

Cisco Catalyst 3750 Metro Series Switch Enhanced Services Ports Deliver Advanced QoS Features for Metro Ethernet Service Providers

Overview

Today's enterprise customers are demanding a higher level quality of service (QoS) from service providers. The ability to apply QoS criteria to network traffic is essential for service provider networks. As enterprises and service providers continue the shift to a single converged network carrying mission-critical applications over a unified multiservice (data, voice, and video) architecture, the ability to manage traffic flows and delivery terms becomes increasingly critical.

Enterprises that outsource their network services to service providers will turn to providers who can supply QoS features to ensure that on-demand, bandwidth-intensive applications and time-sensitive, information-delivery applications receive the bandwidth necessary to maintain performance. The ability to apply QoS to customer traffic will become vital to meeting service-level agreements (SLAs) and maintaining network performance. Support for advanced QoS features will also allow service providers to create premium products and service bundles.

Cisco® Catalyst® 3750 Metro Series switches meet the demand for a high level of QoS. The series' QoS technology facilitates the management of network performance with respect to bandwidth, delay, jitter, and

packet loss, which are critical to optimizing application performance and meeting SLAs. The switch supports the Cisco Hierarchical Queuing Framework (HQF), which includes traffic classification; two-rate, three-color policing; Class-Based Weighted Fair Queuing (CBWFQ); Low-Latency Queueing (LLQ); and traffic shaping on the output of the Enhanced Services ports.

The hierarchical implementation of this framework enables service providers to manage their QoS at three different levels—at the physical interface level (Enhanced Services ports), the VLAN level (VLANs within each port), and marked class level (classes with within each VLAN), providing the most granular and flexible architecture to fulfill SLAs. In addition, congestion-avoidance techniques such as Weighted Random Early Detection (WRED) and tail drop are supported to better regulate network traffic and prevent congestion.

Product Overview

The Cisco Catalyst 3750 Metro Series is a new premier multilayer switch that brings greater intelligence to the metro Ethernet edge, enabling the delivery of more differentiated metro Ethernet services. The switch features a hierarchical QoS framework and traffic shaping, intelligent 802.1Q tunneling, VLAN mapping, Multiprotocol Label Switching (MPLS)



and Ethernet over MPLS (EoMPLS) support, and redundant AC or DC power. The switch is ideal for service providers seeking to deliver profitable business services, such as Layer 2, Layer 3, and MPLS VPNs, in several bandwidths and with different SLAs. Moreover, with flexible software options, the Cisco Catalyst 3750 Metro Series offers service providers a cost-effective path for meeting current and future service requirements from enterprises and commercial businesses.

Cisco Catalyst 3750 Metro Series switches are built with two types of ports—customer ports and Enhanced Services ports. Customer ports connect to end customers and have the same features as Cisco Catalyst 3750 Series switches. The main QoS features available on customer ports include single-rate ingress policer, egress queue shapers, and weighted tail drop.

Enhanced services ports connect to the network. These ports have all the advanced QoS features on hardware and software that service provider customers demand. For example, in today's Ethernet access switches, the egress port usually has only four or eight queues. This is not a problem if each switch is only servicing a single customer, but when the second customer is added to the switch, service providers will have to mix all customer traffic together and provide similar QoS profiles. This limits the different QoS profiles demanded by each customer.

The solution is to increase the number of queues supported on the network port. Cisco Catalyst 3750 Metro Series switches use Cisco HQF technology. The Catalyst 3750 Metro Series hardware supports 4000 queues on each of its Enhanced Services ports. Numerous egress queues allow service providers to configure a set of queues for each customer. This removes the constraint of sharing multiple customers' traffic in the same queues and provides a custom QoS profile for each customer.

This white paper discusses Hierarchical Queuing Framework features only on egress of enhanced services ports.

