QoS

Implementing Cisco Quality of Service (QOS) v2.0

Student Guide

Version 2.0

Course Introduction

Overview

Implementing Cisco Quality of Service (QOS) v2.0 provides students with in-depth knowledge of IP quality of service (QoS) requirements, conceptual models such as Best Effort, Integrated Services (IntServ) and Differentiated Services (DiffServ), and the implementation of IP QoS on Cisco IOS platforms.

The curriculum covers the theory of IP QoS, design issues, and configuration of various QoS mechanisms to facilitate the creation of effective administrative policies providing QoS. Case studies and lab exercises included in the course help students apply the concepts that are mastered in individual modules to real-life scenarios.

The course also gives students design and usage rules for various advanced IP QoS features and the integration of IP QoS with underlying Layer 2 QoS mechanisms, allowing them to design and implement efficient, optimal, and trouble-free multiservice networks.

Course Goal and Objectives

This section describes the course goal and objectives.

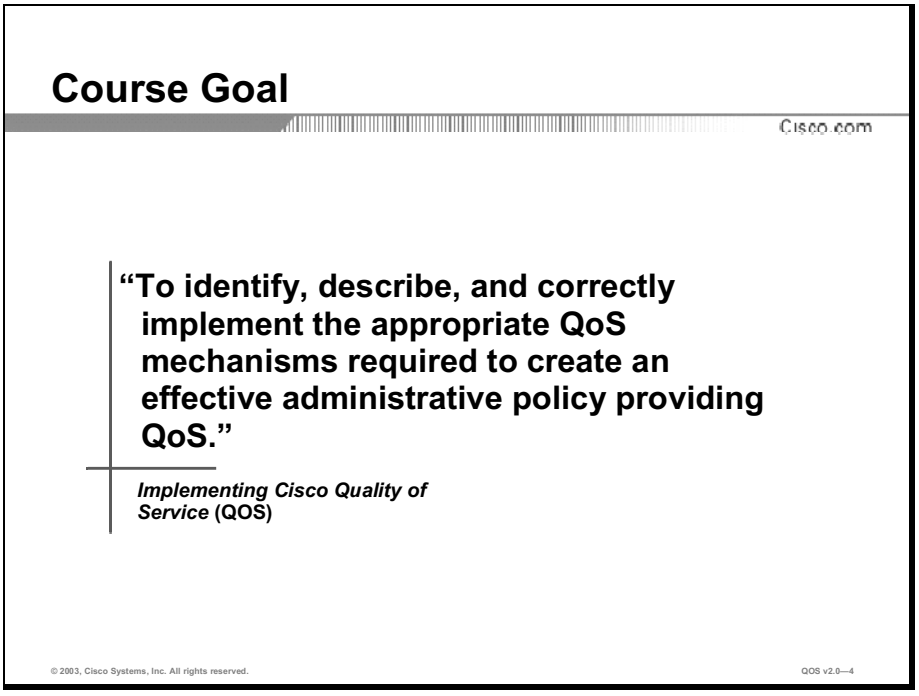

Upon completing this course, you will be able to meet these objectives:

- Explain the need to implement QoS and explain methods for implementing and managing QoS
- Identify and describe different models used for ensuring QoS in a network and explain key IP QoS mechanisms used to implement the models
- Explain the use of MQC and AutoQoS to implement QoS on the network
- Classify and mark network traffic to implement a policy defining QoS requirements
- Use Cisco QoS queuing mechanisms to manage network congestion
- Use Cisco QoS congestion avoidance mechanisms to reduce the effects of congestion on the network
- Use Cisco QoS traffic policing and traffic shaping mechanisms to effectively limit the rate of network traffic
- Use Cisco link efficiency mechanisms to improve the bandwidth efficiency of low-speed WAN links
- Correctly select the most appropriate QoS mechanisms for providing QoS using Cisco "best practices" in service provider and enterprise networks

Course Outline

The outline lists the modules included in this course.

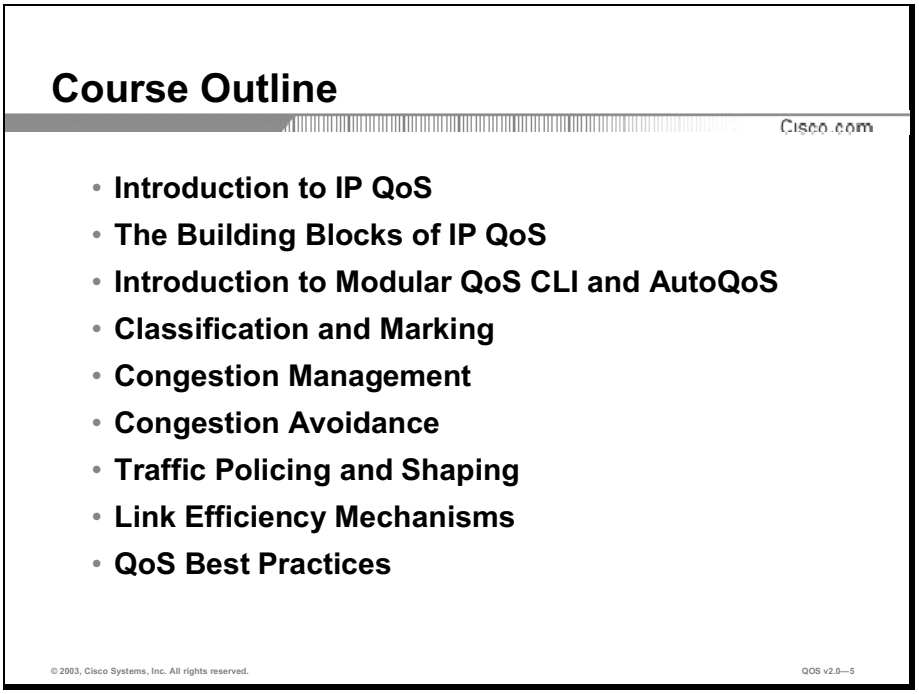

Cisco Certifications

This topic discusses Cisco career certifications and paths.

Cisco provides three levels of general career certifications for IT professionals with several different tracks to meet individual needs. Cisco also provides focused Cisco Qualified Specialist (CQS) certifications for designated areas such as cable communications, voice, and security.

There are many paths to Cisco certification, but only one requirement—passing one or more exams demonstrating knowledge and skill. For details, go to http://www.cisco.com/go/certifications.

Learner Skills and Knowledge

This topic lists the course prerequisites.

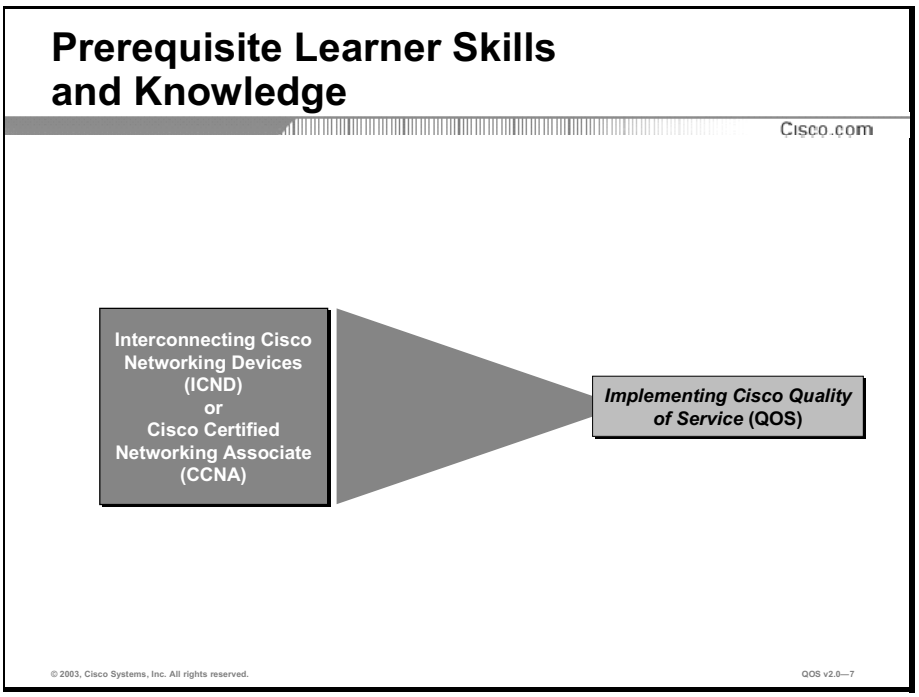

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Completion of the *Interconnecting Cisco Networking Devices* (ICND) course or Cisco CCNA® certification.

Learner Responsibilities

This topic discusses the responsibilities of the learners.

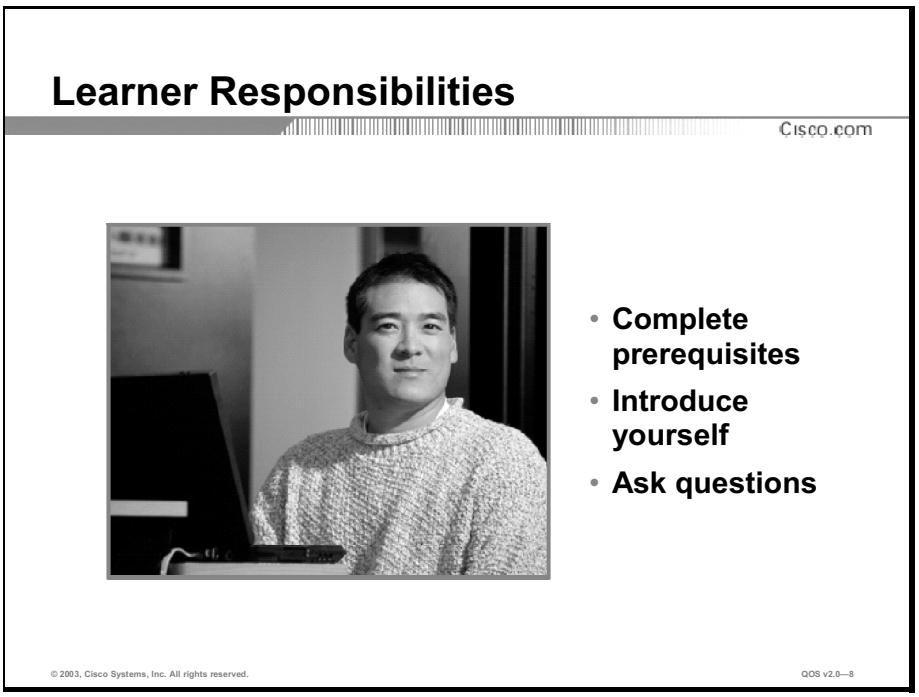

To take full advantage of the information presented in this course, you must have completed the prerequisite requirements.

In class, you are expected to participate in all lesson exercises and assessments.

In addition, you are encouraged to ask any questions relevant to the course materials.

If you have pertinent information or questions concerning future Cisco product releases and product features, please discuss these topics during breaks or after class. The instructor will answer your questions or direct you to an appropriate information source.

General Administration

This topic lists the administrative issues for the course.

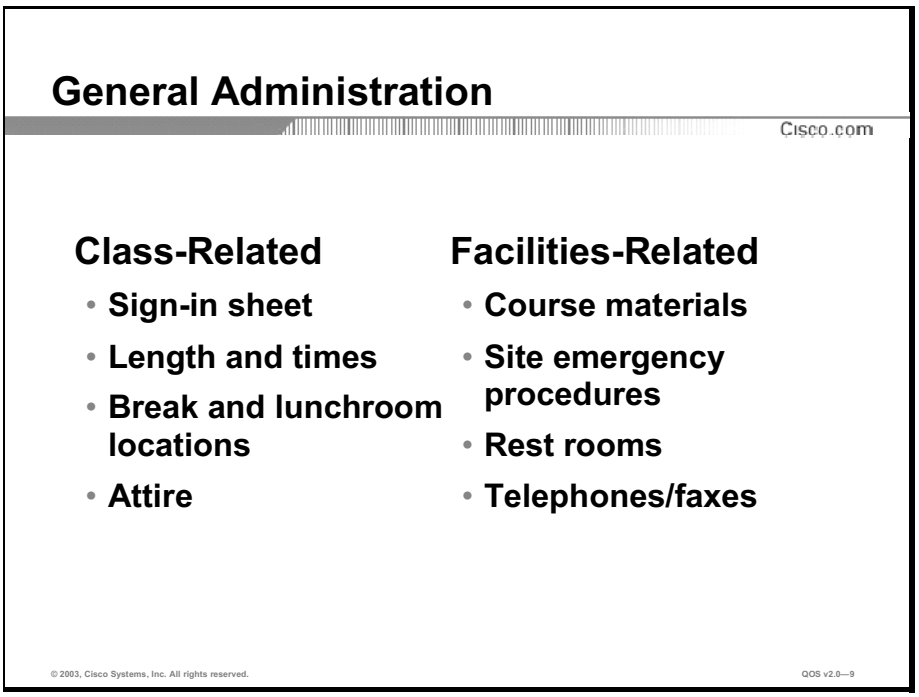

The instructor will discuss these administrative issues:

- Sign-in process
- Starting and anticipated ending times of each class day
- Class breaks and lunch facilities
- Appropriate attire during class
- Materials that you can expect to receive during class
- What to do in the event of an emergency
- Location of the rest rooms
- How to send and receive telephone and fax messages

Course Flow

This topic covers the suggested flow of the course materials.

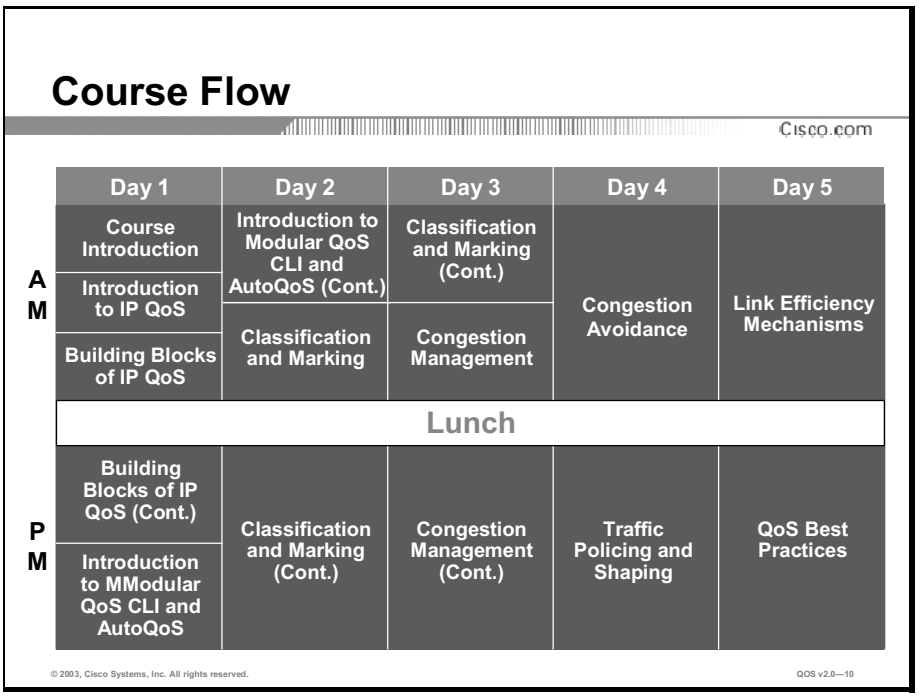

The schedule reflects the recommended structure for this course. This structure allows enough time for the instructor to present the course information and for you to work through the laboratory exercises. The exact timing of the subject materials and labs depends on the pace of your specific class.

