAudio connections are terminated and all other connections remain open, how is the bandwidth redistributed?

A.) The guarantee is not implemented, and the RealAudio Rule's weight is not included in the total weight of all rules with open connections, thus:

$(30/60) * 500 \text{ MBps} = 250 \text{ MBps}$ for the IMAP Rule

$(20/60) * 500 \text{ MBps} = 166.7 \text{ MBps}$ for the Execu_Net Rule

$(10/60) * 500 \text{ MBps} = 83.3 \text{ MBps}$ for the Default Rule

Q&A

Q.) If all 4 H.323 connections remained open, and the traffic flows matching the Default Rule all closed, how would bandwidth be allocated?

A.) The remaining 6,144 KBps would be shared among the 4 connections matching the H.323 Rule.

Q.) Since Debbie did not allow additional connections in the Action properties, what happens if someone attempts to establish a 5th H.323 connection?

A.) No bandwidth will be allocated to the 5th H.323 connection.

If Debbie checks “Accept Additional Connections” in the H.323 Rule’s Action properties, bandwidth will be allocated to additional connections by weight. So, if 4 H.323 connections have guaranteed bandwidth of 512 KBps, and 2 additional connections are opened, and connections match the Default Rule, bandwidth will be allocated:

1. The 4 guaranteed H.323 connections will be given their 512 KBps.
2. The Default Rule will receive it’s 90% of the remaining bandwidth (5,529.6 KBps).
3. The remaining 10% of the bandwidth (614.4 KBps) will be shared among the 6 H.323 connections:

614.4 KBps for the first 4 H.323 connections and

102.4 KBps for the next 2 H.323 connections

Since Debbie’s research indicates 384-512 KBps are required to sustain minimum quality for the H.323 sessions, she should consider either increasing the weight of the H.323 Rule or not allowing additional connections.



The sum of guarantees in rules in the upper level should not exceed 90 percent of the capacity of the link.

QOS BY SOURCE AND SERVICE WITH LIMITS AND WEIGHTS

Edward, the Security Administrator for MNO Corporation, must ensure that the total bandwidth consumed by FTP traffic never exceeds 250 KBps. Also, traffic from Execu_Net should receive higher priority than other traffic.

Edward configures a QoS Policy with the following rules:

Name	FTP_Limit	Execu_Net	Default
Source	Any	Execu_Net	Any
Destination	Any	Any	Any
Service	FTP	Any	Any
Action	Weight 10 Limit 250 KBps	Weight 20	Weight 10
Track	None	None	None
Install On	Any	Any	Any
Time	Any	Any	Any
Comment	Blank	Blank	Blank

The total bandwidth available to the QoS Module is 2 MBps (2,048 KBps).

If connections match all 3 rules in the QoS Policy, bandwidth will be allocated as follows:

250 KBps for the FTP_Limit Rule, because its allocation cannot exceed the defined limit

$(2/3) * 1,798 = 1,198.7$ KBps for the Execu_Net Rule

$(1/3) * 1,798 = 599.3$ KBps for the Default Rule



If the bandwidth allocated by weight is less than the limit, the rule will receive the lower allocation. In the example above, if Edward's QoS Module had only 512 KBps total available bandwidth, and all three rules had open connections matching them, the FTP_Limit Rule would have been allocated 128 KBps.



QOS BY SOURCE WITH WEIGHTS AND PER-CONNECTION LIMITS

Edward analyzes the traffic at MNO Corporation, and determines that limiting each FTP connection to no more than 20 KBps will improve his compliance directives from management. He changes the FTP_Limit Rule to read as follows:

Name	FTP_Limit
Source	Any
Destination	Any
Service	FTP
Action	Weight 10 Limit 250 KBps LC 20 KBps
Track	None
Install On	Any
Time	Any
Comment	Blank

Now any single FTP connection will be able to consume no more than 20 KBps, and the sum total of all FTP connections will be able to consume no more than 250 KBps.

When per-connection limits are imposed, the number of connections matching the rule with per-connection limits has an impact on the total bandwidth remaining for allocation by weight. For example, if only 2 FTP connections were open, then only 40 KBps would be allocated to the FTP_Limit Rule. The remaining 2,008 KBps would be distributed among connections matching the Execu_Net and Default Rules as follows.