Cisco HQF Overview

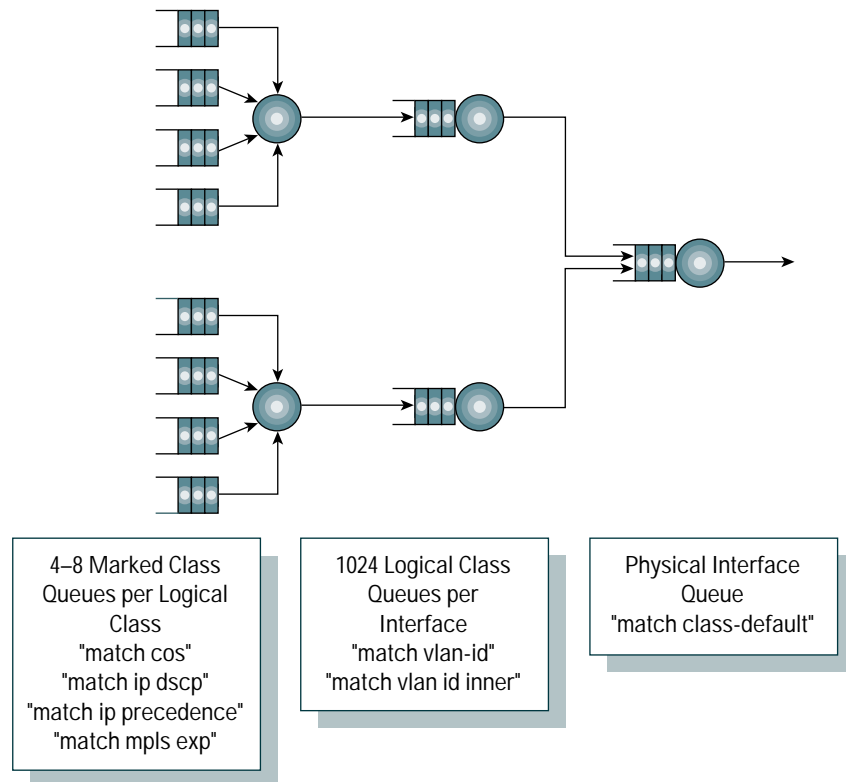
Cisco HQF provides hierarchical queuing management. Unlike traditional multilayer switches, which support a flat level of queues, the Cisco Catalyst 3750 Metro Series switch's enhanced services ports provide three levels of packet queues—physical, logical, and class level. The physical level of hierarchy is used for policing or shaping the physical interface. It can only be identified by class default.

The logical level of hierarchy is classified by VLAN ID. At this level, the classification can match on VLAN tag (802.1Q or Inter-Switch Link [ISL]) or 802.1Q tunnel tags (Q-in-Q).

Traffic at the class level is classified by class of service (CoS), IP Precedence, IP differentiated services code point (DSCP), and MPLS EXP bits. The three levels of queues include 4096 class level queues, 1024 logical level queues, and 1 physical level queue available on each Enhanced Services port. For each logical level queue, there can be 1 to 8 class level queues. These queues can also be serviced by multiple QoS functions such as policing, shaping, and sharing.



Figure 1
Hierarchical Queuing Framework Model



Cisco Catalyst 3750 Metro Series switches perform the following functions within Cisco HQF:

- Queue scheduling
 - CBWFQ
 - Strict priority
- Traffic shaping
- Policing—Two-rate, three-color policer
- WRED

Queue Scheduling—CBWFQ and LLQ

Cisco Catalyst 3750 Metro Series switches support two types of queue scheduling—CBWFQ and LLQ. CBWFQ is used for bandwidth allocation. It allows users to allocate minimum bandwidth to each class of traffic. In addition, it allows queues to share bandwidth when one class is not using all its allocated bandwidth.

LLQ requires the configured queue to have strict priority. As long as there is traffic in this queue, it will always be serviced first. Due to the behavior, LLQ has very low latency, which is good for voice or video traffic.