Icons and Symbols

This topic shows the Cisco icons and symbols used in this course.

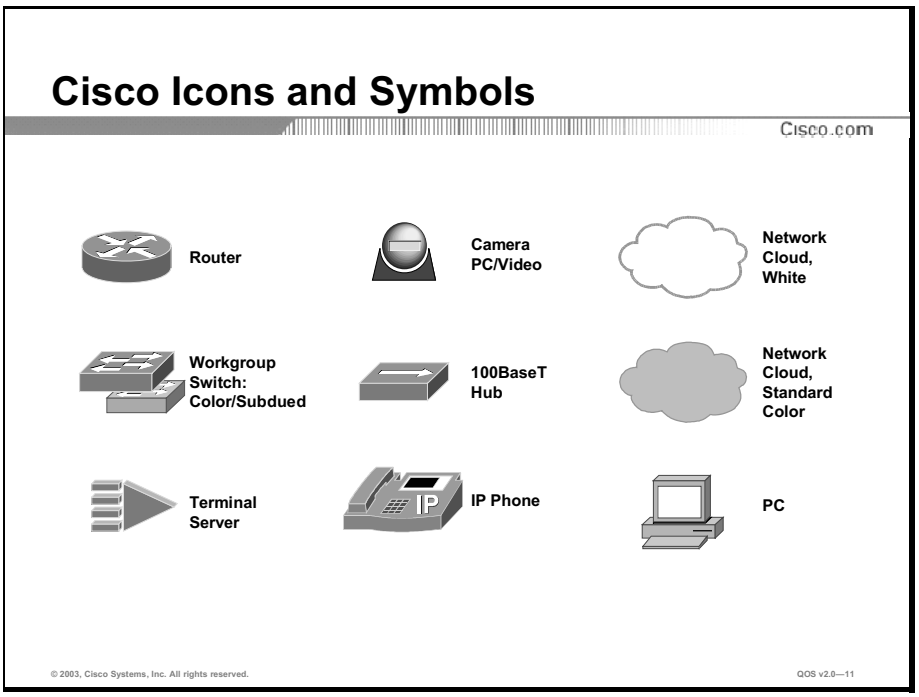

Learner Introductions

This is the point in the course where you introduce yourself.

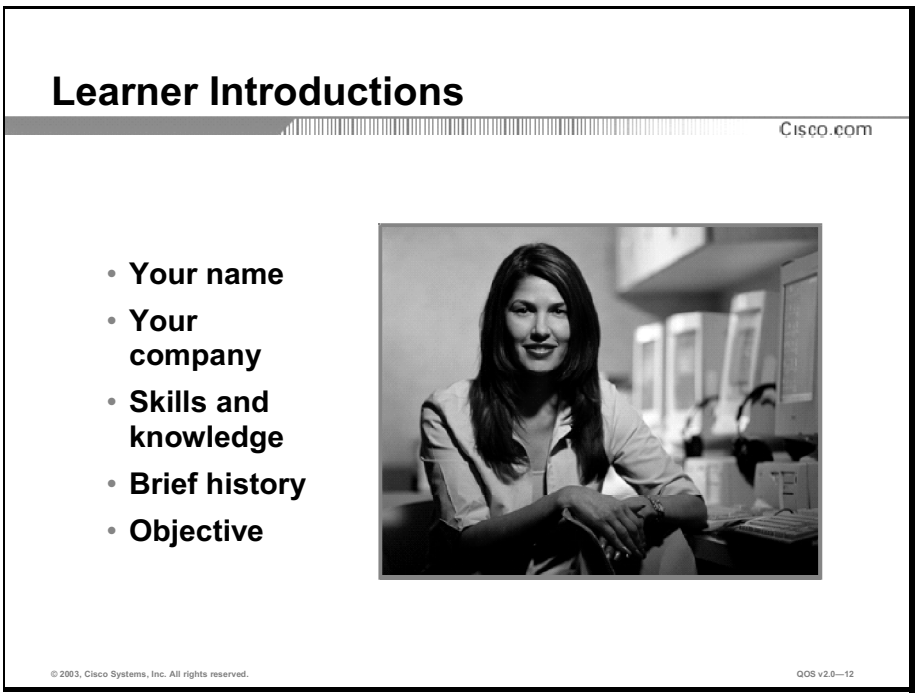

Prepare to share the following information:

- Your name
- Your company
- If you have most or all of the prerequisite skills
- A profile of your experience
- What you would like to learn from this course

Course Evaluations

Cisco relies on customer feedback to make improvements and guide business decisions. Your valuable input will help shape future Cisco learning products and program offerings.

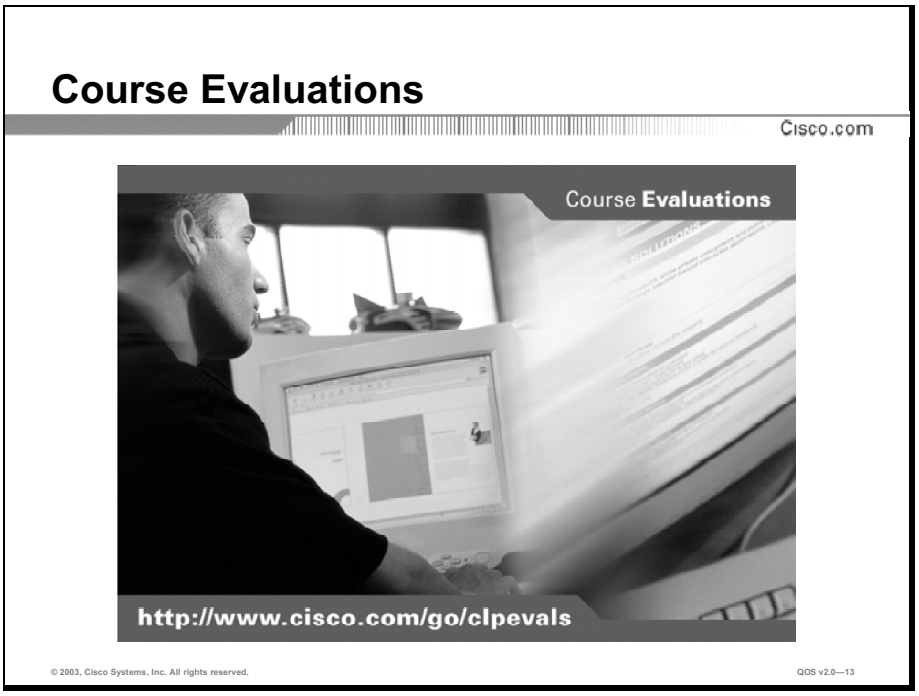

On the first and final days of class, your instructor will provide the following information needed to fill out the evaluation:

- Course acronym *(printed on student kit side label)* _____
- Course version number *(printed on student kit side label)* _____
- Cisco Learning Partner ID # _____
- Instructor ID # _____
- Course ID # *(for courses registered in Cisco Learning Locator)* _____

Please use this information to complete a brief (approximately 10 minutes) online evaluation concerning your instructor and the course materials in the student kit. To access the evaluation, go to http://www.cisco.com/go/clpevals.

After the completed survey has been submitted, you will be able to access links to a variety of Cisco resources, including information on the Cisco Career Certification programs and future Cisco Networkers events.

If you encounter any difficulties accessing the course evaluation URL or submitting your evaluation, please contact Cisco via email at clpevals_support@external.cisco.com.

Module 1

Introduction to IP QoS

Overview

As user applications continue to drive network growth and evolution, demands to support different types of traffic is also increasing. Different types of applications with differing network requirements create a need for administrative policies mandating how individual applications are to be treated by the network. Network traffic from business-critical applications must be protected from other types of traffic. Requests from business-critical and delay-sensitive applications must be serviced with priority. The employment and enforcement of quality of service (QoS) policies within a network plays an essential role in enabling network administrators and architects to meet networked application demands. QoS is a crucial element of any administrative policy that mandates how to handle application traffic on a network. This module introduces the concept of QoS, explains key issues of networked applications, and describes different methods for implementing QoS.

Module Objectives

Upon completing this module, you will be able to explain the need to implement QoS and explain methods for implementing and managing QoS.

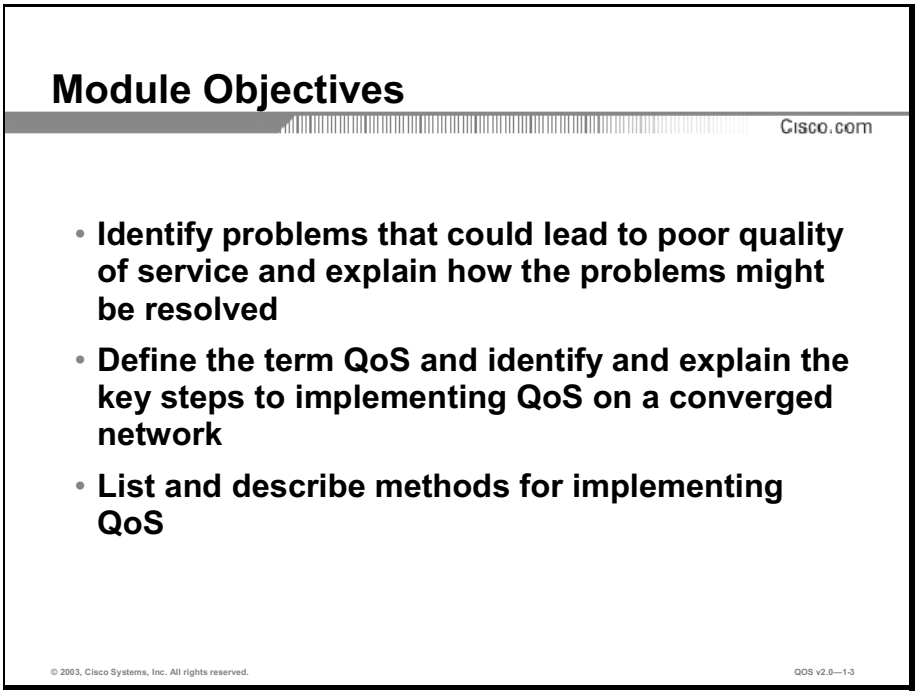

Module Outline

The outline lists the components of this module.

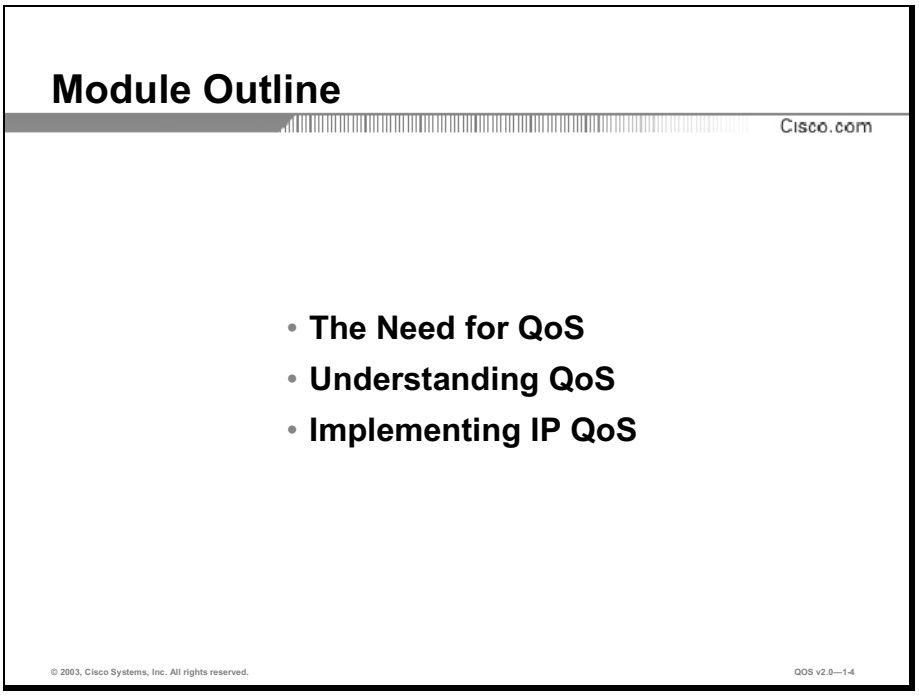

The Need for QoS

Overview

A communications network forms the backbone of any successful organization. These networks transport a multitude of applications and data, including high-quality video and delay-sensitive data such as real-time voice. Bandwidth-intensive applications stretch network capabilities and resources, but also complement, add value, and enhance every business process. Networks must provide secure, predictable, measurable, and sometimes guaranteed services. Achieving the required QoS by managing delay, delay variation (jitter), bandwidth, and packet loss parameters on a network becomes the secret to a successful end-to-end business solution. QoS is the set of techniques used to manage network resources.

Relevance

With the emergence of networks incorporating many types of traffic with different requirements, QoS has become an essential component of networking. As more converged networks are implemented, QoS becomes more important.

Objectives

Upon completing this lesson, you will be able to identify problems that could lead to poor quality of service and explain how the problems might be resolved. This includes being able to meet these objectives:

- Describe a converged IP network supporting voice, video, and data traffic
- Identify the four key quality issues with converged networks
- Explain how a lack of bandwidth can cause quality problems and ways to resolve those problems
- Explain how end-to-end delay can cause quality problems and ways to resolve those problems
- Explain how packet loss can cause quality problems and ways to resolve those problems

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Basic knowledge of internetworking with TCP/IP concepts

Outline

The outline lists the topics included in this lesson.