GUARANTEE AND LIMIT INTERACTION

If a rule limit and a guarantee per rule are defined in a rule, the limit should not be smaller than the guarantee. A combination of guarantees and limits can be used instead of very low weights, to ensure connections get the bandwidth they need but do not consume a large amount of the bandwidth allocated in the weight distribution.

Recall the example with Debbie and JKL corporation. Debbie’s research indicated that each H.323 connection needed from 384-512 KBps to sustain optimal quality. Debbie could use a per-connection guarantee to make sure each H.323 connection received at least 384 KBps, and a per-connection limit to make sure none received more than 512 KBps.

Debbie configures a QoS Policy with the following rules:

Name	Video_Conferencing	Default
Source	Any	Any
Destination	Any	Any
Service	H.323	Any
Action	Weight 10 PC 384 KBps LC 512 KBps	Weight 10
Track	None	None
Install On	Any	Any
Time	Any	Any
Comment	Blank	Blank

The total bandwidth available to the QoS Module is 8 MBps (8,192 KBps).

Debbie only provides the per-connection guarantee for 4 connections, but she allows additional connections in the Action properties. Notice that she also changed the weight for the Video_Conferencing Rule from 1 to 10, so it is now equal to the weight for the Default Rule. If there are connections matching both the Video_Conferencing Rule and the Default Rule, and there are 6 H.323 sessions open, bandwidth will be allocated as follows:

1. 384 KBps to each of the first 4 H.323 sessions (total of 1,536 KBps)
2. All 6 H.323 connections will be taken up to 512 KBps, provided it does not exceed the bandwidth allocation by weight. In this case, the by-weight allocation would have been 3,328 KBps, but only 1,536 KBps is required to take each connection to its 512 KBps limit.
3. The remaining 5,120 KBps is allocated to open connections matching the Default Rule.

BANDWIDTH ALLOCATION AND SUB-RULES

Sub-rules are rules within a rule. Sub-rules allow a Security Administrator even more granular control over bandwidth allocation, but permit weights, guarantees, and limits within top-level rules. Sub-rules are only available in Traditional mode.



When a sub-rule is created for a top-level rule, a default rule for the sub-rule set is also generated. The default rule for the sub-rule set determines the fate of traffic that does not match any of the defined sub-rules in the set.

George is the Security Administrator at PQR Corporation. PQR Corporation provides a wide variety of services to its customers, via Web applications over HTTP. George needs to ensure that HTTP traffic receives higher priority than all other traffic, but he also must also make sure that inbound HTTP traffic receives higher priority than outbound HTTP traffic.

George configures a QoS Policy with the following rules:

Name	Web_Traffic	Default
Source	Any	Any
Destination	Any	Any
Service	HTTP	Any
Action	Weight 20	Weight 10
Track	None	None
Install On	Any	Any
Time	Any	Any
Comment	Blank	Blank

The total bandwidth available to the QoS Module is 8 MBps (8,192 KBps).

George creates a network object, PQR_Net, to represent the supernet that encompasses all of PQR's internal networks.

George then configures two sub-rules for the Web_Traffic Rule. (The Default Rule for the Web_Traffic sub-rule set was added automatically by Check Point QoS when the first sub-rule was created.)

Name	Inbound_Web	Outbound_Web	Default
Source	Any	PQR_Net	Any
Destination	PQR_Net	Any	Any
Service	Any	Any	Any
Action	Weight 50	Weight 20	Weight 10
Track	None	None	None
Install On	Any	Any	Any
Time	Any	Any	Any
Comment	Blank	Blank	Blank



There is no need to define the service for these sub-rules, as only HTTP traffic will be passed down from the top-level rule. The Source, Destination, and Service fields of a sub-rule must always be a subset of the parent rule, or the sub-rule will not be effective.

Connections match both the top-level rules (Web_Traffic and Default), so bandwidth is allocated between the two rules according to weight:

$$(20/30) * 8,192 = 5,461.3 \text{ KBps for the Web_Traffic Rule}$$

$$(10/30) * 8,192 = 2,730.7 \text{ KBps for the Default Rule}$$

Connections match all three sub-rules of the Web_Traffic Rule. The bandwidth allocated to the Web_Traffic Rule is further subdivided according to the weights of the sub-rules:

$$(50/80) * 5,461.3 = 3,413.3 \text{ KBps for the Inbound_Web sub-rule}$$

$$(20/80) * 5,461.3 = 1,365.3 \text{ KBps for the Outbound_Web sub-rule}$$

$$(10/80) * 5,461.3 = 682.7 \text{ KBps for the Default sub-rule}$$

Guarantees and limits within sub-rule sets are handled the same way as with top-level rules. The bandwidth is limited to that allocated to the sub-rule set's top-level rule.