CBWFQ

CBWFQ extends standard WFQ to provide support for user-defined traffic classes. For CBWFQ, users define traffic classes based on match criteria, including CoS, IP Precedence, IP DSCP, and MPLS EXP bits. Packets satisfying the match criteria for a class constitute the traffic for that class. A queue is reserved for each class, and traffic belonging to a class is directed to the queue for that class. After a class has been defined according to its match criteria, users can assign it bandwidth to guarantee a minimum amount of bandwidth that is available to it during times of congestion.

LLQ

LLQ brings strict priority queuing to CBWFQ. Strict priority queuing allows delay-sensitive data such as voice to be dequeued and sent first (before packets in other queues are dequeued), giving delay-sensitive data preferential treatment over other traffic. This feature reduces latency and jitter in configured queues.

Traffic Shaping

Traffic shaping buffers the data frames from a particular traffic flow. The data frames are then sent into the network in regulated amounts to ensure that the traffic fits within the promised traffic envelope for the particular connection. Traffic shaping decreases the “burstiness” of User Datagram Protocol (UDP) and TCP traffic, decreasing the load on network elements. The traffic shaping function on the Cisco Catalyst 3750 Metro Series allows for the configuration of the average rate at which the traffic is sent out of each queue. In addition to providing a smoother traffic flow, traffic shaping can also be used to control the maximum amount of bandwidth allowed for that queue.

Policing

Cisco Catalyst 3750 Metro Series switches support a two-rate, three color policer, providing advanced bandwidth management through rate limiting. The Cisco Catalyst 3750 Metro Series' policing function allows users to control the maximum rate of traffic received on each Enhanced Services port queue. This rate-limiting function is based on user-defined criteria; it marks packets by setting the CoS value, IP Precedence value, IP DSCP value, and MPLS experimental value.

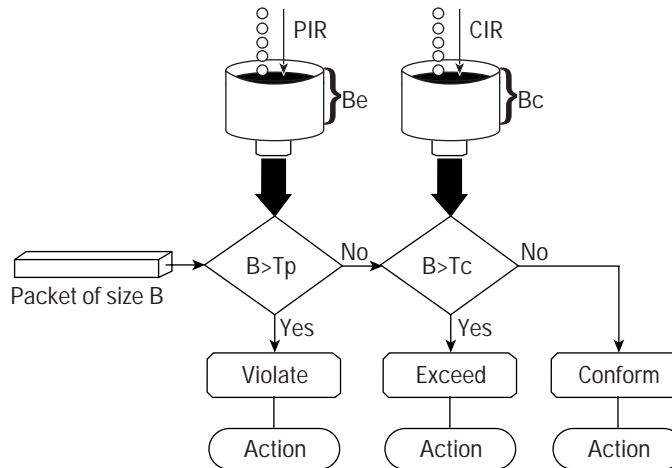
With the two-rate policer, users can enforce traffic policing according to two separate rates—Committed Information Rate (CIR) and Peak Information Rate (PIR). The two-rate policer manages the maximum rate of traffic through the dual-token bucket algorithm, which uses user-configured values to determine the maximum rate of traffic allowed on a queue at a given moment.

The dual-token bucket algorithm provides users with three actions for each packet—a conform action, an exceed action, and an optional violate action. Traffic entering the queue with the two-rate policer configured is placed into one of these categories. Within these three categories, users can decide packet treatments. For instance, packets that conform can be configured to be sent, packets that exceed can be configured to be sent with a decreased priority, and packets that violate can be configured to be dropped.

In addition to rate limiting, the two-rate policer allows users to independently mark a packet according to whether the packet conforms, exceeds, or violates a specified rate. Networking devices within the network can then use this setting to determine how the traffic should be treated. For example, the WRED feature uses the IP Precedence value to determine the probability that a packet will be dropped.



Figure 2
Dual-token bucket Algorithm



WRED

RED is a congestion avoidance mechanism that takes advantage of TCP's congestion control mechanism. By randomly dropping packets prior to periods of high congestion, RED tells the packet source to decrease its transmission rate. Assuming the packet source is using TCP, it will decrease its transmission rate until all the packets reach their destination, indicating that the congestion is cleared.

