



Network Architecture and Service Model

A service provider network architecture contains access routers, distributed routers and core routers or ATM switches. Access routers terminate customer connections at the edge of the network.

IP QoS provisioning with the Cisco IP Solution Center (ISC) is configured on the access circuit that involves the access router (provider edge devices, or PEs) in the service provider network and the customer equipment (CEs) in the customer network.

Ethernet QoS provisioning with ISC supports a subset of the features required for the Metro Ethernet 3.1 Solution. With Ethernet QoS, ISC can deploy QoS Policies on Cisco Catalyst switches in a provider's network.

This appendix includes the following sections:

- [IP QoS, page A-1](#)
- [Ethernet QoS, page A-12](#)

IP QoS

This section describes network architecture and service model for IP QoS in ISC.

It includes the following sections:

The section contains the following subsections:

- [IP QoS Service Provider Network Architecture, page A-1](#)
- [IP QoS Service Model Overview, page A-2](#)
- [Service Model Components, page A-3](#)
- [QoS Link Definition, page A-4](#)
- [Service Level IP QoS Policy, page A-4](#)
- [Link Level IP QoS Policy, page A-7](#)
- [IP QoS Service Requests, page A-8](#)
- [IP QoS Provisioning Strategies, page A-9](#)

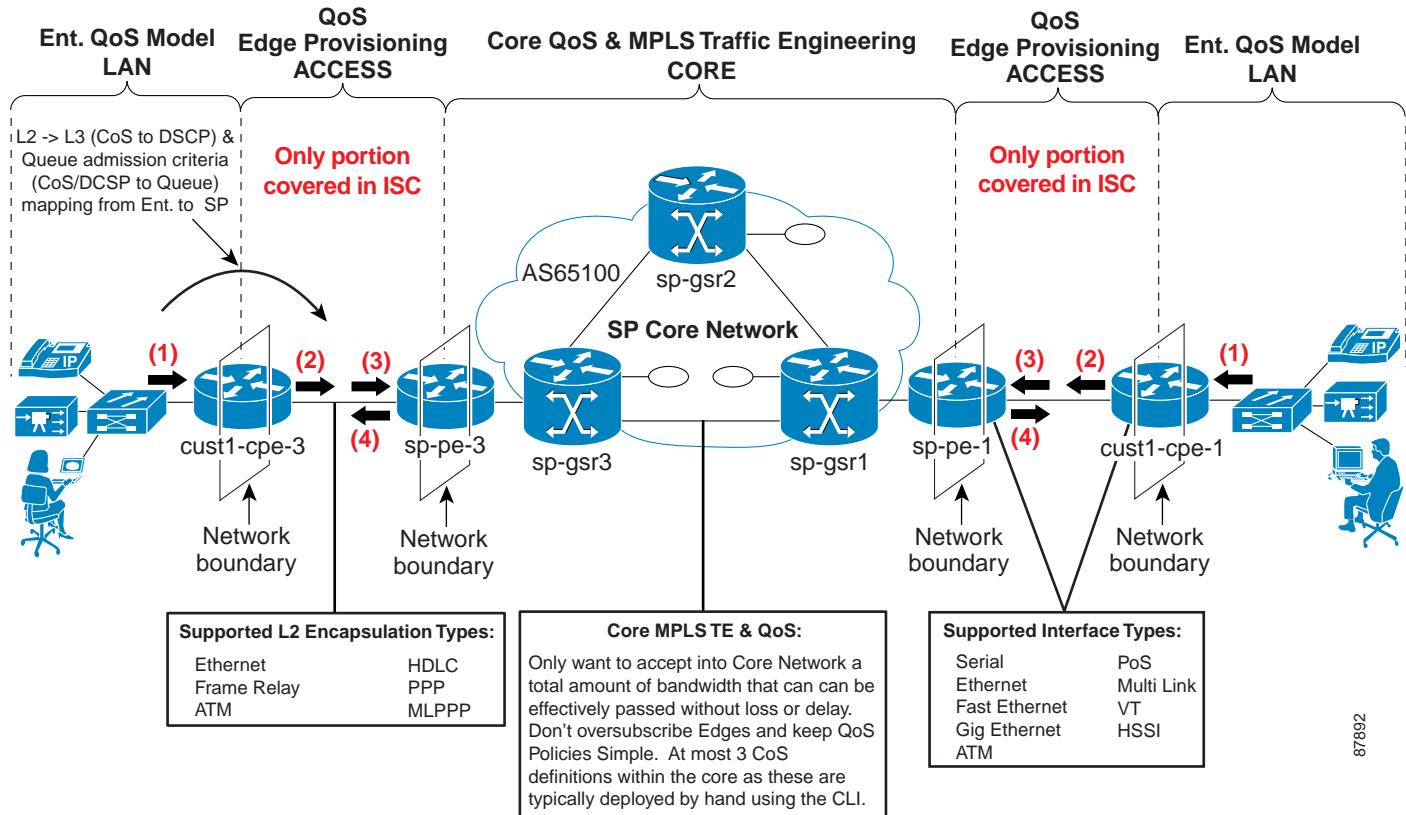
IP QoS Service Provider Network Architecture