ADDITIONAL QoS RULE CONSIDERATIONS

Some additional considerations for QoS rules include:

- If a guarantee is defined in a sub-rule, then a guarantee must be defined for the rule above it.
- The guarantee of a sub-rule cannot be greater than the guarantee defined for the rule above it.
- A rule guarantee must not be smaller than the sum of guarantees defined in its sub-rules.
- If a rule's weight is low, some connections may receive very little bandwidth.
- The sum of guarantees in rules in the upper level should not exceed 90 percent of the capacity of the link.
- If per-connection guarantees are defined both for a rule and its sub-rules, the sum of the per-connection guarantees for the sub-rules should not be greater than the per-connection guarantee of the top level rule.
- If both a rule limit and a per-connection limit are defined for a rule, 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 in all the sub-rules, the rule limit should not be greater than the sum of limits defined in the sub-rules.
- If a rule limit and a guarantee per rule are defined in a rule, then the limit should not be smaller than the guarantee.



If both a limit and a guarantee are defined in a rule, and the limit is equal to the guarantee, connections may receive no bandwidth. This situation can occur when the rule has sub-rules with total rule guarantees that add up to the total rule guarantee for the entire rule, and the rule also has sub-rules with no guarantee.

DIFFERENTIATED SERVICES

.....

Differentiated Services (DiffServ) is an architecture for providing different types or levels of service for network traffic. Inside the enterprise network, packets are marked in the IP header TOS byte as belonging to a certain Class of Service, or QoS Class. These packets are then granted priority on the public network. DiffServ packets have meaning on the public network, not inside the enterprise network. Effective 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 from one location to another in one of two ways:

```
:ipsec.copy_TOS_to_inner
```

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

```
:ipsec.copy_TOS_to_outer
```

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

This property should be set, per QoS Module, in `$FWDIR/conf/objects_5_0.c`.

The default setting is:

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

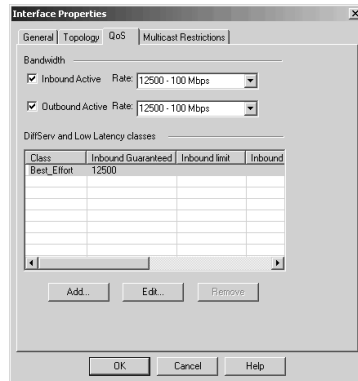
```
:ipsec.copy_TOS_to_outer (true)
```

Interaction Between DiffServ Rules and Other Rules

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 rules of this class are installed.



For example, suppose 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.



Interface QoS Properties Tab

On other interfaces, the rule is not installed and FTP connections routed through those other interfaces do not receive the weight specified in the rule. To specify a weight for all FTP connections, add a rule under “Best Effort.”

DiffServ rules can be installed only on interfaces for which the relevant QoS Class has been defined in the QoS tab of the Interface Properties screen. “Best Effort” rules (that is, non-DiffServ rules) can be installed on all interfaces of gateways with QoS Modules installed. Only rules installed on the same interface interact with each other.

LOW LATENCY QUEUING

.....

For most traffic, **Weighted Fair Queuing** (WFQ) is adequate. WFQ helps avoid drops caused by congestion. Avoiding drops can mean holding long queues. Long queues may lead to non-negligible delays.

Long queues are inappropriate for voice and video traffic. For most delay-sensitive applications, packets need not be dropped from queues to keep them short. The streams of these applications have a known, bounded bit rate. Check Point QoS can be configured to forward as much traffic as the stream delivers, ensuring only a small number of packets accumulate in the queues.

Low Latency Queuing allows a Security Administrator to define special Classes of Service for delay-sensitive applications. Rules under these classes are used with other rules in the QoS Policy Rule Base.