WRED on Cisco Catalyst 3750 Metro Series provides the ability to configure RED features selectively, based on IP Precedence or IP DSCP. Based on these two traffic markings, the user has the options to configure minimum threshold, maximum threshold, and mark probability.

When the average queue depth is above the minimum threshold, RED starts dropping packets. The rate of packet drop increases linearly as the average queue size increases, until the average queue size reaches the maximum threshold. When the average queue size is above the maximum threshold, all packets are dropped.

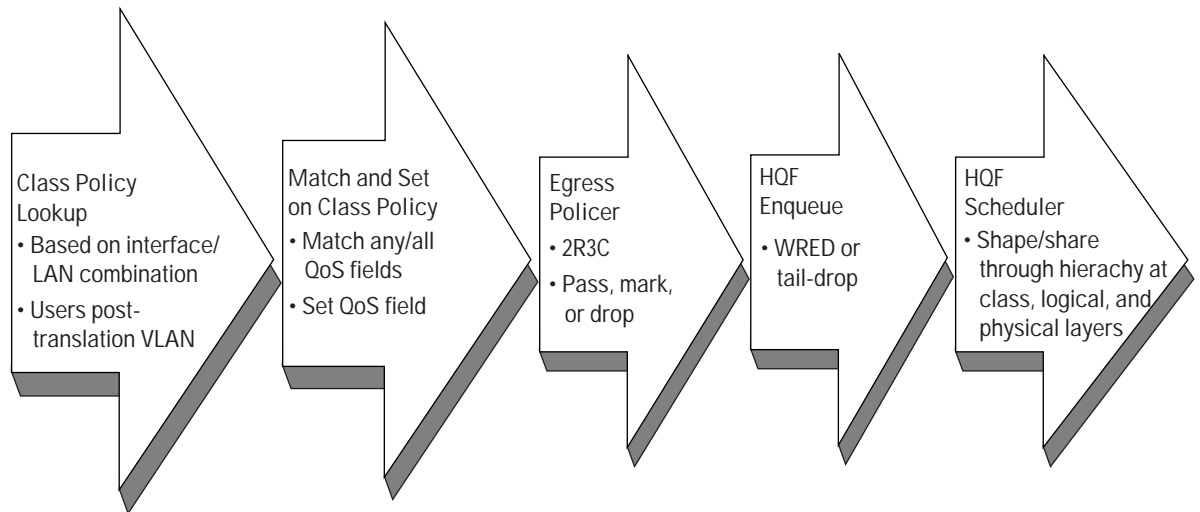
The mark probability denominator is the fraction of packets dropped when the average queue depth is at the maximum threshold. For example, if the denominator is 512, one out of every 512 packets is dropped when the average queue is at the maximum threshold.

Order of Operation

On Cisco Catalyst 3750 Metro Series Enhanced Services ports, all QoS functions can be enabled at the same time. This provides the flexibility to configure many types of QoS profiles. Depending on the order of the functions, the resulting profiles can be very different—it is important to understand the order of processing as the frames enter and leave the queues. The order of the main QoS functions is first—policing, then CBWFQ, then shaping. Figure 3 briefly describes the QoS stages that an Enhanced Services port goes through for each frame it receives from customer ports.



Figure 3
Network Processor processing stages



Cisco HQF Configuration

Cisco HQF is configured on the Cisco Catalyst 3750 Metro Series using Cisco Modular QoS CLI (MQC). Cisco MQC standardizes the command-line interface (CLI) and semantics for configuring QoS features. It provides standardized QoS provisioning across all Cisco IOS® Software platforms, and hides the underlying hardware architecture. Customers can thus upgrade to the next-generation Cisco platform and copy existing Cisco MQC QoS configurations, regardless of the deployment platforms. Cisco MQC also supports the concept of hierarchical policies to deliver an extremely powerful QoS solution.