The service provider network architecture, supported within the scope of the IP QoS provisioning model in ISC is:

- Access circuit (CEs)
- Distribution routers (PEs)
- Core (routers and ATM switches)

These QoS components and concepts are represented in [Figure A-1](#).

Figure A-1 QoS Components and Concepts



IP QoS Service Model Overview

The IP QoS service model in ISC is designed so that QoS provisioning is implemented for traffic that enters the access circuit at the network edge (CE), and through the distribution portion (the CE-PE link).

This section provides an overview of the IP QoS service model in ISC.

The ISC implementation of IP QoS provisioning involves traffic through several different types of devices, link speeds, and encapsulation types. For this reason, an ISC QoS policy is divided into two categories:

- Service level IP QoS policy—The service level policy corresponds to IP QoS service classes. QoS service classes provide a method for classifying traffic.

Typically, a service provider creates three or four service classes for each QoS policy. For example, a service provider might have a platinum, gold, silver, and bronze QoS policies, and each of these policies might contain 3 service classes; a VoIP, a management, and a data service class (Best-Effort or Business-Data-1).

- Link level QoS policy—The link level QoS policy is a grouping of QoS parameters that are sensitive to the CE-PE link's bandwidth and layer 2 encapsulation. Link level parameters include Frame Relay traffic shaper, ATM shaper, FRF.12, LFI over MLPPP, cRTP, and interface-based rate limiting.

Typically, a service provider creates several link level QoS settings. For example, a service provider might create a link QoS setting for different bandwidths and encapsulation types, such as; FR_64K_gold, FR_64K_silver, FR_128K_bronze, ATM_1MBPS_gold, and Ethernet_2mbps_Silver.

ISC provides two levels of QoS policies because a QoS service request might contain one or more links with different circuit bandwidths and encapsulation types. The service level policy is designed for a type of service, like voice, but can apply to more than one link type. The link level policy is designed for different link speeds, like 1 Mbps, and can apply QoS provisioning per link.

To provision QoS parameters for devices in a service request, a network operator must:

- Select the appropriate service level QoS policy
- Associate a corresponding link level policy with each link in the service request.

For example, if the QoS service request is comprised of two links; a Frame Relay link with a bandwidth of 64kbps, and an ATM link, with a bandwidth of 1 mbps, and the service level agreement (SLA) purchased by the enterprise customer is the Gold policy, the following settings might be associated with the QoS service request:

- Gold service level QoS policy
- FR_64K_gold link level QoS settings tied to the Frame Relay link
- ATM_1MBPS_gold link level QoS settings tied to the ATM link

QoS policies can either be customer-oriented or provider-oriented. Typically, service provider networks have a combination of both service level and link level QoS policies.

Service Model Components

A QoS policy is a set of parameters that control and condition the traffic flowing through a service provider network.

The Cisco IP Solution Center (ISC) configures QoS at the access circuit, which involves the PEs in the service provider network and the CEs in the customer network. A QoS policy is applied to the selected set of access circuits using a QoS service request. The ISC provisioning engine generates the QoS configuration from the service request and downloads the configuration to the specified CE and PE devices.

A QoS service request can be integrated with VPN provisioning accomplished through ISC or deployed on its own if VPN services are not provisioned through ISC.

In ISC, the IP QoS service model is comprised of:

- IP QoS Link—contains device (CE and PE) interface information.
- IP QoS Policy—ISC offers two different levels of policies. Most networks have a combination of both policy types.
 - Service level QoS policy—The part of the QoS policy where you define service classes for the different service level agreements (SLA) purchased by customers.
 - Link level QoS policy—The part of the QoS policy that specifies QoS parameters that are specific to the CE-PE link.
- IP QoS Service Request—a container for the link objects and QoS policies to be applied to the device.