For each Low Latency class defined on an interface, specify a constant bit rate and maximal delay for active directions. Check Point QoS checks packets matched to Low Latency class rules, to prevent them from being delayed longer than their maximal delay. If the maximal delay of a packet is exceeded, the packet is dropped. Otherwise, packets are transmitted at the constant bit rate defined for the Low Latency class to which it belongs. If the Constant Bit Rate of the class is defined correctly (meaning that it is not smaller than the expected arrival rate of the matched traffic), packets are not dropped. When the arrival rate is higher than the specified Constant Bit Rate, packets exceeding this constant rate are dropped, to ensure that those transmitted are within the maximal delay limitations.

Low Latency Classes

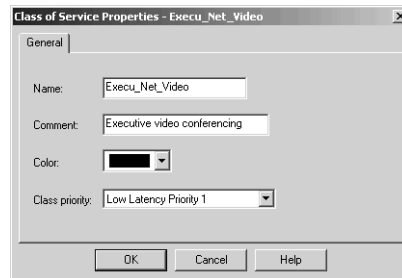
Low Latency classes specify the maximal delay that is tolerated and a Constant Bit Rate. Check Point QoS guarantees traffic matching rules of this type is forwarded within the limits of the bounded delay.



DEFINING A LOW LATENCY CLASS


To define a Low Latency Class:

1. Select Manage > QoS > QoS Classes from the top menu bar in SmartDashboard.
2. Click the New button, and select Low Latency Class of Service.
3. Give the class a name, and assign it a priority.



Low Latency Class Properties

4. Select the target QoS Module. Go to the Topology screen, and select the interface to edit.
5. Click the Add button, and select Low Latency Classes.
6. Populate the Add Low Latency QoS Class Properties screen, and click OK to exit the properties screens.



Low Latency QoS Class Properties

7. Right-click the QoS Rule Base, and select Add Class of Service > Above.

- Select the target Class of Service from the Add Class of Service screen, and click OK. The Class of Service Rule is added to the QoS Rule Base. A Best Effort Rule is also added:

NAME	SOURCE	DESTINATION	SERVICE	ACTION	TRACK	INSTALL ON	TIME	COMMENT
Execu_Net_Video								
Best_Effort								
Web_Traffic	* Any	* Any	tcp http	Weight 10	- None	* All	* Any	
Default	* Any	* Any	* Any	Weight 10	- None	* All	* Any	

Low Latency Class Rule and Best Effort Rule

- To activate the Low Latency Class, define at least one rule under it in the QoS Policy Rule Base:

NAME	SOURCE	DESTINATION	SERVICE	ACTION	TRACK	INSTALL ON	TIME	COMMENT
Execu_Net_Video								
Video_Rule	* Any	* Any	* Any	Weight 10	- None	* All	* Any	
Best_Effort								
Web_Traffic	* Any	* Any	tcp http	Weight 10	- None	* All	* Any	
Default	* Any	* Any	* Any	Weight 10	- None	* All	* Any	

Low Latency Class Rule with QoS Policy Rule

The traffic matching any Low Latency Class rule receives the delay and Constant Bit Rate properties defined for the specified class, and is also allocated bandwidth according to the rule properties (weight, limit, and guarantee).



The maximal delay 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

It is advisable to define more than one Low Latency Class, if different types of traffic require different maximal delays. Low Latency Classes are assigned one out of five priority levels. These priority levels are relative to other Low Latency Classes. The class with the lower maximal delay should 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, and must be taken into consideration when setting the minimal delay for the class. Initially set the priorities for all Low Latency Classes according to maximal delay, and then define the classes by descending priority.



For each direction of an interface (inbound and outbound), the sum of the constant bit rates of all the Low Latency Classes cannot exceed 20 percent of the total designated bandwidth rate. The 20-percent limit is set to ensure that “Best Effort” traffic does not suffer substantial delay and jitter, as a result of the Low Latency Classes.

COMPUTING CBR

To accurately set the CBR of a Low Latency Class, a Security Administrator must know:

- The bit rate of a single application stream of traffic matching the class.
- The expected number of simultaneously open streams.

The CBR of the class should be the bit rate of a single application, multiplied by the expected number of simultaneous streams:

$$(\text{Single Stream}) * (\text{Number of Streams Expected}) = \text{CBR}$$



If the number of streams exceeds the number expected when setting the CBR, the total incoming bit rate will exceed the CBR, and many drops will occur. Avoid this situation, by limiting the number of concurrent streams.