Cisco MQC allows the configuration of QoS functions in the Cisco Catalyst 3750 Metro Series for traffic classification, modification, policing, queuing, and scheduling, which can be used at various levels in the hierarchy. The available commands under Cisco MQC are as follows:

- Bandwidth—Specify a minimum bandwidth allocation to a traffic class
- Priority—Specify that a traffic class requires low latency
- Shape—Specify a maximum rate for a traffic class
- Police—Specify a policer for the traffic class
- Queue-Limit—Tune the limit on queue depth for a traffic class
- Random-Detect—Enable WRED drop policy on a traffic class
- Set—Set CoS, DSCP, IP Precedence, and MPLS EXP bits in a traffic class
- Match—Match any, VLAN, inner VLAN, CoS, DSCP, IP Precedence, and MPLS EXP bits to classify traffic
- Service-Policy—Attach a service policy to either the physical interface or logical level



Not all these commands can be configured at all hierarchy levels. On the physical level, the commands available are Police and Shape. On the logical level, the commands available include Bandwidth, Shape, Police, and Match (any, VLAN, and inner VLAN). On the class level, the commands available are Bandwidth, Priority, Shape, Police, Queue-Limit, Random-Detect, Match (any, CoS, DSCP, IP Precedence, MPLS EXP), and Set (CoS, DSCP, IP Precedence, MPLS EXP). Modular QoS CLI configuration includes the following steps:

- Step 1. Define a traffic class with the class-map command
- Step 2. Create a traffic policy by associating the traffic class with one or more QoS features (using the policy-map command)
- Step 3. Attach the traffic policy to the interface with the service-policy command

The class-map command is used to define a traffic class, which classifies traffic. A traffic class contains three major elements—a name, a series of match commands, and, if more than one match command exists in the traffic class, an instruction on how these match commands should be evaluated. The traffic class is named in the class-map command line; for example, if you enter the class-map cisco command while configuring the traffic class in the CLI, the traffic class would be named cisco.

The match commands are used to specify various criteria for classifying packets. Packets are checked to determine whether they match the criteria specified in the match commands; if a packet matches the specified criteria, that packet is considered a member of the class and is forwarded according to the QoS specifications set in the traffic policy. Packets that fail to meet any of the matching criteria are classified as members of the default traffic class.

```
C3750M(config)# class-map cisco
C3750M(config-cmap)# match ip precedence 1
```

The policy-map command is used to create a traffic policy, which configures the QoS features that should be associated with the traffic that has been classified in a user-specified traffic class or classes. A traffic policy contains three elements—a name, a traffic class (specified with the class command), and the QoS policies.

The name of a traffic policy is specified in the policy-map CLI (for example, issuing the policy-map class1 command would create a traffic policy named class1). The traffic class that is used to classify traffic to the specified traffic policy is defined in policy map configuration mode, which is the automatic mode after naming the traffic policy. After choosing the traffic class that is used to classify traffic to the traffic policy, the user can enter the QoS features to apply to the classified traffic. This is done in policy-map class configuration mode.

```
C3750M(config)# policy-map shape
C3750M(config-pmap)# class cisco
C3750M(config-pmap)# shape average 64000
```

Cisco MQC does not require that users associate only one traffic class to one traffic policy. When packets match to more than one match criterion, multiple traffic classes can be associated with a single traffic policy. The service-policy command is used to attach the traffic policy, as specified with the policy-map command, to an interface or a class to create hierarchy.

```
C3750M(config)# policy-map parent
C3750M(config-pmap)# class cisco
C3750M(config-pmap)# service-policy child
```



On Cisco Catalyst 3750 Metro Series switches, the Cisco HQF can only be attached to the egress of ES ports. The service-policy output class1 command would attach all the characteristics of the traffic policy named class1 to the specified interface. All packets leaving the specified interface are evaluated according to the criteria specified in the traffic policy named class1.

```
Router(config)# interface g1/1/1
Router(config-if)# service-policy output class1
```

Cisco HQF Configuration Guidelines

Bandwidth

The following rules and methods are used while allocating bandwidth for the logical and class level classes.