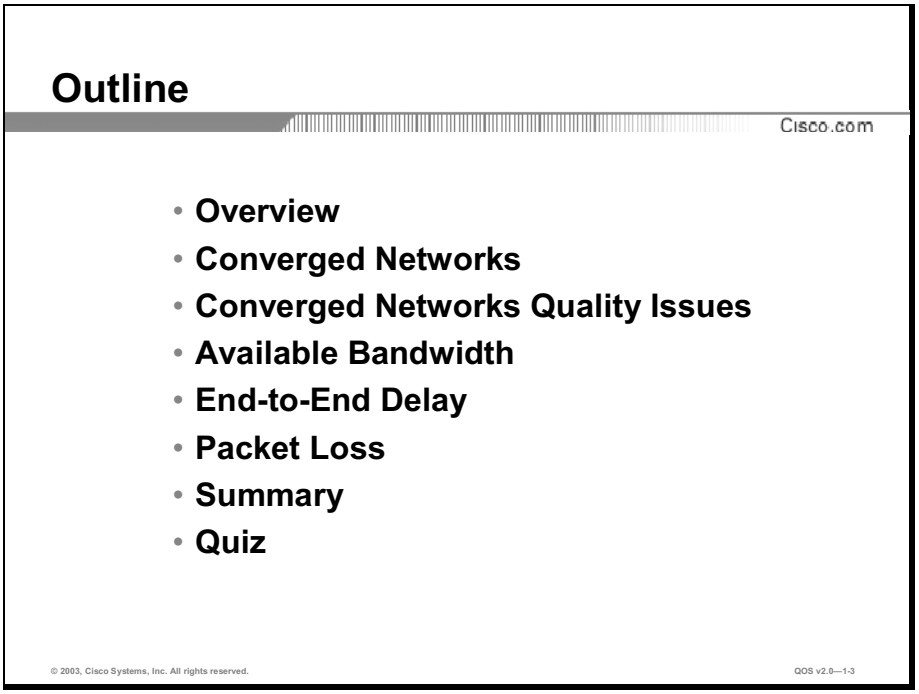

Converged Networks

This topic explains why QoS was not important before networks converged.

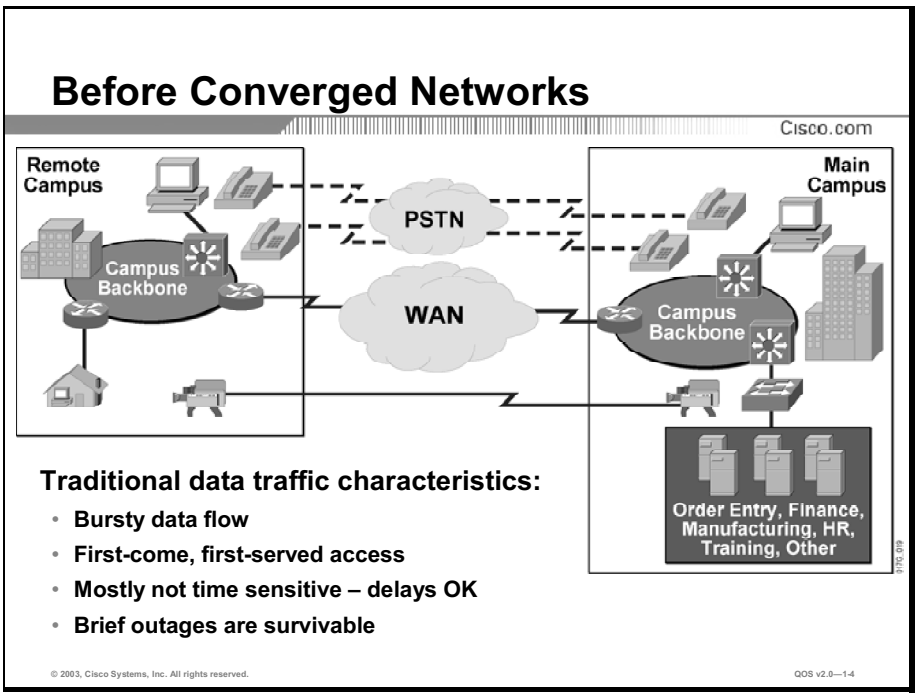

Before networks converged, network engineering focused on connectivity. The rates at which data came onto the network resulted in bursty data flows. Data, arriving in packets, tried to grab as much bandwidth as it could at any given time. The access was very egalitarian—a first-come, first-served basis. Whoever arrived first got the bandwidth.

As a result of this somewhat anarchic way of attacking the network, the data rate is adaptive to network conditions.

The protocols that have been developed have adapted to the bursty nature of data networks, and brief outages are survivable. For example, if retrieving e-mail, a delay of a few seconds is generally not noticeable. A delay of minutes is annoying, but not serious.

Traditional networks also had requirements for applications such as data, video, and systems network architecture (SNA). Since each application had different traffic characteristics and requirements, network designers deployed nonintegrated networks designed for carrying a specific type of traffic: data network, SNA network, voice network, and video network.

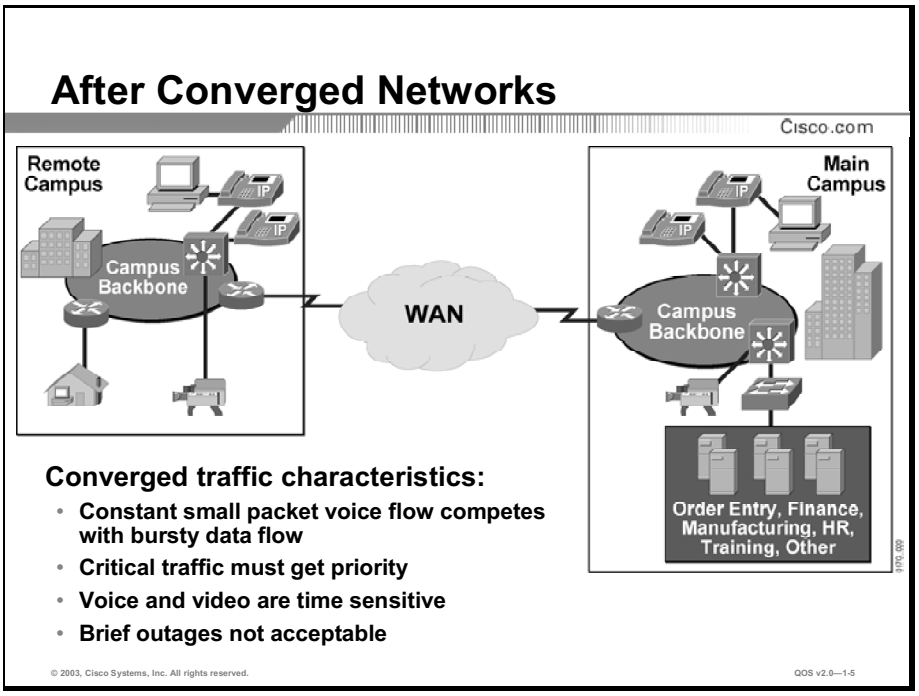

The figure illustrates a converged network in which voice, video, and data traffic use the same network facilities. Merging these different traffic streams with dramatically differing requirements can lead to a number of problems.

Although packets carrying voice traffic are typically very small, they cannot tolerate delay and delay variation as they traverse the network. Voices will break up and words will become incomprehensible.

On the other hand, packets carrying file transfer data are typically large and can survive delays and drops. It is possible to retransmit part of a dropped (data) file, but it is not feasible to retransmit a part of a (voice) conversation.

The constant, but small packet voice flow competes with bursty data flows. Unless some mechanism mediates the overall flow, voice quality will be severely compromised at times of network congestion. The critical voice traffic must get priority.

Voice and video traffic are very time sensitive. They cannot be delayed and they cannot be dropped or the resulting quality of voice and video will suffer.

And, finally, converged networks cannot fail. While a file transfer or e-mail packet can wait until the network recovers, voice and video packets cannot. Even a brief network outage on a converged network can seriously disrupt business operations.

Converged Networks Quality Issues

This topic describes the basic quality issues presented by converged networks.

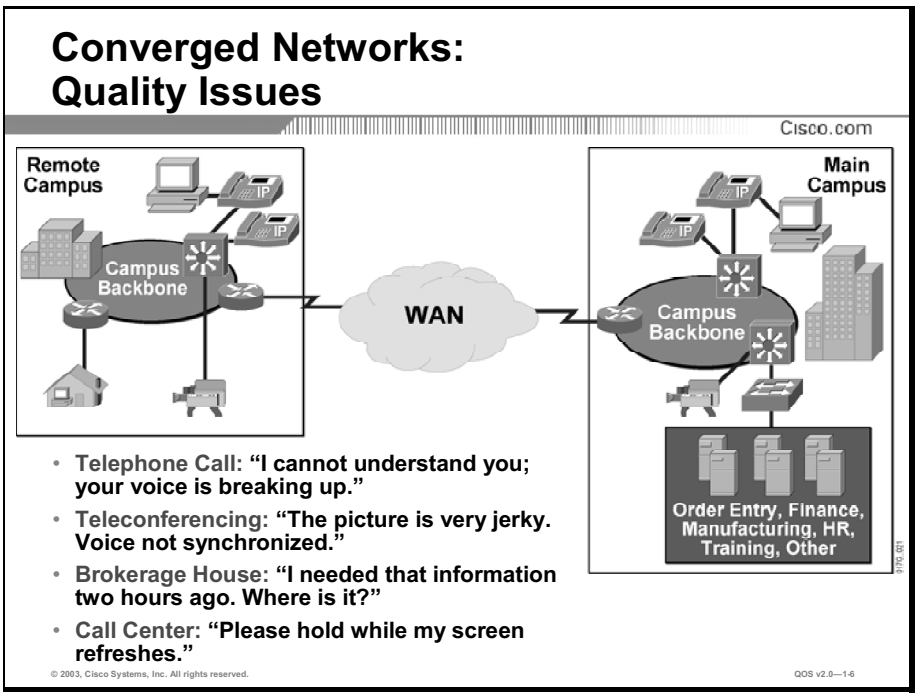

With inadequate preparation of the network, voice transmission is choppy or unintelligible. Gaps in speech are particularly troublesome where pieces of speech are interspersed with silence, and speech literally disappears. In voice-mail systems this silence is a problem. For example, you dial 68614. In a situation where the gaps in speech are actually gaps in the tone, 68614 becomes 6688661144, because the gaps in speech are perceived as pauses in the touch tones.

Poor caller interactivity is the consequence of delay. It causes two problems—echo and talker overlap.

- Echo is caused by the signal reflecting the speaker voice from the far-end telephone equipment back into the speaker ear.
- Talker overlap is caused when one-way delay becomes greater than 250 ms. When this occurs, one talker steps in on the speech of the other talker resulting in a "walkie-talkie" call mode.

Disconnected calls are the worst cases: If there are long gaps in speech, the parties will hang up; if there are signaling problems, the calls are disconnected. Such events are completely unacceptable in the voice world yet are quite common for an inadequately prepared data network that is attempting to carry voice.

Multimedia streams, such as those used in IP telephony or videoconferencing, may be extremely sensitive to delivery delays and create unique QoS demands on the underlying networks that carry them. When packets are delivered using the "best-effort" delivery model, they may not arrive in order, in a timely manner, or at all. The result is unclear pictures, jerky and slow movement, and sound that is out of synchronization with the image.

Converged Networks: Quality Issues (Cont.)

Cisco.com

- **Lack of bandwidth: multiple flows compete for a limited amount of bandwidth**
- **End-to-end delay (fixed and variable): packets have to traverse many network devices and links that add up to the overall delay**
- **Variation of delay (jitter): sometimes there is a lot of other traffic, which results in more delay**
- **Packet Loss: packets may have to be dropped when a link is congested**

The four big problems facing converged enterprise networks are bandwidth capacity, delay (both fixed and variable), variation of delay (also called jitter), and packet loss.

Large graphic files, multimedia uses, and increasing use for voice and video cause bandwidth capacity problems over data networks.

Delay is the time it takes for a packet to reach the receiving endpoint after being transmitted from the sending endpoint. This time is termed the “end-to-end delay,” and consists of two components: fixed network delay and variable network delay. Jitter is the delta, or difference, in the total end-to-end delay values of two voice packets in the voice flow.

Two types of fixed delay are serialization and propagation delays. Serialization is the process of placing bits on the circuit. The higher the circuit speed, the less time it takes to place the bits on the circuit. Therefore, the higher the speed of the link, the less serialization delay is incurred. Propagation delay is the time it takes for frames to transit the physical media.

Processing delay is a type of variable delay, and is the time required by a networking device to look up the route, change the header, and complete other switching tasks. In some cases, the packet also must be manipulated. For example, the encapsulation type or the hop count must be changed. Each of these steps can contribute to the processing delay.

Loss of packets is usually caused by congestion in the WAN, resulting in speech dropouts or a stutter effect if the play-out side tries to accommodate by repeating previous packets.

Available Bandwidth

This topic explains how a lack of bandwidth can adversely impact QoS in a network and describes ways to effectively increase bandwidth on a link.

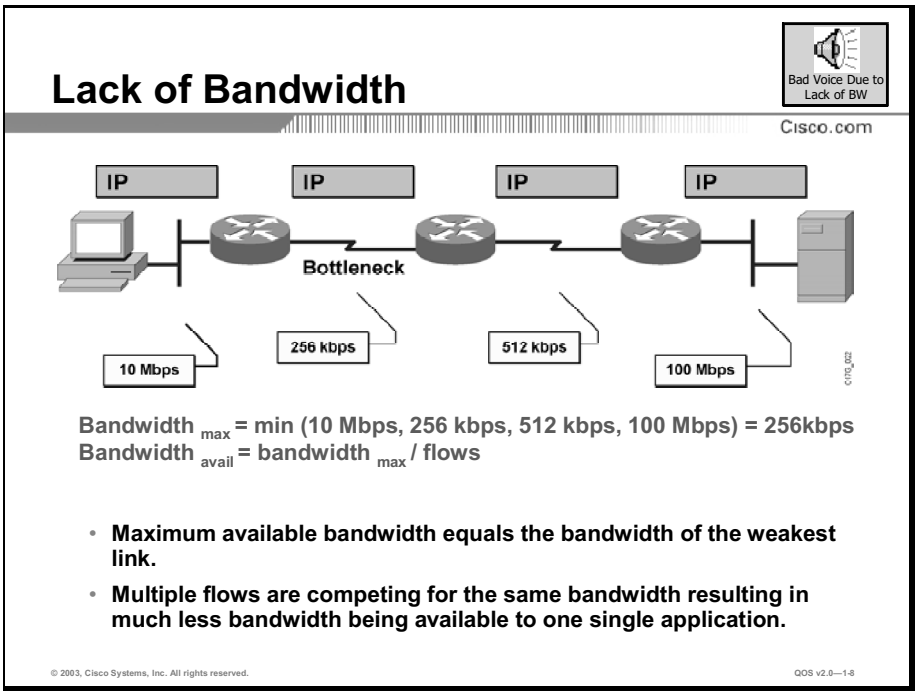

The example illustrates an empty network with four hops between a server and a client. Each hop is using different media with a different bandwidth. The maximum available bandwidth is equal to the bandwidth of the slowest link.