COMPUTING MAXIMAL DELAY

The delay value defined for a class determines the number of packets that can be queued in the Low Latency queue before drops begin. When computing the maximal delay of a Low Latency class, consider both the maximal delay that streams matching the class can tolerate, and the minimal delay that Check Point QoS can guarantee.

Use the following method to estimate the greatest delay the class can tolerate:

1. Refer to the technical details of the streaming application, and find the delay it can tolerate.
2. Find or estimate the delay the external network imposes.
3. Subtract the delay imposed by the external network from the maximum delay the streaming application can tolerate.

Use the following method to estimate the smallest delay for the class:

1. Find the bit rate of the streaming application.



If you cannot find the bit rate of the streaming application in the application properties, you can use Check Point SmartView Monitor to observe the application and discover the bit rate.

2. Estimate the typical packet size in the stream.



If you do not know the packet size, you can use `fw monitor` to capture traffic and record packet size, or you can use the LAN's MTU. (1500 bytes for Ethernet)

3. Estimate burst size, by monitoring the internal interfaces that precede the QoS Module.
4. If no burstiness is detected, the minimum delay of the class should be no smaller than this:

$$(3 * \textit{packet size}) / \textit{bit rate}$$

5. If burstiness is detected, the minimum delay of the class should be no smaller than this:

$$[(\textit{burst size} + 1) * \textit{packet size}] / \textit{bit rate}$$

The maximal delay for a class should be between the greatest and smallest delay. Set the delay close to the greatest value, only if the application occasionally bursts.



When you set the maximal delay, you will see an error message, if the value is lower than the value set for other Low Latency Classes of higher priority. Check Point QoS will also display an error message, if the maximal delay is incompatible with the interface speed. Both error messages include a minimum acceptable maximal-delay value. Set the maximal delay to a value no smaller than the one printed in the message.

PREVENTING UNWANTED DROPS

If the aggregate bit rate going through the Low Latency Class exceeds the CBR of the class, drops occur. This situation may occur when the number of streams actually opened exceeds the number you expected, when you set the CBR.

Limit the number of connections allowed to the number of connections used to compute the class's CBR, by modifying the QoS Policy as follows:

1. Define a single rule under the class, with a per-connection guarantee as its Action.
2. In the Per Connection Guarantee field of the QoS Action Properties screen, define the per-connection bit rate you expect.
3. In the Number of guaranteed connections field, define the maximum number of connections allowed in this class.
4. Do not check the Accept additional non-guaranteed connections option.

When to Use Low Latency Queuing

Use Low Latency Queuing when the bit rate of the stream is known and controlling delay is important. Low Latency Classes do not receive TOS markers. If preferential treatment is required beyond the QoS Module, DiffServ should be used instead.

ADVANCED FEATURES

.....

Check Point QoS includes advanced features to allow authenticated QoS, support Citrix MetaFrame, and integrate with Load Sharing configurations.

Authenticated QoS

Check Point Authenticated QoS provides Quality of Service (QoS) for end-users in dynamic IP environments, such as remote-access and DHCP environments. Authenticated QoS enables 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. Check Point QoS supports Client Authentication, including Encrypted Client Authentication, and SecuRemote/SecureClient Authentication. User and Session Authentication are not supported.

Citrix MetaFrame Support

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. One of the disadvantages of using Citrix ICA is that uncontrolled printing traffic can consume all available bandwidth, leaving mission-critical applications struggling.

Check Point QoS solves the problem by:

- Classifying 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.

LIMITATIONS

The following limitations apply when using Check Point QoS to control Citrix MetaFrame traffic:

- The Citrix TCP services are supported in Traditional mode QoS Policies only.
- Session Sharing must be disabled.
- The number of applications that are detected by the inspection infrastructure is limited to 2,048. Console errors will be sent if this limit is exceeded. These errors are harmless, and will not affect your system. Simply 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 single rule.

Load Sharing

Load Sharing is a mechanism that distributes traffic within a cluster of gateways, so that the total throughput of multiple machines is increased. Check Point QoS's architecture guarantees that Load Sharing will provide either:

- Two-way Stickiness; all packets of a single connection use the same machine in both directions.
- Conversation Stickiness; all packets of control/data connections within a conversation use the same machine in both directions.