- The shaper at the physical level restricts the bandwidth visible to the logical level classes; the shaper rate configured at the physical interface becomes the total bandwidth available to all the logical level classes.
- At all logical level classes, the shaper rate also becomes the visible bandwidth if the user for that class does not specify the bandwidth.
- At all logical and class level classes, the shaper rate must be greater than or equal to the bandwidth allocated at that level.
- The total bandwidth at the logical level classes, including the default class, cannot exceed the visible bandwidth at the physical level.
- The total bandwidth at the class level classes, including the default class, cannot exceed the visible bandwidth at the parent logical level.
- At the logical and class levels, the user-configured bandwidth is summed for all siblings. If the summed-up bandwidth is less than the parent visible bandwidth, remaining bandwidth is distributed equally among all the siblings.

Priority

- Priority queuing can only be configured at the class level
- For each logical level queue, one class priority queue can be configured
- Priority queuing on the Cisco Catalyst 3750 Metro Series is strict priority; no bandwidth or percent command is allowed

Cisco HQF Application Examples

Ethernet to the Business—TLS with Differentiated Services

Metro Ethernet Provider A is providing Transparent LAN Service (TLS) to two customers in a building. Both customers are connected to a Cisco Catalyst 3750 Metro Series switch in the basement. The provider wants to limit the uplink to 500 MB but to also provide both customers with their own QoS profiles.

Customer A

- SP ID = 10
- TLS bandwidth = 150 MB shaped, must be accessible by all classes



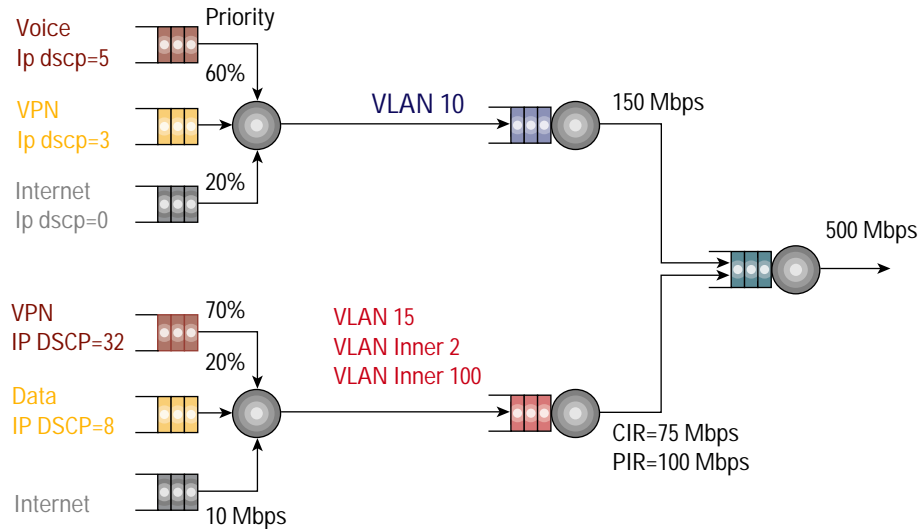
- Three levels of class:
 - Priority class: For voice, identified by IP Precedence = 5
 - VPN class: Regular data traffic, identified by IP Precedence = 3; 60 percent of bandwidth
 - Internet class: Internet traffic, identified by IP Precedence = 0; 20 percent of bandwidth

Customer B

- SP ID = 15
- TLS bandwidth = CIR: 75 MB, PIR: 100 MB; must be accessible by all classes
- Two VLANs:
 - VLAN 2
 - VPN class: IP DSCP = 32 in case of congestion; 70 percent of available bandwidth should be guaranteed for this class
 - Data class: IP DSCP = 8; 20 percent of available bandwidth should be guaranteed for this class
 - VLAN 100
 - Internet class: Should never exceed 10 MB