The calculation of the available bandwidth, however, is much more complex in cases where multiple flows are traversing the network. The calculation of the available bandwidth in the illustration is a rough approximation.

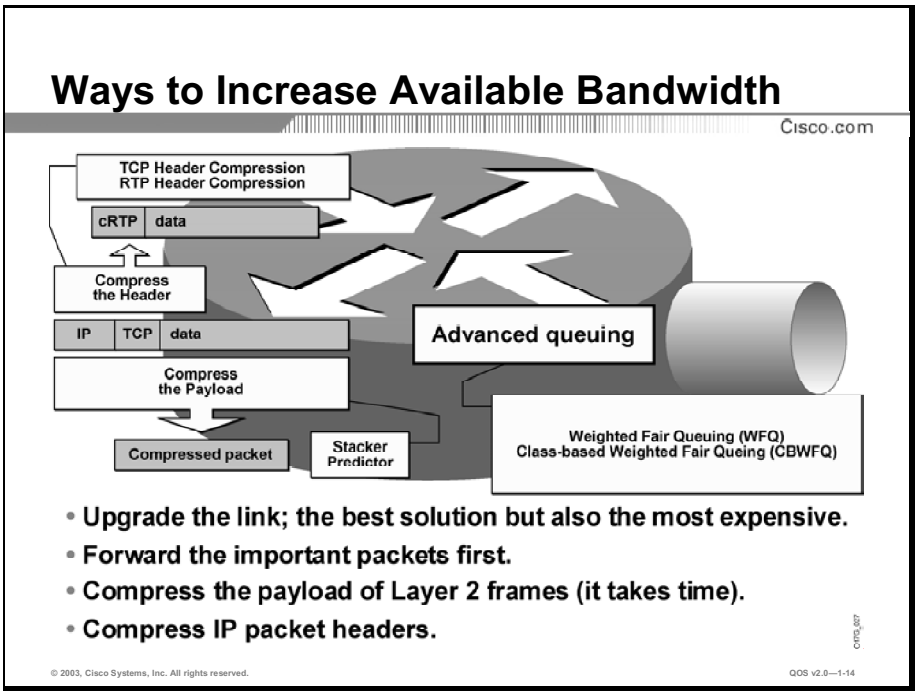

The best approach is to increase the link capacity to accommodate all applications and users, with some extra bandwidth to spare. Although this solution sounds simple, it brings a high cost in terms of the expense and time it takes to implement. Very often there are also technological limitations for upgrading to a higher bandwidth.

Another option is to classify traffic into QoS classes and prioritize it according to importance. (Voice and business-critical traffic should get sufficient bandwidth to support their application requirements, voice should get prioritized forwarding, and the least important traffic should get whatever unallocated bandwidth is remaining.) A wide variety of mechanisms are available in Cisco IOS QoS software that provide bandwidth guarantees:

- Priority queuing (PQ) or custom queuing (CQ)
- Modified deficit round robin (MDRR) (on Cisco 12000 series routers)
- Distributed type of service (ToS)-based and QoS group-based weighted fair queuing (WFQ) (on Cisco 7*x*00 series routers)
- Class-based weighted fair queuing (CBWFQ)
- Low-latency queuing (LLQ)

Optimizing link usage by compressing the payload of frames (virtually) increases the link bandwidth. Compression, on the other hand, also increases delay because of the complexity of compression algorithms. Using hardware compression can accelerate packet payload compressions. Stacker and Predictor are two compression algorithms that are available in Cisco IOS software.

Another link efficiency mechanism is header compression. This mechanism is especially effective in networks where most packets carry small amounts of data (that is, where payload-to-header ratio is small). Typical examples of header compression are TCP header compression and Real-Time Transport Protocol (RTP) header compression.

End-to-End Delay

This topic explains how end-to-end delay can adversely impact QoS in a network and describes ways to effectively reduce delay.

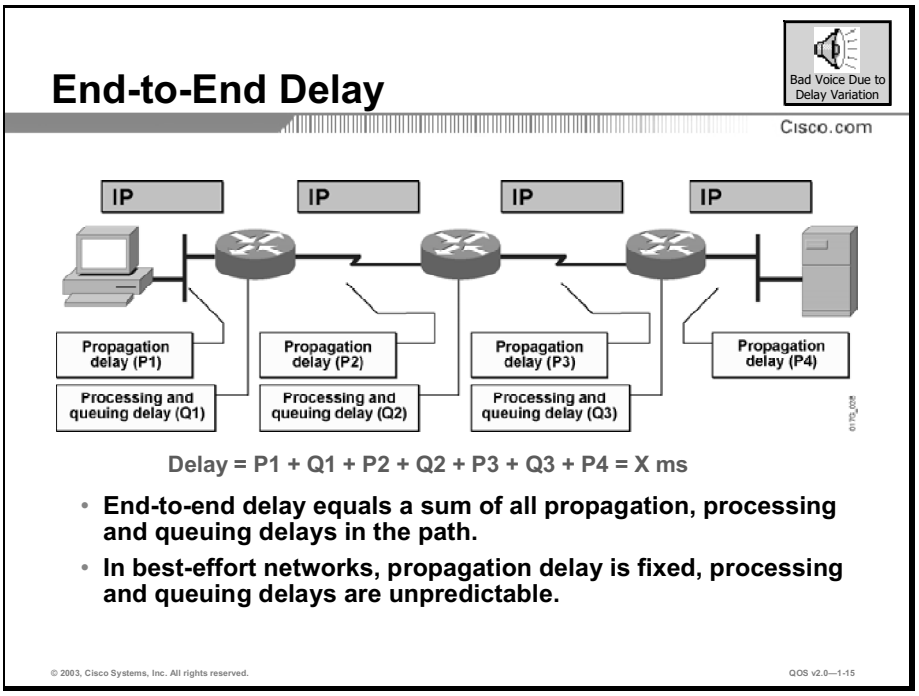

This figure illustrates the impact a network has on the end-to-end delay of packets going from one end to the other. Each hop in the network adds to the overall delay because of these factors:

- Propagation delay is caused by the speed of light traveling in the media; for example, the speed of light traveling in fiber optics or copper media.
- Serialization delay is the time it takes to clock all the bits in a packet onto the wire. This is a fixed value that is a function of the link bandwidth.
- There are processing and queuing delays within a router, which can be caused by a wide variety of conditions.

Propagation delay is generally ignored but it can be significant; for example, about 40 ms coast-to-coast, over optical. Internet Control Message Protocol (ICMP) echo (ping) is one way to measure the round-trip time of IP packets in a network.

Example: Effects of Delay

A customer has a router in New York and a router in San Francisco, each connected by a 128-kbps WAN link. The customer sends a 66-byte voice frame across the link. To transmit the frame (528 bits) it will take 4.125 ms to clock out (serialization delay). However, the last bit will not arrive until 40 ms *after* it clocks out (propagation delay). The total delay equals 44.125 ms.

Now, change the circuit to a T1. To transmit the frame (528 bits) it will take 0.344 ms to clock out (serialization delay). However, the last bit will not arrive until 40 ms after transmission (propagation delay) for a total delay of 40.344 ms. In this case, the significant factor is propagation delay. In the same situation—but for a link between Seattle and San Francisco—serialization delay remains the same and propagation delay drops to around 6 ms, making 528 bits take 10.125 (128-kbps link) and 6.344 (T1 link).

As is evident, you must take both serialization and propagation delays into account.

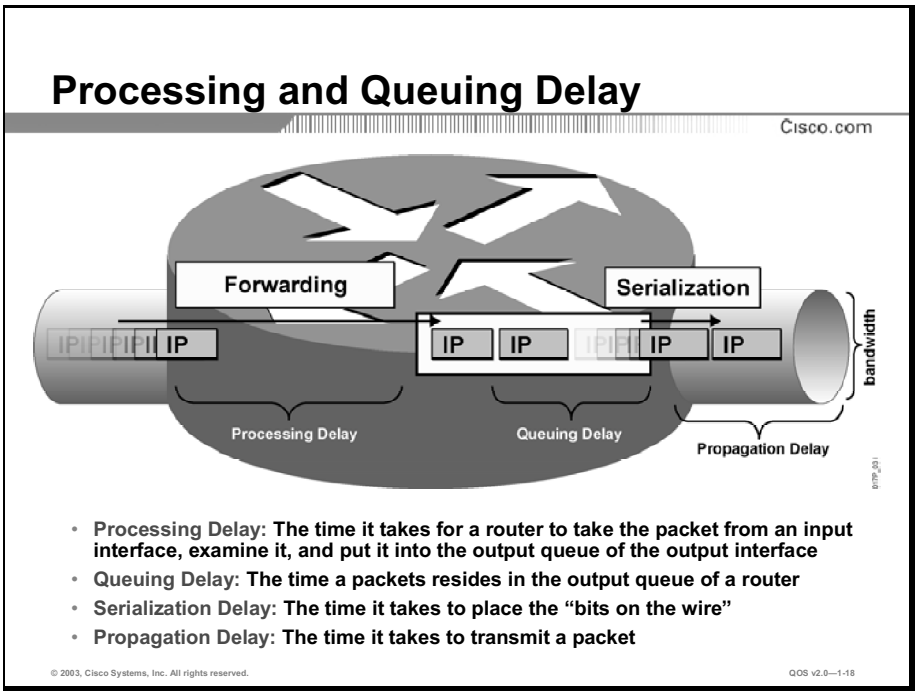

There are four kinds of delay:

- **Processing Delay:** The time it takes for a router to take the packet from an input interface and put it into the output queue of the output interface. The processing delay depends on various factors, such as:
  - CPU speed
  - CPU utilization
  - IP switching mode
  - Router architecture
  - Configured features on both input and output interface
- **Queuing Delay:** The time a packet resides in the output queue of a router. It depends on the number and sizes of packets already in the queue and on the bandwidth of the interface. It also depends on the queuing mechanism.
- **Serialization Delay:** The time it takes to place a frame on the physical medium for transport.
- **Propagation Delay:** The time it takes to transmit a packet. (This usually depends on the type of media interface.)

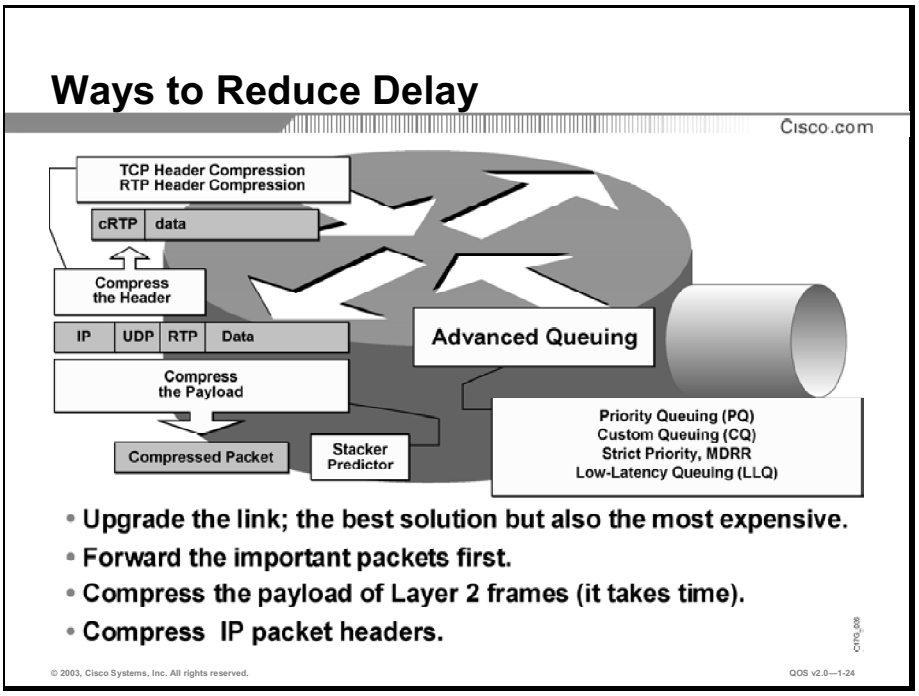

Assuming that a router is powerful enough to make a forwarding decision in negligible time, it can be said that most processing, queuing, and serialization delay is influenced by the following factors:

- Average length of the queue
- Average length of packets in the queue
- Link bandwidth

There are several approaches to accelerate packet dispatching of delay-sensitive flows:

- **Increase link capacity:** Sufficient bandwidth causes queues to shrink, making sure packets do not have to wait long before they can be transmitted. Additionally, more bandwidth reduces serialization time. On the other hand, this might be an unrealistic approach because of the costs associated with the upgrade.
- **Prioritize delay-sensitive packets:** This is a more cost-effective approach. There are a wide variety of queuing mechanisms available in Cisco IOS software that have pre-emptive queuing capabilities, for example:
  - PQ
  - CQ
  - Strict-priority or alternate priority queuing within the MDRR (on Cisco 12000 series routers)
  - LLQ
- **Compress payload:** Payload compression reduces the size of packets and, therefore, virtually increases link bandwidth. Additionally, compressed packets are smaller and need less time to be transmitted. On the other hand, compression uses complex algorithms that take time and add to the delay. This approach is, therefore, not used to provide low-delay propagation of packets.

- **Header compression:** Header compression is not as CPU-intensive as payload compression and can be used in combination with other mechanisms to reduce delay. It is especially useful for voice packets that have a bad payload-to-header ratio, which is improved by reducing the header of the packet (RTP header compression). By minimizing delay, jitter is also reduced (delay is more predictable).

Packet Loss

This topic explains how packet loss can adversely impact QoS in a network and describes ways to manage packet loss so that QoS is not affected.