Check Point QoS provides a fault-tolerant QoS solution for cluster Load Sharing that deploys a unique, distributed, WFQ bandwidth-management technology. The user is able to specify a unified QoS Policy per virtual interface of the cluster. The resulting bandwidth allocation is therefore identical to that obtained by installing the same Policy on a single server.



Under a load state, there are a few connections that are backlogged active for short periods of time. In such cases, the Load Sharing function in ClusterXL is not spread evenly. But in this case, there is no congestion and therefore no need for QoS.

MONITORING QOS POLICY

The QoS Policy and its impact on traffic can be monitored using SmartView Tracker, Eventia Reporter, and SmartView Monitor. Each SmartConsole tool provides a rich set of data, to evaluate the effectiveness of the QoS Policy.

SmartView Tracker

The following two conditions must be met for a match on a QoS Policy rule to be logged:

- The QoS logging check box must be selected on the Gateway Properties > Additional Logging Configuration screen.
- The connection's matching rule must be marked with either Log or Account, in the rule's track field.

QOS TRACK FIELD SET TO LOG

QoS rules with the track field set to Log can generate the following types of log events:

- **Connection Rejection**
QoS rejects a connection when the number of guaranteed connections is exceeded, and/or when the rule's action properties are not set to accept additional connections.
- **Running Out of Packet Buffers**
QoS generates an Out of Packet Buffers string when QoS global packet buffers are exhausted, or one of the interface direction's packet buffers is exhausted. This report is generated no more frequently than once every twelve hours.
- **LLQ Packet Drop**
When a packet is dropped from an LLQ connection, a report is generated. This report is generated no more frequently than once every five minutes.


QOS TRACK FIELD SET TO ACCOUNT

QoS rules set to Account generate the following types of events:

- Total bytes transmitted through QoS for each relevant interface and direction
- Total bytes dropped from the connection as a result of QoS's drop Policy
- Count of bytes dropped from the connection, because the maximum used memory fragments for a single connection was exceeded
- Number of bytes dropped from the connection, due to delay expiration for LLQ connections
- Average packet delay for LLQ connections
- Jitter (maximum delay difference between two consecutive packets) for LLQ connections

SmartView Monitor

SmartView Monitor allows a Security Administrator to view the top QoS rules by interface. Information provided in SmartView Monitor can help a Security Administrator optimize the QoS Rule Base, by showing which rules should be placed near the top:



Sample of SmartView Monitor QoS Information

Eventia Reporter

The Eventia Reporter Network Activity report is very useful to a Security Administrator attempting to define a QoS Policy. The Network Activity report presents data about traffic accepted by the Security Gateway. This report can be used to see network activity, to determine effective allocation of bandwidth.

Specific sections include information regarding:

- Overall traffic characteristics, as well as a breakdown by hour and by date
- The top network users
- Top services used
- Top sources and top destinations of network traffic

If a Security Administrator is concerned that an abundance of HTTP and RealAudio traffic is consuming excessive bandwidth and starving IMAP connections, the Top Network Activity - Services section of the Network Activity report could help determine if proper application of QoS could improve bandwidth allocation:

Top Network Activity Services					
Index	Service	Traffic Size			
		Number of Connections	% of Total Connections	MB	% of Total Bytes
1	http	4,954	50.81%	164.80	25.42%
2	pop-3	2,332	23.92%	74.70	11.63%
3	https	1,428	14.64%	47.28	7.29%
4	ftp	684	7.01%	351.82	54.28%
5	smtp	353	3.62%	9.58	1.48%
Total (5)		9,751	100.00%	648.18	100.00%
Average		1,950	20.00%	129.64	20.00%

Top Network Activity - Services Report



OPTIMIZING CHECK POINT QOS

Check Point QoS performance can be improved by following the suggestions below:

- Upgrade to the newest Check Point QoS version available.
- Install Check Point QoS only on the external interfaces of the QoS Module. Unless you are using limits for inbound traffic, installing Check Point QoS only in the outbound direction will provide you with most of the functionality and improvements.
- Put more frequent rules at the top of your Rule Base. You can use SmartView Monitor to analyze how much a rule is used.
- Turn per-connection limits into per-rule limits.
- Turn per-connection guarantees into per-rule guarantees.

- *Optimizing Check Point QoS*
-
-
-
-