Note

To set up QoS provisioning for MPLS VPN services, see [IP QoS for MPLS VPNs, page 3-28](#).

QoS Link Definition

A QoS Link can contain two interfaces (for both the CE and PE) or one interface (CE only or PE only). For QoS provisioning, you can select both interfaces in the CE-PE link. A typical device interface selection is as follows:

- For the CE device:
 - The provider-facing device interface is selected as the link endpoint.
 - The customer-facing LAN interface is selected for marking and rate limiting.



Note

Marking and rate limiting on the customer-facing LAN interface is optional and is done using a CE device editor as part of defining QoS link candidates.

- For the PE device:
 - The customer-facing interface is marked as a link endpoint.

The interfaces selected as link endpoints can be provisioned with QoS parameters such as policing, traffic shaping, congestion management, congestion avoidance, link efficiency, and CAR. You apply these parameters later in the provisioning process.

See [Creating QoS Link Candidate Objects, page 3-2](#) for information on defining QoS link candidate interfaces in the ISC user interface.

Service Level IP QoS Policy

The service level portion of the QoS policy corresponds to service classes. A QoS service class provides a method for classifying traffic flows into classes so that you can apply the appropriate QoS parameters to a class of traffic instead of applying them to all traffic. For example, all TCP traffic might be grouped into a single class so that bandwidth is allocated for the class and not for individual traffic flows.

A QoS service class can include:

- Methods for classifying traffic (protocol, DSCP value, IP precedence value, source address)
- Methods for marking traffic (DSCP or IP precedence values)

- Traffic shaping parameters (average/peak, rate)
- Rate limiting parameters (mean/peak rate, burst size, conform/exceed/violate actions)
- Congestion management parameters (bandwidth and queue limit)
- Congestion avoidance parameters (drop, exponential weighing constant)

A typical service provider network might create different QoS policies, and each QoS policy might contain three to five service classes. For example, a service provider might have a gold, silver, and bronze QoS policies, each specifying different service level agreements (SLA), and each of those QoS policies might contain one or more service classes. Most networks require at least a voice, a management, and a data service class.

ISC provides five default or template service classes for you to modify and use for a service level QoS policy:

- VoIP—voice service class
- Routing Protocol—routing protocol service class
- Management—management service class
- Business-Data-1—data service class
- Best Effort—data service class

See [Creating the Service Level IP QoS Policy, page 3-9](#) for information on defining the service level QoS policy in the ISC user interface.

The following section describes the five service classes provided with ISC.

QoS Service Classes

A QoS service class defines how each QoS parameter is applied.

Network traffic can be categorized into voice traffic, data traffic, and control traffic. Voice and data traffic are common in enterprise networks. Control traffic refers to routing protocol traffic and management traffic, which are commonly used in the service provider portion of the network.

The five default service classes provided with ISC cover most networks, which require at least one for interactive voice traffic, one for management traffic, and at least one service class for data traffic.

You can either remove or add more service classes if required. ISC supports the number of service classes defined by the Cisco differentiated services (DiffServ) architecture; up to 64 classes for DSCP traffic, and up to 8 service classes for IP Precedence traffic.

See [Adding a Data Service Class, page B-20](#) for more information.

VoIP Service Class

Interactive voice traffic in ISC refers to any voice traffic (telephone calls, faxes) that is IP-encapsulated and sent over the network, such as Voice-over-IP (VoIP).

Mandatory QoS components for this service class:

- Traffic classification
- Marking
- Congestion management
- Rate-limiting (optional)

Routing Protocol Service Class

Routing protocol traffic refers to traffic control messages, such as route update messages, hellos, database descriptors, keepalives, and database refresh messages. We recommended the minimum bandwidth, one percent, for your routing protocol service class.

Mandatory QoS components for this service class:

- Traffic classification
- Congestion management

Management Service Class

Management traffic refers to the traffic between the management station at the provider core and the access routers. We recommended the minimum bandwidth, one percent, for your management service class.

Mandatory QoS components for this service class:

- Traffic classification
- Marking
- Congestion management

QoS parameters for the VoIP, Routing Protocol, and Management service classes are described in [VoIP, Routing Protocol, and Management Service Classes, page B-2](#).