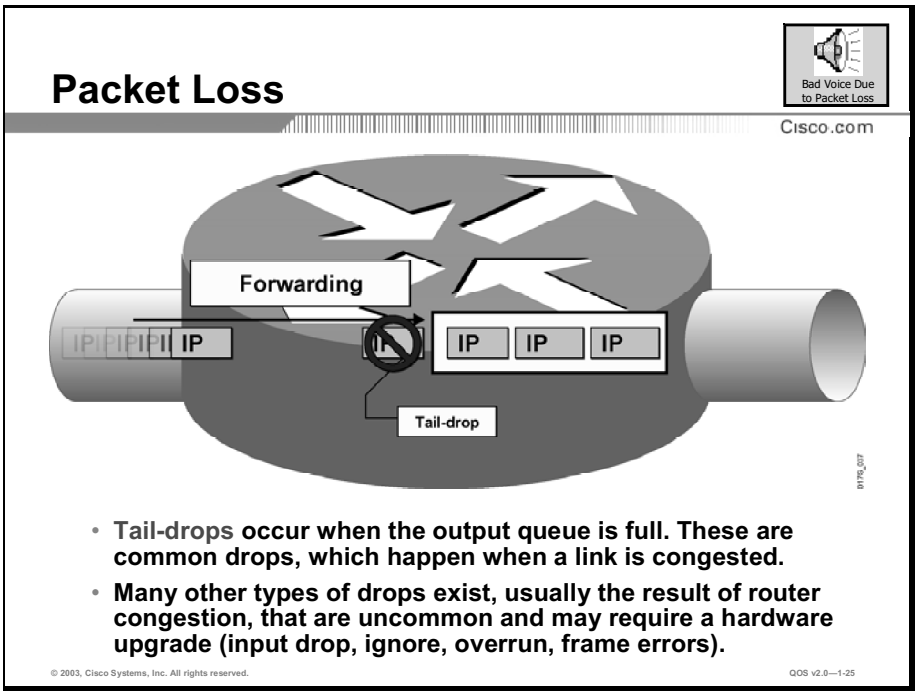

The usual packet loss occurs when routers run out of buffer space for a particular interface (output queue). The figure illustrates a full interface output queue, which causes newly arriving packets to be dropped. The term used for such drops is simply “output drop” or “tail-drop” (packets are dropped at the tail of the queue).

Routers might also drop packets for other (less common) reasons, for example:

- **Input queue drop:** Main CPU is congested and cannot process packets (the input queue is full).
- **Ignore:** Router ran out of buffer space.
- **Overrun:** CPU is congested and cannot assign a free buffer to the new packet.
- **Frame errors:** Hardware-detected error in a frame—cyclic redundancy check (CRC), runt, giant.

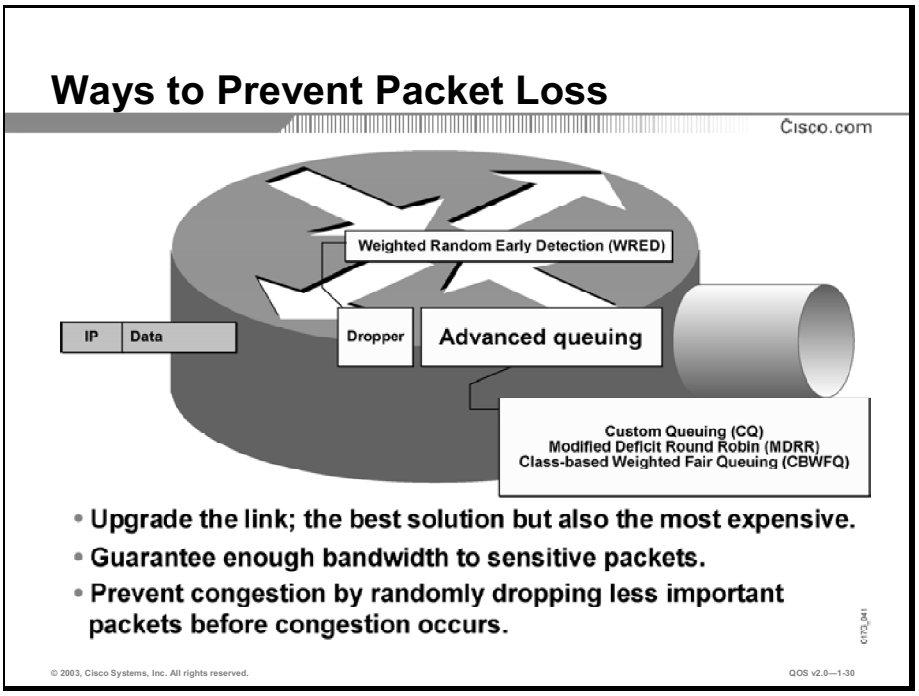

Packet loss is usually the result of congestion on an interface. Most applications that use TCP experience slowdown because TCP adjusts to the network resources. (Dropped TCP segments cause TCP sessions to reduce their window sizes.) There are some other applications that do not use TCP and cannot handle drops (fragile flows).

The following approaches can be taken to prevent drops of sensitive applications:

- Increased link capacity to ease or prevent congestion.
- Guarantee enough bandwidth and increase buffer space to accommodate bursts of fragile applications. There are several mechanisms available in Cisco IOS QoS software that can guarantee bandwidth and provide prioritized forwarding to drop-sensitive applications, for example:
  - PQ
  - CQ
  - MDDR (on Cisco 12000 series routers)
  - IP RTP prioritization
  - CBWFQ
  - LLQ
- Prevent congestion by dropping other packets before congestion occurs. Weighted random early detection (WRED) can be used to start dropping other packets before congestion occurs.

There are some other mechanisms that can also be used to prevent congestion:

- Traffic shaping delays packets instead of dropping them (generic traffic shaping [GTS], Frame Relay traffic shaping [FRTS], and class-based shaping).

Traffic policing can limit the rate of less important packets to provide better service to drop-sensitive packets (committed access rate [CAR] and class-based policing).

Summary

This topic summarizes the key points discussed in this lesson.

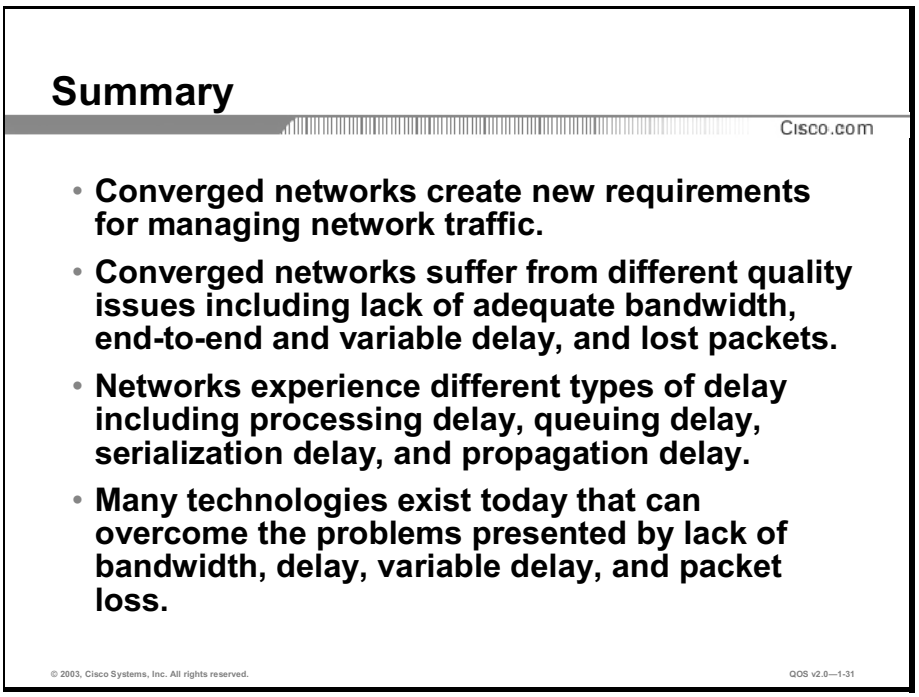

References

For additional information, refer to this resource:

- To learn more about QoS, refer to “Cisco IOS Quality of Service (QoS)” at the following URL: http://www.cisco.com/warp/public/732/Tech/qos/

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) Which three types of traffic can typically tolerate packets being dropped? (Choose three.)

A) file transfers

B) voice

C) e-mail

D) HTTP

Q2) Which three of the following are key quality issues specifically for converged networks? (Choose three.)

A) lack of bandwidth

B) end-to-end delay

C) variable delay

D) propagation delay

Q3) The maximum bandwidth available between two points is _____.

A) the bandwidth of the slowest link

B) the bandwidth of the fastest link

C) the average bandwidth across the links

D) the average of the slowest and fastest links

Q4) What are three ways to reduce delay for time-sensitive packets in a network? (Choose three.)

A) compress headers

B) forward the most important packets first

C) upgrade bandwidth on the links

D) aggressively drop packets

Q5) Which is the most common drop to occur when a link is congested?

A) tail drop

B) input drop

C) overrun drop

D) no buffer drop

Quiz Answer Key

Q1) A, C, D

Relates to: Converged Networks

Q2) A, B, C

Relates to: Converged Networks Quality Issues

Q3) A

Relates to: Available Bandwidth

Q4) A, B, C

Relates to: End-to-End Delay

Q5) A

Relates to: Packet Loss

Understanding QoS

Overview

The basic concepts and key terminology of QoS are explained in this lesson. The three key steps involved in implementing a QoS policy are described.

Relevance

To understand the more technical aspects of network QoS, it is first important to understand the basic concepts of QoS and to be able to define some key QoS terms.

Objectives

Upon completing this lesson, you will be able to define the term QoS and identify and explain the key steps to implementing QoS on a converged network. This includes being able to meet these objectives:

- Define the term “QoS” with respect to traffic in a network
- List and explain the key steps involved in implementing a QoS policy on a network
- Explain the QoS requirements of the common types of network applications
- Define the term “QoS service policy”
- Define the term “QoS policy”

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Basic knowledge of internetworking with TCP/IP concepts

Outline

The outline lists the topics included in this lesson.

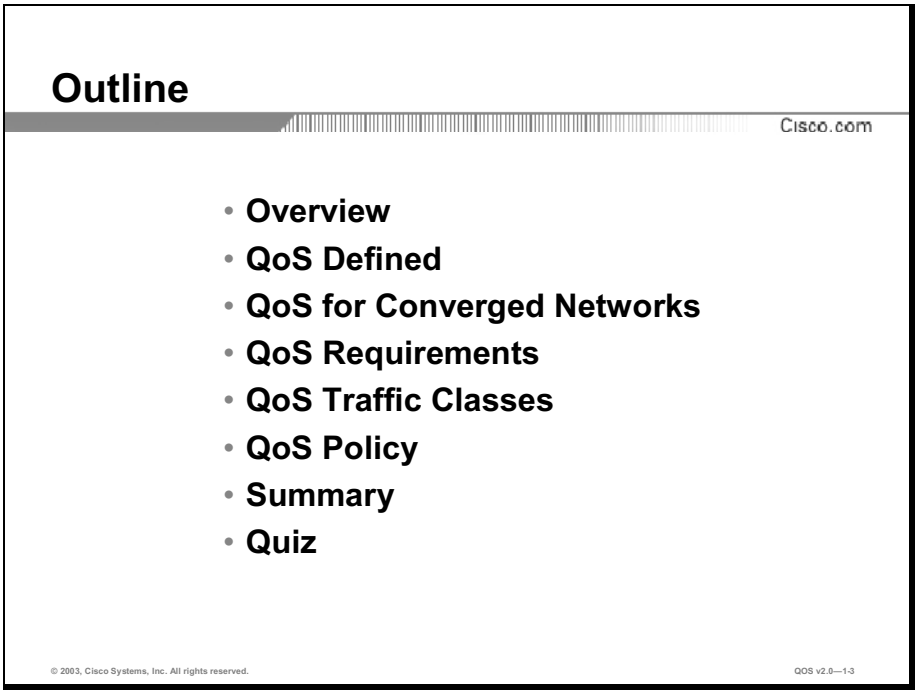

QoS Defined

This topic defines the term “QoS.”

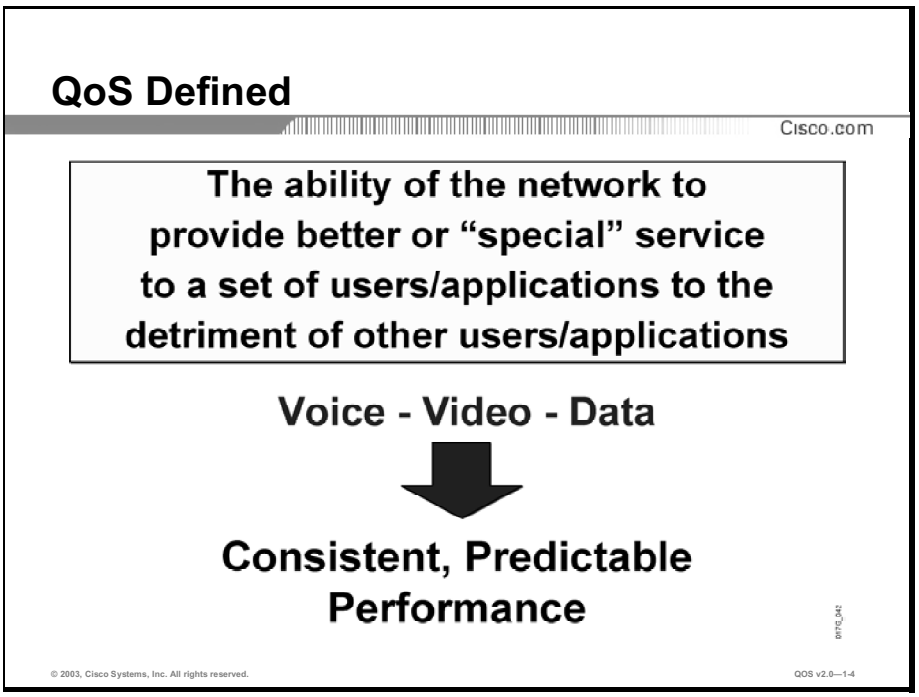

QoS is the ability of the network to provide better or “special” service to selected users and applications to the detriment of other users and applications.

Cisco IOS QoS features enable network administrators to control and predictably service a variety of networked applications and traffic types, thus allowing network managers to take advantage of a new generation of media-rich and mission-critical applications.

The goal of QoS is to provide better and more predictable network service by providing dedicated bandwidth, controlled jitter and latency, and improved loss characteristics. QoS achieves these goals by providing tools for managing network congestion, shaping network traffic, using expensive wide-area links more efficiently, and setting traffic policies across the network. QoS offers intelligent network services that, when correctly applied, help to provide consistent, predictable performance.