Figure 4



Configuration

```

class-map match-all custAvlan10                                --- identify Customer A
  match vlan 10
class-map match-all custAipprec5                             -- Customer As priority traffic
  match ip precedence 5
class-map match-all custAipprec0                             -- Customer As Internet traffic
  match ip precedence 0
class-map match-all custAipprec3                             -- Customer As VPN traffic
  match ip precedence 3
class-map match-all custBvlan15                               -- identify Customer B
  match vlan 15
  match vlan inner 2
  match vlan inner 100
class-map match-all custBipdscp32                            --- Customer Bs VPN traffic
  match ip dscp 32
class-map match-all custBipdscp8                             -- Customer Bs data traffic
  match ip dscp 8
policy-map custAclasspolicy                                   -- policy for Customer As class queues
  class custAipprec5
    priority                                                  -- assign strict priority to this queue
  class custAipprec3
  class custAipprec0
  bandwidth percent 60
  
```



```
random-detect
random-detect precedence 3 200 400 3000
class custAipprec0
bandwidth percent 20
policy-map custBclasspolicy -- policy for Customer B's class queues
class custBipdsdp32
bandwidth percent 70
class custBipdsdp8
bandwidth percent 20
class class-default
shape average 10000000
policy-map vlanpolicy -- logical level policy attaching class level policies
class custAvlan10
shape average 150000000
service-policy custAclasspolicy
class custBvlan15
police cir 75000000 bc 8192 pir 100000000 be 8192
service-policy custBclasspolicy
policy-map ES1 -- physical level policy attaching logical level policy
class class-default
shape average 500000000
service-policy vlanpolicy
interfaces gi1/1/1 --- attach the HQF policy to the interface
service-policy output ES1
```

Scalability Numbers

Table 1 Scalability Numbers

Queues	8000
Traffic shapers	8000
Policers	2000
Class-map	4000
Policy-map	2000
WRED	1000

Supported MIB

The Class-Based QoS MIB provides read access to QoS configurations. This MIB also provides QoS statistics information based on Cisco MQC, including information regarding class map and policy map parameters.

Summary

Cisco® Catalyst® 3750 Metro Series switches with the Hierarchical Queuing Framework features on the Enhanced Service ports is able to deliver traffic shaping, 2 rates 3 colors policers, and other advance QoS features on its 8000 queues. With the set of rich features and flexibility, the Service Providers will be able to offers differentiated services to multiple customers serviced by a Catalyst 3750 Metro Series switch. This will help Service Provider increase the revenues and receive higher profit margin from the Metro Ethernet Services.



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

European Headquarters
Cisco Systems International BV
Haarlerbergpark
Haarlerbergweg 13-19
1101 CH Amsterdam
The Netherlands
www-europe.cisco.com
Tel: 31 0 20 357 1000
Fax: 31 0 20 357 1100

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-7660
Fax: 408 527-0883

Asia Pacific Headquarters
Cisco Systems, Inc.
Capital Tower
168 Robinson Road
#22-01 to #29-01
Singapore 068912
www.cisco.com
Tel: +65 6317 7777
Fax: +65 6317 7799

Cisco Systems has more than 200 offices in the following countries and regions. Addresses, phone numbers, and fax numbers are listed on the Cisco Web site at www.cisco.com/go/offices

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China PRC • Colombia • Costa Rica • Croatia
Czech Republic • Denmark • Dubai, UAE • Finland • France • Germany • Greece • Hong Kong SAR • Hungary • India • Indonesia • Ireland
Israel • Italy • Japan • Korea • Luxembourg • Malaysia • Mexico • The Netherlands • New Zealand • Norway • Peru • Philippines • Poland
Portugal • Puerto Rico • Romania • Russia • Saudi Arabia • Scotland • Singapore • Slovakia • Slovenia • South Africa • Spain • Sweden
Switzerland • Taiwan • Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela • Vietnam • Zimbabwe

All contents are Copyright © 1992–2004 Cisco Systems, Inc. All rights reserved. Catalyst, Cisco, Cisco Systems, the Cisco Systems logo, and Cisco IOS are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and certain other countries.

All other trademarks mentioned in this document or Web site are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company.
(0304R) ETMG 203144—JR 01.04