Business-Data-1 and Best Effort Service Classes

The two data service classes, Business-Data-1 and Best Effort, are nearly identical. The only difference between them is the Traffic Classification parameter. For Business-Data-1, traffic is classified by protocol. Best-Effort classifies all traffic.

The QoS requirements for data applications can vary. Each data application should be profiled before you determine the appropriate classification and scheduling treatment.

Mandatory QoS components for this service class:

- Traffic classification
- Marking
- Congestion management

Optional components:

- Traffic shaping or rate limiting
- Congestion avoidance



Note A typical network requires traffic shaping or rate limiting, but not both.

QoS parameters for the Business-Data-1 and Best-Effort service classes are described in [Business Data and Best Effort Service Classes, page B-10](#).

Link Level IP QoS Policy

The link level portion of the QoS policy corresponds to QoS parameters that are sensitive to link bandwidth and the CE-PE link's encapsulation type. A link level QoS policy, called link QoS settings in the ISC user interface, provides a method for defining policies specific to the CE-PE link. For example, you might require different policies for Frame Relay and ATM links because of the different encapsulation involved.

Link level QoS parameters in ISC include:

- Link bandwidth (bandwidth specified in kbps)
- Aggregated traffic shaper types (Frame Relay traffic shapers, ATM traffic shapers, and parent level traffic shapers for nested policies)



Note Aggregated traffic shapers are different from class-based traffic shapers. Aggregated traffic shapers apply to traffic through a particular CE-PE link. Class-based traffic shapers apply to all traffic specified in the service class.

- Link efficiency settings (FRF.12, LFI on MLPPP, and cRTP)
- Interface-based aggregated rate limiters (traffic classification, direction, mean rate, burst size, and conform/exceed action)



Note Interface-based aggregated rate limiters are different from class-based rate limiters. Interface-based aggregated rate limiters apply to traffic through a particular CE-PE link. Class-based rate limiters apply to all traffic specified in the service class.

Aggregated Traffic Shapers

Aggregated traffic shaping allows you to control the traffic leaving an interface. You can select an aggregated traffic shaper for each CE-PE link.

Aggregated traffic shapers are optional. ISC supports the following aggregated traffic shapers:

- Frame Relay traffic shaper, or FRTS
- FRTS (non-MQC Based)
- Parent level Class-based Shaper
- ATM traffic shaper (VBR-rt)
- ATM traffic shaper (VBR-nrt)
- ATM traffic shaper (CBR)
- ATM traffic shaper (ABR)

See [Aggregated Traffic Shapers, page B-21](#) for more information on defining the aggregated traffic shapers parameters in the ISC user interface.

Link Efficiency

Link efficiency settings are based on the bandwidth of the CE-PE link itself and are used to minimize serialization delay on the link. ISC uses methods of fragmentation and compression to minimize this delay.

ISC supports the following link efficiency settings:

- LFI on Frame Relay (FRF.12)—Supports the transport of real-time voice and data traffic on Frame Relay virtual circuits (VCs) without causing excessive delay to the real-time traffic.
- LFI on MLPPP—Multilink PPP (MLPPP) provides a method of splitting, recombining, and sequencing datagrams across multiple logical data links. MLPPP allows packets to be fragmented and the fragments to be sent at the same time over multiple point-to-point links to the same remote address.
- cRTP header compression—cRTP compresses the IP/UDP/RTP header in an RTP data packet from 40 bytes to approximately 2 to 5 bytes. Use cRTP on any WAN interface where bandwidth is at a premium and much of the traffic is RTP traffic.

See [Link Efficiency Settings, page B-30](#) for information on defining the link efficiency parameters in the ISC user interface.

Interface-Based Aggregated Rate Limiters

Interface-based aggregated rate limiters allow you to control the maximum rate of traffic sent or received on an interface for the CE-PE link. You can also specify traffic handling policies for when the traffic conforms or exceeds the specified rate limit.

Aggregate rate limits match all packets or a subset of packets on an interface or subinterface. To specify class-based rate limiting parameters, see [Creating the Service Level IP QoS Policy, page 3-9](#).

ISC supports the following interface-based rate limiter parameters:

- Traffic classification
- Direction
- Mean rate
- Burst sizes (conformed and extended)
- Conform and exceed actions

See [Interface-Based Aggregated Rate Limiters, page B-31](#) for information on defining the interface-based rate limiters in the ISC user interface.