QoS for Converged Networks

This topic describes the three steps necessary for implementing QoS on a network.

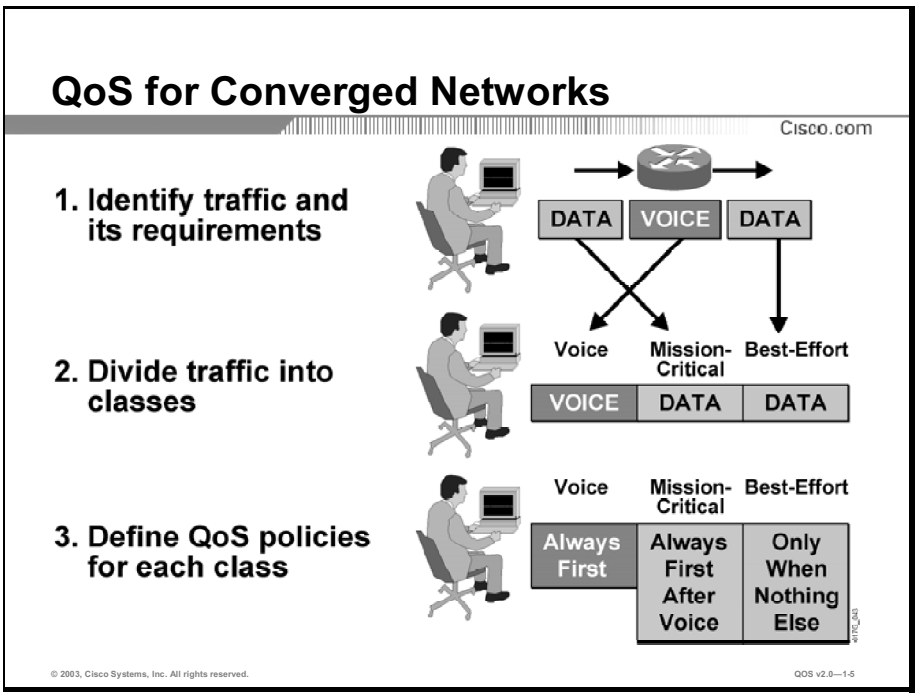

There are three basic steps involved in implementing QoS on a network:

Step 1 Identify traffic and its requirements. Study the network to determine the type of traffic running on the network and then determine the QoS requirements for the different types of traffic.

Step 2 Group the traffic into classes with similar QoS requirements. In the example below, four classes of traffic could be defined: v*oice, high priority, low priority,* and *browser*.

Step 3 Define QoS policies that will meet the QoS requirements for each traffic class.

Example: Three Steps to Implementing QoS on a Network

In a typical network, voice will always require absolute minimal delay. Some data associated with key applications will require very low delay (transaction-based data used in airline reservations or online banking applications). Other types of data can tolerate a great deal of delay (file transfers and e-mail). Nonbusiness network surfing can also be delayed or even prohibited.

A one-to-one mapping between traffic classes and QoS policies need not be made. For example, three QoS policies could be implemented to meet the requirements of the four traffic classes defined above:

- ***NoDelay:*** Assign to voice traffic
- ***BestService*:** Assign to high-priority traffic
- ***Whenever:*** Assign to both the low priority and browser traffic

QoS Requirements

This topic explains how traffic is identified on a network and describes elemental QoS requirements.

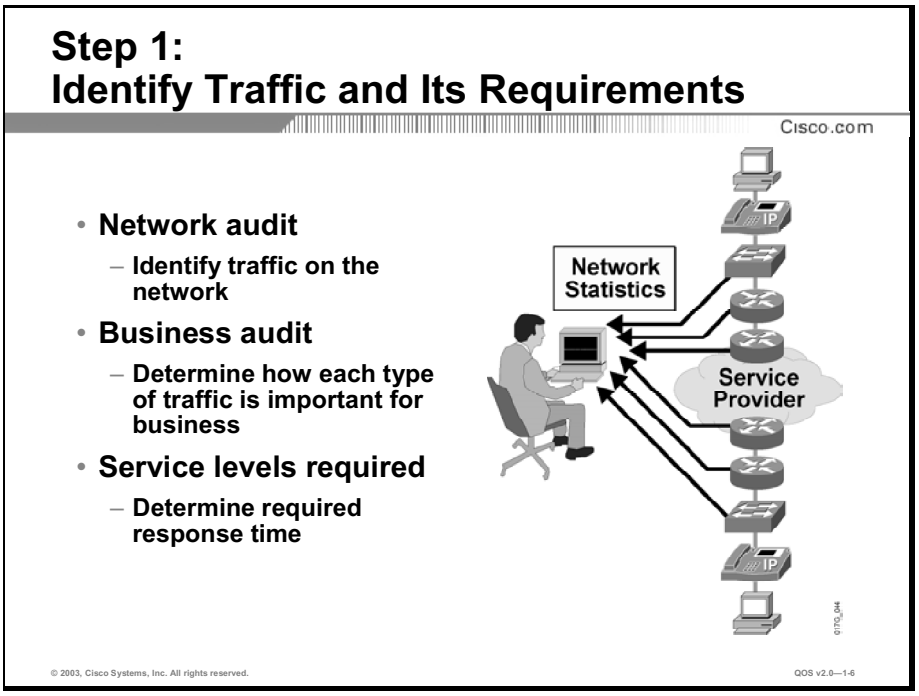

The first step in implementing QoS is identifying the traffic on the network and determining QoS requirements for the traffic.

Determine users QoS problems. Measure the traffic on the network during congested periods. Conduct CPU utilization assessment on each of their network devices during busy periods to determine where problems might be occurring.

Determine the business model, business goals, and obtain a list of business requirements. This will help you define the number of classes and determine the business requirements for each traffic class.

Define the service levels required by different traffic classes in terms of response time and availability. What is the impact on business if a transaction is delayed by two or three seconds? Can file transfers wait until the network is quiescent?

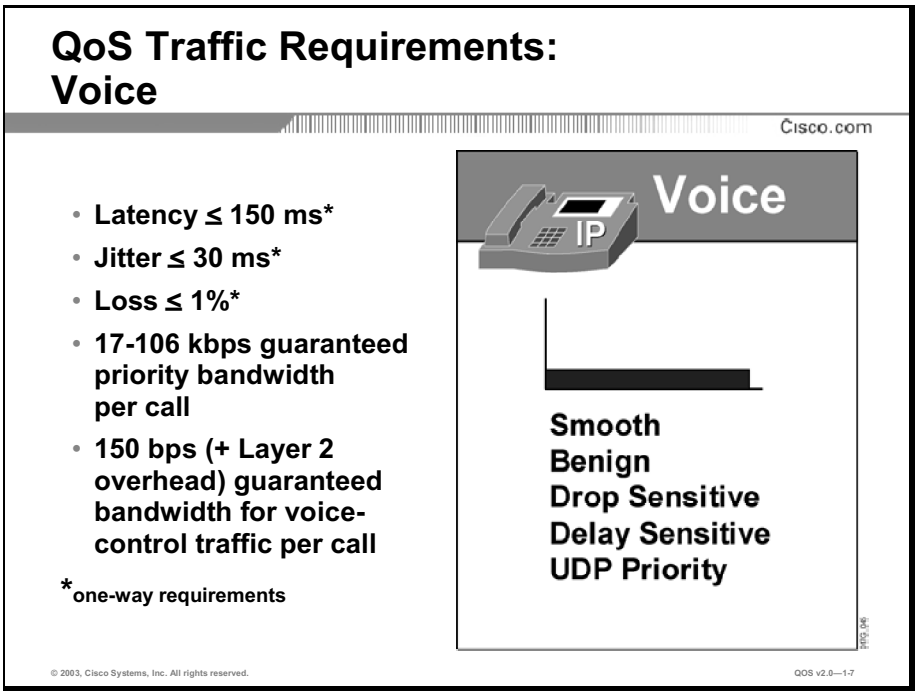

Voice traffic has extremely stringent QoS requirements. Voice traffic generally generates a smooth demand on bandwidth and has minimal impact on other traffic as long as it is managed.

While voice packets are typically small (60 to 120 bytes), they cannot tolerate delay or drops. The result of delays and drops are poor—and often unacceptable—voice quality. Because drops cannot be tolerated, User Datagram Protocol (UDP) is used to package voice packets because TCP retransmit capabilities have no value.

Voice packets can tolerate no more than a 15-ms delay (one-way requirement) and less than 1 percent packet loss.

A typical voice call will require 17 to 106 kbps of guaranteed priority bandwidth plus an additional 150 bps per call for voice-control traffic. Multiplying these bandwidth requirements times the maximum number of calls expected during the busiest time period will provide an indication of the overall bandwidth required for voice traffic.

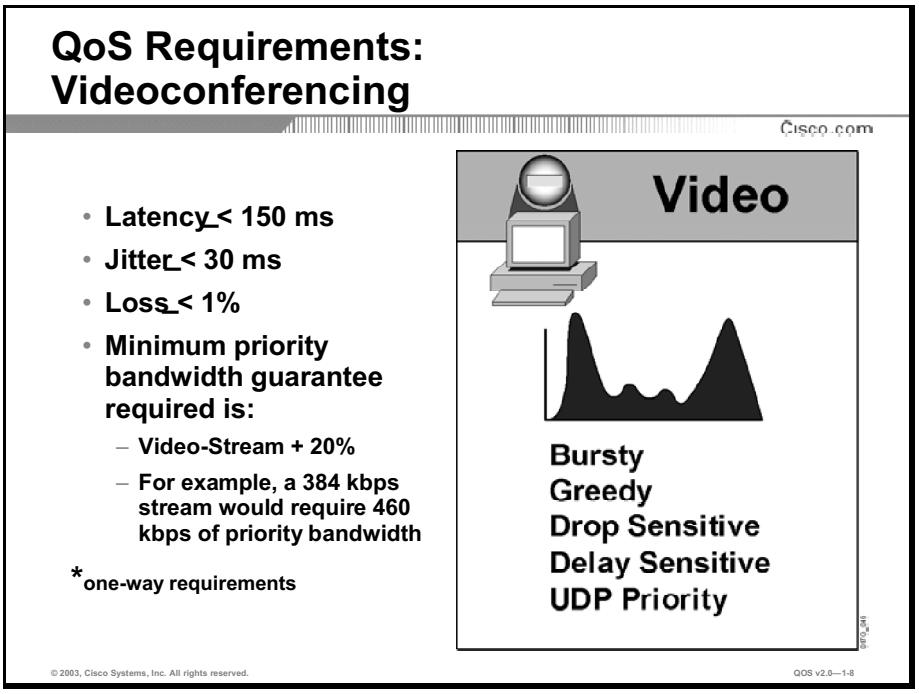

Videoconferencing applications also have very stringent QoS requirements very similar to voice.

But videoconferencing traffic is often bursty and greedy in nature and, as a result, can impact other traffic. Therefore, it is important to understand the videoconferencing requirements for a network and to provision carefully for it.

The minimum bandwidth for a videoconferencing stream would require the actual bandwidth of the stream (dependent upon the type of videoconferencing codec being used) plus some overhead. For example, a 384-kbps video stream would actually require a total of 460 kbps of priority bandwidth.

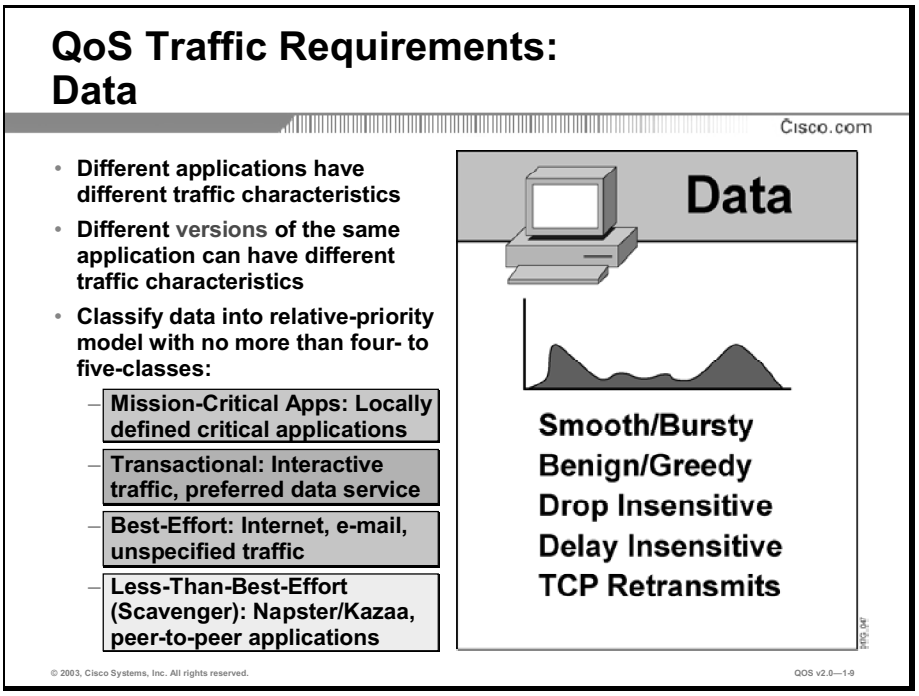

The QoS requirements for data traffic vary greatly.

Different applications (for example, a human resources application versus an automated teller machine application) may make greatly different demands on the network. Even different versions of the same application may have varying network traffic characteristics.

While data traffic can demonstrate either smooth or bursty characteristics depending upon the application, data traffic differs from voice and video in terms of delay and drop sensitivity. Almost all data applications can tolerate some delay and generally can tolerate high drop rates.

Because data traffic can tolerate drops, the retransmit capabilities of TCP become important and, as a result, many data applications use TCP.

In enterprise networks, important (business-critical) applications are usually easy to identify. Most applications can be identified based on TCP or UDP port numbers. Some applications use dynamic port numbers that, to some extent, make classifications more difficult. Cisco IOS software supports network-based application recognition (NBAR), which can be used to recognize dynamic port applications.

It is recommended that data traffic be classified into no more than four to five-classes as described in the graphic above. There will still remain additional classes for voice and video.

QoS Traffic Classes

This topic explains how to divide traffic into traffic classes.