IP QoS Service Requests

An IP QoS service request contains one or more QoS links. Each link can optionally be associated with a link QoS setting. A QoS policy can be associated with a QoS service request.

An IP QoS service request should:

- Contain an IP QoS policy
- All links in the service request can be associated with a link QoS setting

To apply IP QoS policies to network devices, you must deploy the QoS service request. When you deploy a QoS service request, ISC compares the device information in the Repository (the ISC database) with the current device configuration and generates a configlet.

See [Creating an IP QoS Service Request, page 3-20](#) for information on creating the QoS service request using the ISC user interface.

See [Cisco IP Solution Center Infrastructure Reference, 4.0](#) for more information on the ISC Repository.

IP QoS Provisioning Strategies

ISC configures IP QoS at the access circuit, which involves the PE devices in the service provider network and the CE devices in the customer network. A QoS policy is applied to the selected set of access circuits using a QoS service request.

Typically, the points of congestion in the access circuit are:

- The provider-facing interface on the CE, with traffic flowing from the CE to the PE (egress traffic).
- The customer-facing interfaces on the PE, with traffic flowing from the PE to the CE (egress traffic).

This section describes a QoS provisioning strategy: where the congestion points in the network might be, where to apply QoS parameters, and which QoS provisioning components to use.

Managed CE Scenario

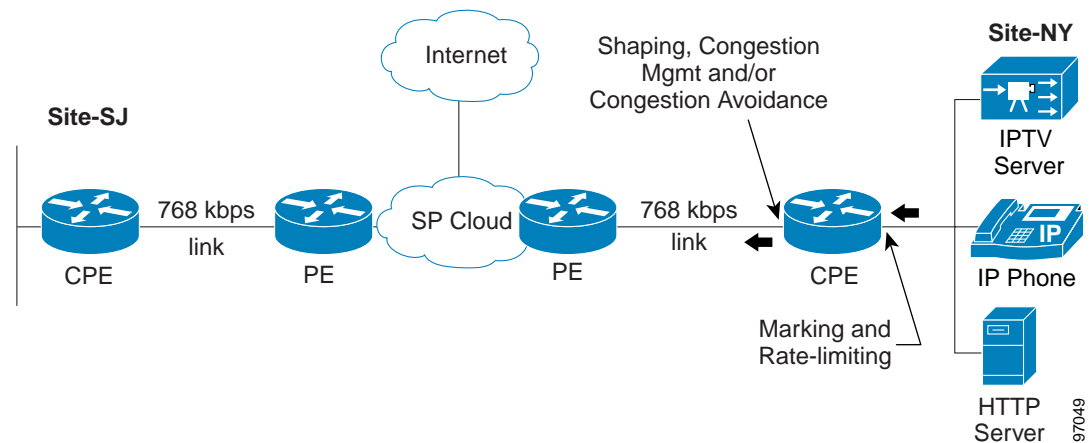
A managed CE scenario occurs when the CE is owned and managed by the service provider. In this network scenario, you can either apply QoS provisioning for the CE only or for both the CE and PE.

This section describes QoS provisioning strategies for both CE only and CE-PE scenarios.

Managed CE Only

Figure A-2 illustrates a network where QoS provisioning is configured only for the managed CE device.

Figure A-2 Managed CE Scenario



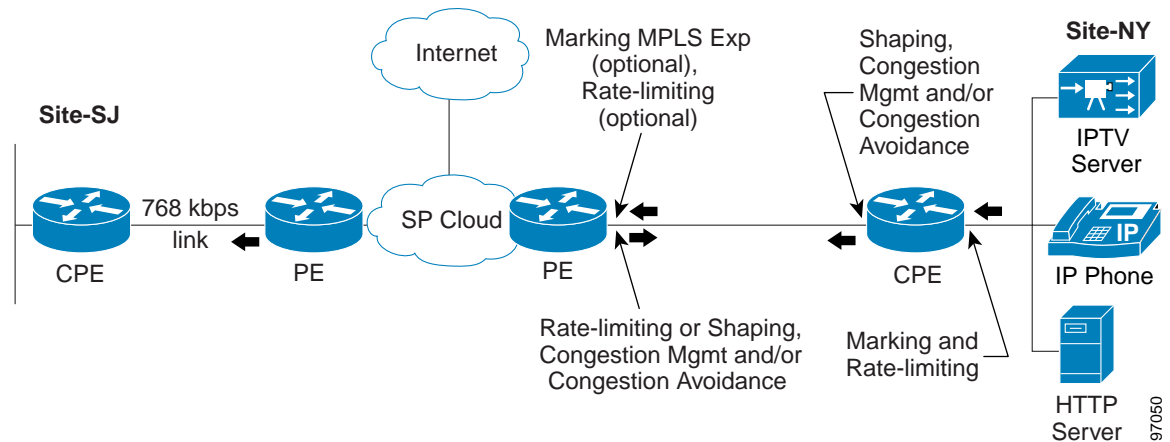
In this QoS provisioning scenario:

- Optionally configure marking and rate limiting at the customer-facing CE interface so that packets can be marked before they are encapsulated by the IPSec, GRE, and L2F and L2TP tunnels. - **IPsec and GRE are not supported in this release.** -
- Configure traffic shaping, congestion management, and congestion avoidance at the provider-facing CE interface and subinterfaces.

Managed CE and PE

Figure A-3 illustrates a network where QoS provisioning is configured for both the managed CE and the PE device.

Figure A-3 *Managed CE and PE Scenario*



In this QoS provisioning scenario:

- For traffic flowing from the CE to the PE, marking and rate-limiting are configured at the customer-facing CE interfaces, while traffic shaping, congestion management, and congestion avoidance are configured at the provider-facing CE interfaces and subinterfaces.
- For traffic flowing from the PE to the CE, QoS configuration is applied at customer-facing interfaces and subinterface for the PE. The configuration at the PE interfaces might be symmetrical to what is configured at provider-facing interfaces of a CE, but it is not required.
- If you are provisioning QoS for an MPLS VPN and you enable MPLS marking in the service level QoS policy, marking with MPLS experimental values is configured at the customer-facing PE interfaces and subinterfaces for traffic flowing from CE to PE.

Unmanaged CE Scenario

An unmanaged CE scenario occurs when the CE is not owned by the service provider, but ISC is aware of the device configuration and interface information. This information must be provided by the owner of the CE device.

In this QoS provisioning scenario:

- You must first create the CE device in ISC so that the device configuration and interface information can be stored in the ISC repository. A configlet is generated for the unmanaged CE, however, the configlet is not downloaded to the unmanaged CE device. A configuration audit is not performed for the unmanaged CE device.

See [Cisco IP Solution Center Infrastructure Reference, 4.0](#) for more information on manually creating CE devices.

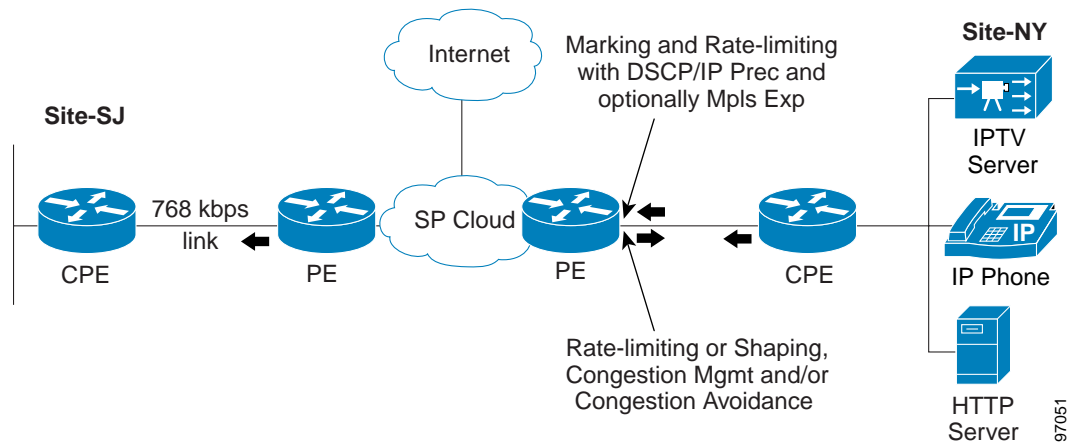
- An untrusted CE is either not managed by the service provider or is only partially-managed by the service provider. We recommend that you re-mark and re-rate limit at the provider ingress interface for the untrusted CE device. Configure the re-marking and re-rate limiting parameters in the service level policy.

See [Creating the Service Level IP QoS Policy, page 3-9](#) for more information.