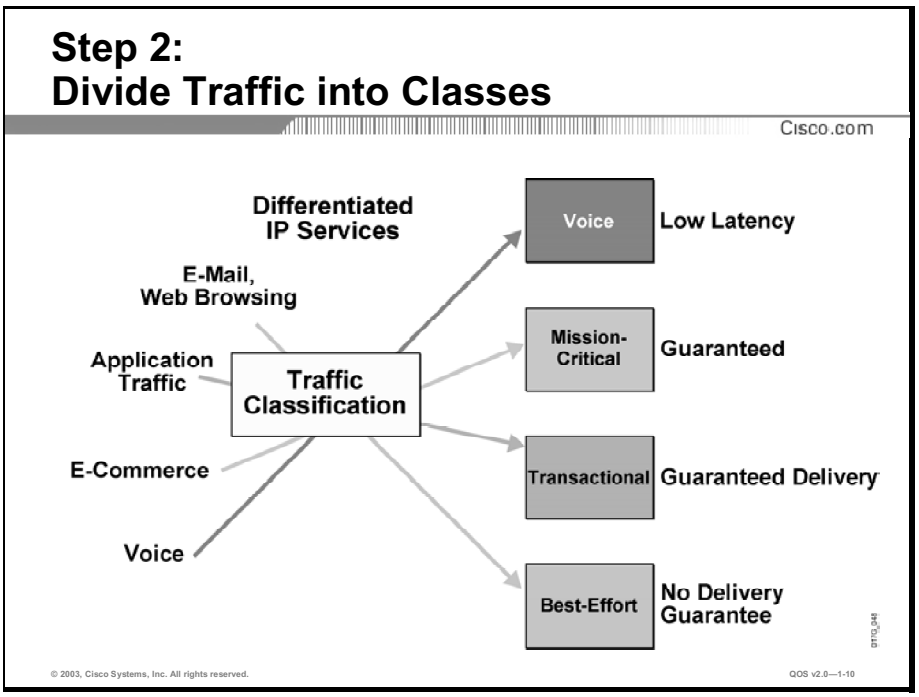

After the majority of network traffic has been identified and measured, use the business requirements to define traffic classes.

Because of its stringent QoS requirements, voice traffic will almost always exist in a class by itself. Cisco has developed specific QoS mechanisms such as LLQ that ensure that voice always receives priority treatment over all other traffic.

After the applications with the most critical requirements have been defined, the remaining traffic classes are defined using the business requirements.

Example: Traffic Classification

A typical enterprise might define five traffic classes as:

- **Voice:** Absolute priority for Voice over IP (VoIP) traffic
- **Mission critical:** Small set of locally defined critical business applications
- **Transactional:** Database access, transaction services, interactive traffic, preferred data services
- **Best effort:** Internet, e-mail
- **Scavenger (less-than-best-effort):** Napster/Kazaa and other point-to-point applications

QoS Policy

This topic describes how to define QoS policies after traffic classes have been defined.

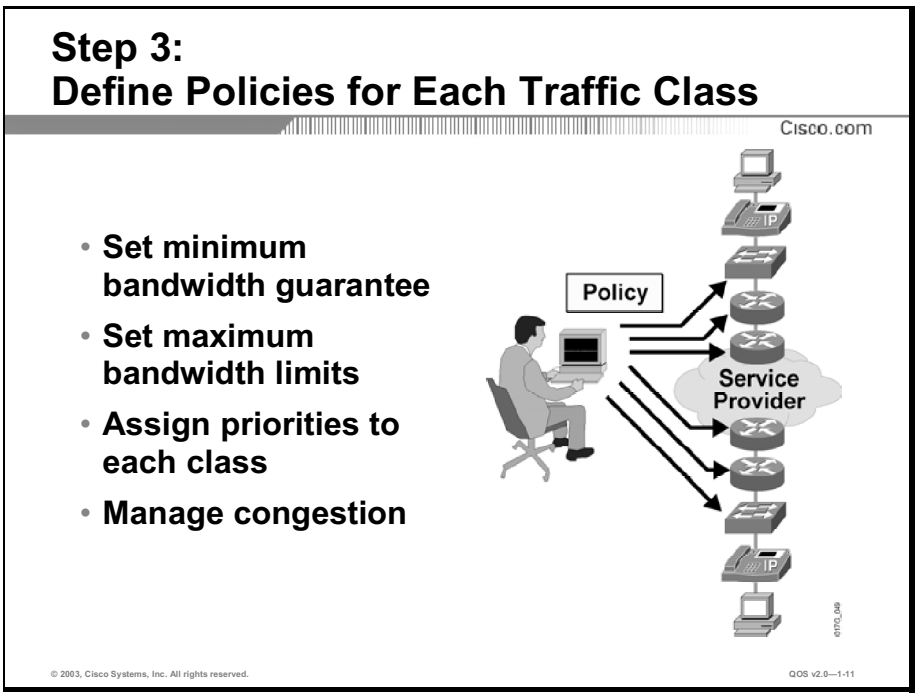

Finally, define a QoS policy for each traffic class. Defining a QoS policy involves:

- Setting a minimum bandwidth guarantee
- Setting a maximum bandwidth limit
- Assigning priorities to each class
- Using QoS technologies, such as advanced queuing, to manage congestion

Example: Defining QoS Policies

Using the traffic classes previously defined, QoS policies could be determined as:

- **Voice:** Minimum bandwidth 1 Mbps. Use QoS marking to mark voice packets as priority 5; use LLQ to always give voice priority.
- **Mission critical:** Minimum bandwidth 1 Mbps. Use QoS marking to mark critical data packets as priority 4; use CBWFQ to prioritize critical class traffic flows.
- **Best effort:** Maximum bandwidth 500 kbps. Use QoS marking to mark these data packets as priority 2; use CBWFQ to prioritize best-effort traffic flows that are below mission-critical and voice.
- **Scavenger:** Maximum bandwidth 100 kbps. Use QoS marking to mark less-than-best-effort (scavenger) data packets as priority 0; use weighted random early detection (WRED) to drop these packets whenever the network has a propensity for congestion.

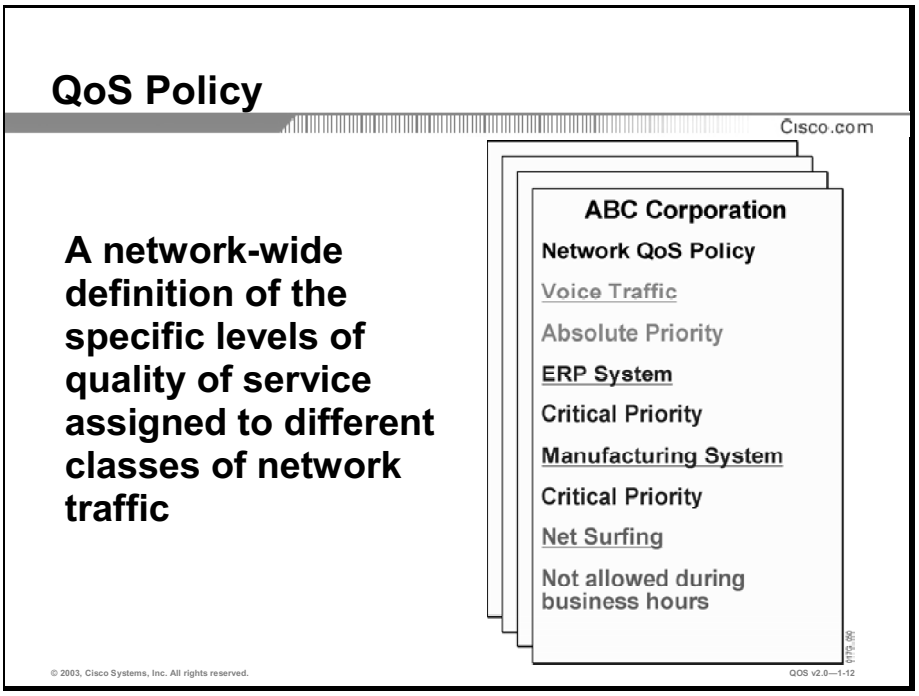

A QoS policy is a network-wide definition of the specific levels of QoS assigned to different classes of network traffic.

Having a QoS policy is just as important in a converged network as a security policy. A written and public QoS policy allows users to understand and negotiate for QoS in the network.

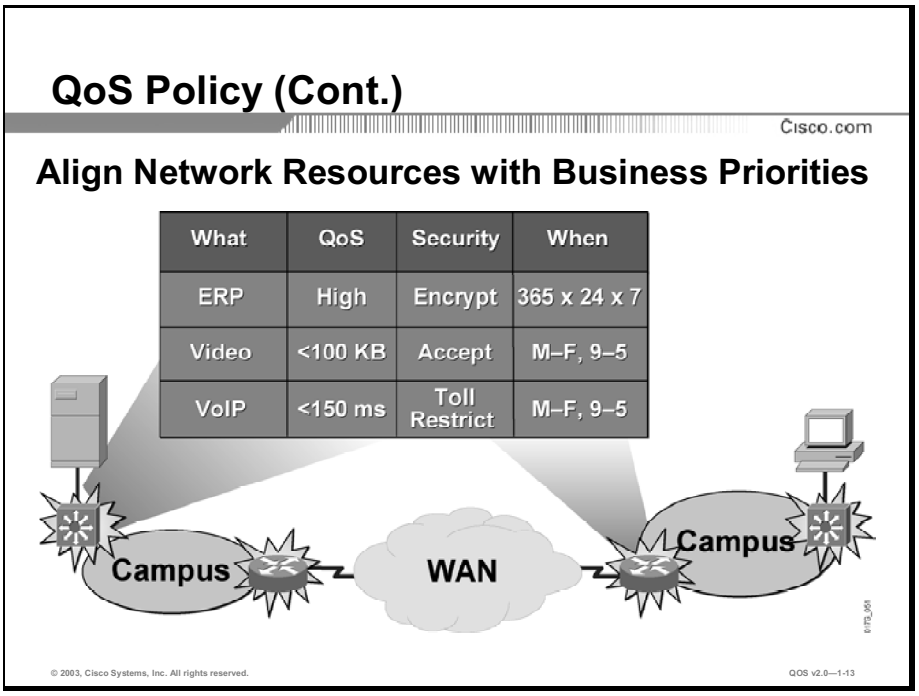

The graphic illustrates how a QoS policy could be defined for a network.

Enterprise Resource Planning (ERP) applications have a high QoS priority and must be available all the time.

Video applications are guaranteed 100 kbps of bandwidth, but can operate only between the hours of 9:00 a.m. to 5:00 p.m. on weekdays.

Voice traffic is guaranteed less than 150 ms delay in each direction but limited to the hours of 9:00 a.m. to 5:00 p.m. on weekdays; toll calls are completely restricted to avoid personal long-distance calls.

Summary

This topic summarizes the key points discussed in this lesson.

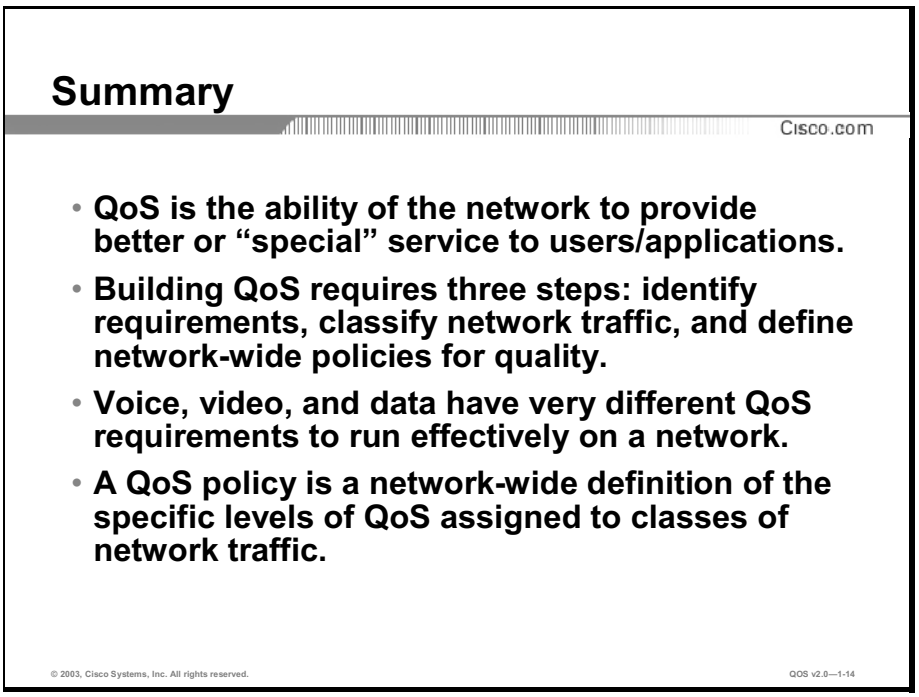

References

For additional information, refer to this resource:

- For more information on QoS, refer to “Implementing Quality of Service” at the following URL: http://www.cisco.com/en/US/partner/tech/tk543/tk757/technologies_white_paper09186a008017f93b.shtml

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) Quality of Service is defined as "the ability of the network to_____."

A) improve the quality of voice transmission

B) offer end-to-end circuits with preferred priority

C) provide special services to user and applications

D) consistently move priority packets to the front of queues

Q2) Which three of the following represent the three steps to implementing Quality of Service in converged networks? (Choose three.)

A) define QoS policies

B) divide traffic into classes

C) identify traffic and its requirements

D) interview users to determine problems

Q3) Which three of the following represent characteristics of voice traffic? (Choose three.)

A) drop sensitive

B) smooth, constant flow

C) benign, does not affect other traffic

D) relies on TCP to handle packet loss

Q4) Which type of application typically uses TCP for a transport protocol?

A) data

B) voice

C) video

D) videoconferencing

Q5) Which three of the following does QoS policy involve? (Choose three.)

A) assigning priorities to each class

B) setting a maximum bandwidth limit

C) setting a minimum bandwidth guarantee

D) combining traffic classes to go under one policy

Quiz Answer Key

Q1) C

Relates to: QoS Defined

Q2) A, B, C

Relates to: QoS for Converged Networks

Q3) A. B, C

Relates to: QoS Requirements

Q4) A

Relates to: QoS Requirements

Q5) A, B, C

Relates to: QoS Policy

Implementing IP QoS

Overview

Cisco recommends using either the Modular QoS command-line interface (MQC) or Cisco AutoQoS for implementing QoS in a network. The MQC offers a highly modular way to fine-tune a network. AutoQoS offers an automated method for almost instantly incorporating consistent voice QoS in a network of routers and switches. In addition, CiscoWorks QoS Policy Manager (QPM) provides centralized QoS design, administration, and traffic monitoring that scales to large QoS deployments.