PE Only Scenario

For a PE only scenario, the service provider's enterprise customer is responsible for applying the QoS configuration at the CE interfaces.

Figure A-4 PE Only Scenario



In this QoS provisioning scenario:

- For traffic flowing from the CE device to the PE device, configure marking and rate-limiting at customer-facing interfaces of the PE device.
- For traffic flowing from the PE device to the CE device, configure traffic shaping, rate-limiting, congestion management, and congestion avoidance at the same customer-facing interfaces and subinterfaces of the PE device.

See [Chapter 3, "Provisioning Process for IP QoS,"](#) for more information on the QoS provisioning process.

Ethernet QoS

This section describes network architecture and service model for Ethernet QoS in ISC.

It includes the following sections:

- [Ethernet QoS Service Model Overview, page A-13](#)
- [QoS Link Definition, page A-13](#)
- [Service Model Components, page A-14](#)
- [Terminology, page A-14](#)
- [Service Level Ethernet QoS Policy, page A-15](#)
- [Link Level Ethernet QoS Policy, page A-16](#)
- [Ethernet QoS Service Requests, page A-16](#)

Service Provider Network Architecture

In the Ethernet QoS implementation of ISC, the 3750-ME and 7600 devices can fulfill different roles as shown in [Figure A-5](#).

Figure A-5 *ISC Ethernet Network Diagram*



The traffic flow is from the CE switches towards the MPLS/IP core. In general, QoS will be applied to the UNI interfaces for the ingress direction.

An exception hereto is UNI-NNI traffic on the 3750 ME switch. In this case, the policies would be situated on the enhanced ES ports on the egress direction.

Ethernet QoS Service Model Overview

Provisioning with Ethernet QoS involves traffic through several different types of devices, link speeds, and encapsulation types. For this reason, an ISC QoS policy is divided into two categories:

- Service level Ethernet QoS policy
- Link level Ethernet QoS policy

The Ethernet QoS service model in ISC is implemented for traffic that enters the User Network Interface(UNI) and either goes out to the Network-to-Network Interface(NNI) or is switched locally within the U-PE.

Ethernet QoS policies correspond to Ethernet QoS service classes. Each QoS service class provides a method to classify, mark, condition, queue, and schedule traffic based on specified criteria, such as Class of Service, VLAN ID, etc.

A typical service provider network might create different QoS policies, and each QoS policy might contain three to four service classes. For example, a service provider might have gold, silver, and bronze QoS policies, each specifying different service level agreements (SLA), and each of those QoS policies might contain one or more service classes. Most networks require at least a voice and a data service class.

To provision Ethernet QoS parameters for devices in a service request, a network operator must:

- Create an Ethernet QoS Policy as described in [Creating an Ethernet QoS Policy, page 4-3](#).
- Create a link QoS policy, if needed.
- Create a QoS service request.
- Select a customer.
- Select a service request for L2VPN, VPLS, or MPLS (the service request must already exist).
- Select a QoS Policy created for Ethernet QoS.
- Select a link QoS policy, if needed.
- Save the service request.
- Deploy the service request.

Ethernet QoS policies can either be customer-owned or provider-owned.

For more information on the Ethernet QoS service model, see [Chapter 4, “Provisioning Process for Ethernet QoS.”](#)

QoS Link Definition

An Ethernet QoS Link can contain two interfaces (for both the U-PE and N-PE) or one interface (U-PE only or N-PE only). For Ethernet QoS provisioning, you can select both interfaces. A typical device interface selection is as follows:

For the U-PE device:

- The provider-facing device interface is selected as the link endpoint.
- The customer-facing LAN interface is selected for marking and rate limiting.

Service Model Components

In ISC, the Ethernet QoS service model is comprised of:

- Ethernet QoS Link—Contains device interface information.
- Ethernet QoS Policy—ISC offers two different levels of policies. Most networks have a combination of both policy types.
 - Service level QoS policy—The part of the QoS policy where you define service classes for the different service level agreements (SLA) purchased by customers.
 - Link level QoS policy—The part of the QoS policy that defines link-specific QoS parameters.
- Ethernet QoS Service Request—A container for the link objects and QoS policies to be applied to the device.

Figure A-6 shows a high-level view of an Ethernet network, sometimes referred to as a Metro Ethernet network, where ISC is used to apply Ethernet QoS.

Figure A-6 *ISC Ethernet Network Diagram*



Terminology

To understand the Ethernet QoS service model used in ISC, one needs to be familiar with the following terminology.

Devices

The following device types are used in Ethernet QoS for ISC:

- U-PE—User Facing Provider Edge within the Access layer.
- PE-AGG—Provider Edge Aggregation within the Aggregation layer.
- N-PE—Network Facing Provider Edge within the Edge layer.
- P—Provider Core within the Core layer.

Interfaces

The following interface types are used in Ethernet QoS for ISC:

- User Network Interface (UNI)—The physical demarcation point between the responsibility of the Service Provider and the responsibility of the Subscriber.
- Network to Network Interface (NNI)—The physical network port which connect between the Service Provider devices.
- Ethernet Network to Network Interface (E-NNI)—An Ethernet NNI.
- MPLS Network to Network Interface (MPLS-NNI)—A MPLS NNI.

Service Definitions

Following the MEF (Metro Ethernet Forum) approach, the services that comprise the Metro Ethernet 3.1 solution can be classified into two general categories:

- Point-to-Point (PTP)—Services of this connection type consist of a single point-to-point Ethernet circuit provisioned between two User Network Interfaces (UNIs).
- Multipoint-to-Multipoint (MPtMP)—Services of this connection type consist of single multipoint-to-multipoint Ethernet circuit provisioned between two or more UNIs. When there are only two UNIs in the circuit, more UNIs can be added to the same EVC if required, which distinguishes this from the point-to-point type.

In the MEF terminology this maps to the following Ethernet service types:

- Ethernet Line Service Type (E-Line)—A point-to-point Ethernet service
- Ethernet LAN Service Type (E-LAN)—A multipoint-to-multipoint Ethernet service

Within these two service types, Metro Ethernet services can be created by assigning values to a set of attributes grouped according to:

- User Network Interface (UNI)—The physical demarcation point between the responsibility of the Service Provider and the responsibility of the Subscriber.
- Ethernet Virtual Connection (EVC)—An association of two or more UNIs that limits the exchange of Service Frames to UNIs within the EVC.

Service Level Ethernet QoS Policy

The service level portion of the QoS policy corresponds to service classes. A QoS service class provides a method for classifying traffic flows into classes so that you can apply the appropriate QoS parameters to a class of traffic instead of applying them to all traffic. For example, all TCP traffic might be grouped into a single class so that bandwidth is allocated for the class and not for individual traffic flows.

A QoS service class can include:

- Methods for classifying traffic (protocol, DSCP value, IP precedence value, source address)
- Methods for marking traffic (DSCP or IP precedence values)
- Rate limiting parameters (mean/peak rate, burst size, conform/exceed/violate actions)
- Congestion management parameters (bandwidth and queue limit)

A typical service provider network might create different QoS policies, and each QoS policy might contain three to five service classes. For example, a service provider might have a gold, silver, and bronze QoS policies, each specifying different service level agreements (SLA), and each of those QoS policies might contain one or more service classes.

You can create real time, business critical, and best effort classes, etc., within a policy.

See [Creating Ethernet QoS Policies, page 4-3](#), for information on defining the service level Ethernet QoS policy in the ISC user interface.

Link Level Ethernet QoS Policy

The link level portion of the QoS policy corresponds to QoS parameters that are sensitive to link bandwidth and the UNI's encapsulation type. A link level QoS policy, called link QoS settings in the ISC user interface, provides a method for defining link-specific policies. For example, you might require different policies for Frame Relay and ATM links because of the different encapsulation involved.

Link level QoS parameters in ISC include:

- VLAN bandwidth (in kbps or %) (3750-ME only)
- VLAN shaper (3750-ME only)
- Port trust
- CoS mutation (7600 only)

Ethernet QoS Service Requests

An Ethernet QoS service request contains one or more QoS links. Each link can optionally be associated with a link QoS setting. A QoS policy can be associated with a QoS service request.

An Ethernet QoS service request should:

- Contain an Ethernet QoS policy
- All links in the service request can be associated with a link QoS setting

To apply Ethernet QoS policies to network devices, you must deploy the QoS service request. When you deploy a QoS service request, ISC compares the device information in the Repository (the ISC database) with the current device configuration and generates a configlet.

See [Creating an Ethernet QoS Service Request, page 4-13](#) for information on creating the QoS service request using the ISC user interface.

See [Cisco IP Solution Center Infrastructure Reference, 4.1](#) for more information on the ISC Repository.