Relevance

It is very important to understand the different ways to implement QoS in a network. There are tradeoffs regarding the use of the MQC versus AutoQoS in implementing QoS and knowing the differences between MQC and AutoQoS can help a network administrator save time and resources.

Objectives

Upon completing this lesson, you will be able to list and describe methods for implementing QoS. This includes being able to meet these objectives:

- List and describe methods for configuring and monitoring QoS on a network
- Explain at a high level, the CLI (nonmodularized) method of configuring QoS
- Explain at a high level, the MQC method of configuring QoS
- Explain at a high level, the AutoQoS method of configuring QoS
- Identify and explain the advantages and disadvantages of using each of the methods of implementing QoS on a network
- Explain how QPM can be used to manage QoS policies on a network
- Explain the purpose of a MIB and how it is used with QPM to monitor network
- Explain how the QoS MIBs can be used with QPM to monitor QoS on a network

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Basic understanding of the Cisco IOS command-line interface (CLI)

Outline

The outline lists the topics included in this lesson.

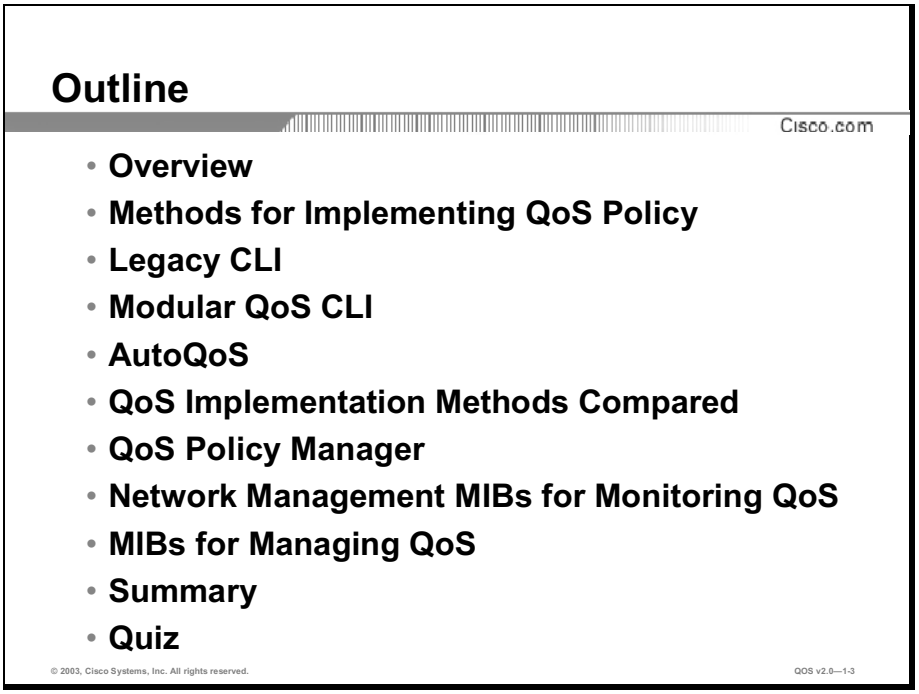

Methods for Implementing QoS Policy

This topic describes four different methods for implementing and managing a QoS policy.

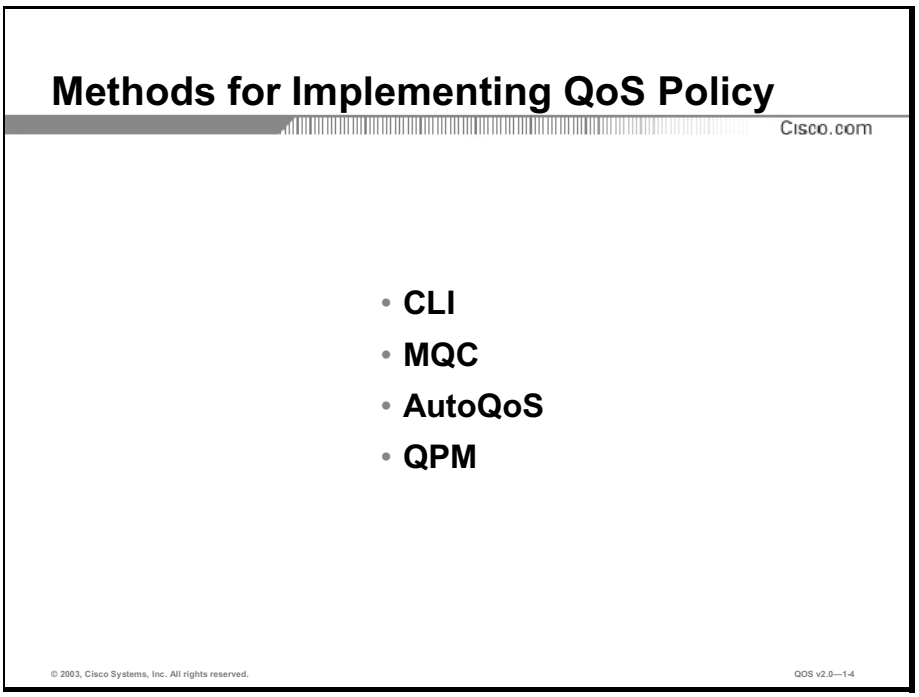

Just a few years ago, the only way to implement QoS in a network was by using the CLI to individually configure QoS policies at each interface. This was a time-consuming, tiresome, and error-prone task involving cutting and pasting configurations from one interface to another.

Cisco introduced the MQC in order to simplify QoS configuration by making configurations modular. Using MQC, QoS can be configured in a building block approach using a single module repeatedly to apply policy to multiple interfaces.

Cisco AutoQoS represents innovative technology that simplifies the challenges of network administration by reducing QoS complexity, deployment time, and cost in enterprise networks. Cisco AutoQoS incorporates value-added intelligence in Cisco IOS software and Cisco Catalyst software to provision and assist in the management of large-scale QoS deployments.

The first phase of Cisco AutoQoS offers straightforward capabilities to automate VoIP deployments for customers who want to deploy IP telephony but lack the expertise and staffing to plan and deploy IP QoS and IP services.

Customers can more easily provision and manage successful QoS deployments by using Cisco AutoQoS together with CiscoWorks QPM. Cisco AutoQoS provides QoS provisioning for individual routers and switches, simplifying deployment and reducing human error. CiscoWorks QPM provides centralized QoS design, administration, and traffic monitoring that scales to large QoS deployments.

Legacy CLI

This topic describes the CLI method for implementing QoS.

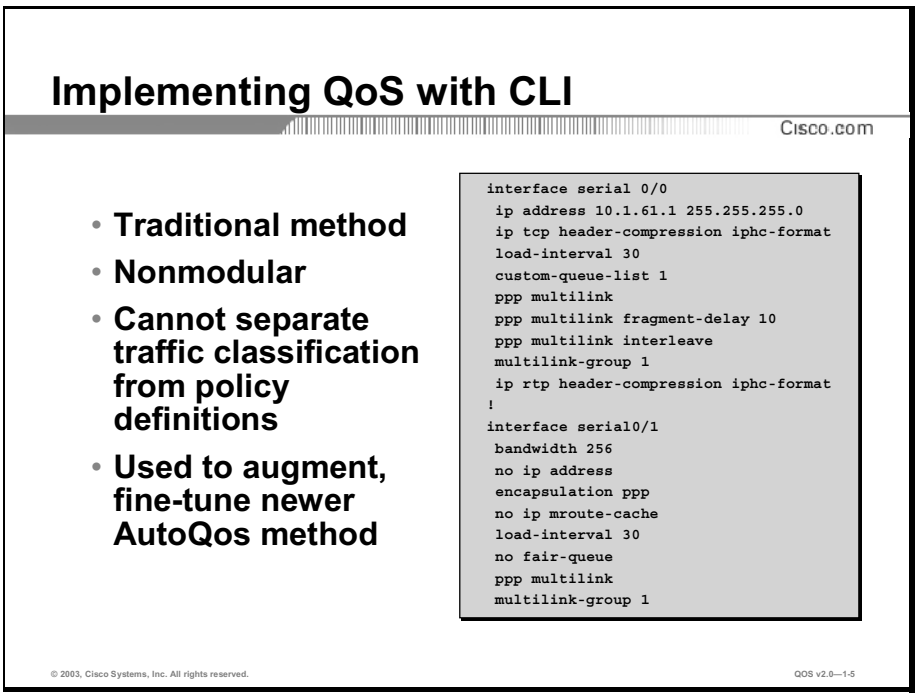

At one time, CLI was the only way to implement QoS in a network. It was a painstaking task involving copying one interface configuration and then pasting it into other interface configurations. It took a lot of time and patience.

The CLI method was nonmodular—there was no way to separate the classification of traffic from the actual definition of policy. Network administrators had to do both on every interface. The figure illustrates an example of the complex configuration tasks involved in using CLI.

While CLI is not recommended for implementing QoS policy, it is still used to fine-tune QoS implementations that have been generated using the Cisco AutoQoS macro.

Modular QoS CLI

This topic describes the MQC method for implementing QoS.

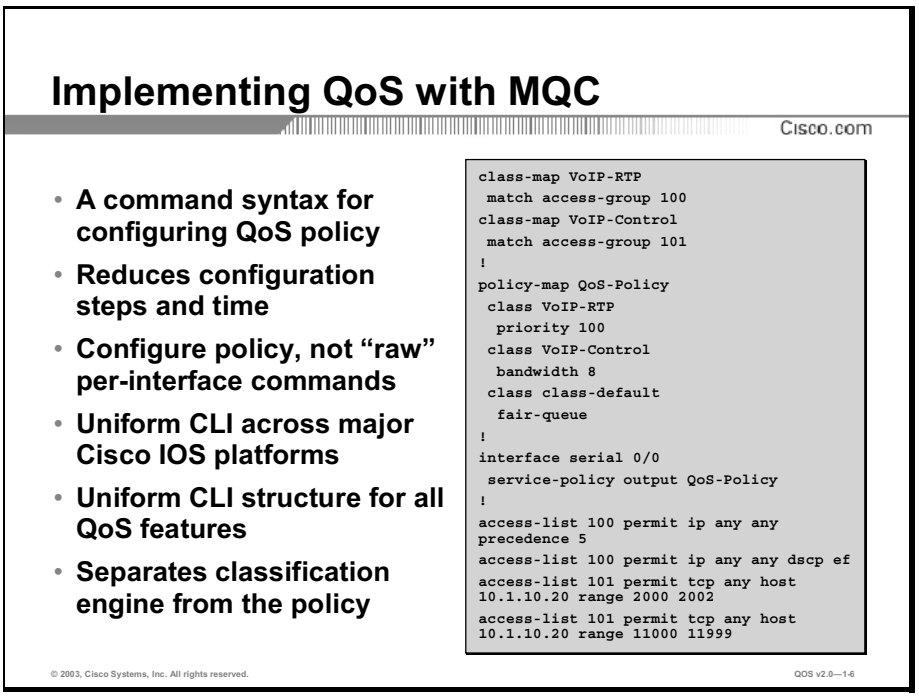

The MQC is a CLI structure that allows users to create traffic polices and then attach these polices to interfaces. A traffic policy contains one or more traffic classes and one or more QoS features. A traffic class is used to classify traffic; the QoS features in the traffic policy determine how to treat the classified traffic.

The MQC offers significant advantages over the legacy CLI method for implementing QoS. By using MQC, a network administrator can significantly reduce the time and effort it takes to configure QoS on a complex network. Rather than configuring "raw" CLI commands interface by interface, the administrator develops a uniform set of traffic classes and QoS policies, which can be applied on interfaces.

The use of the MQC allows the separation of traffic classification from the definition of QoS policy. This enables easier initial QoS implementation and maintenance as new traffic classes emerge and QoS policies for the network evolve.

AutoQoS

This topic describes the use of Cisco AutoQos for implementing QoS in a network.

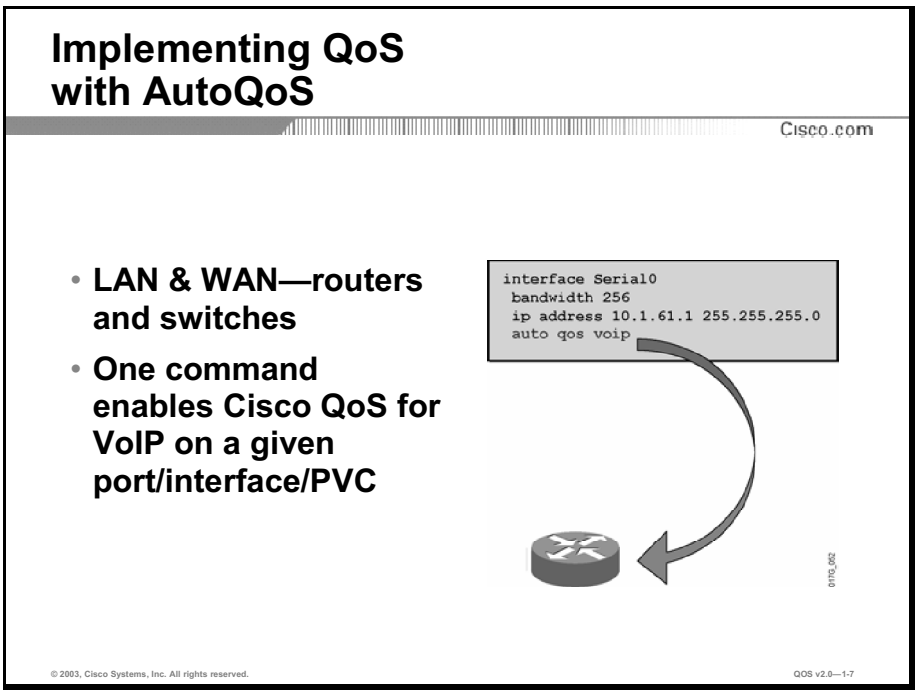

Using Cisco AutoQos, network administrators can implement the QoS features that are required for VoIP traffic without an in-depth knowledge of the following underlying technologies:

- PPP
- Frame Relay
- ATM
- Service policies
- Link efficiency mechanisms, such as link fragmentation and interleaving (LFI)

The AutoQoS VoIP feature simplifies QoS implementation and speeds up the provisioning of QoS technology over a Cisco network. It also reduces human error and lowers training costs. With the AutoQoS VoIP feature, one command (the **auto qos** command) enables QoS for VoIP traffic across every Cisco router and switch.

Network administrators can also use existing Cisco IOS commands to modify the configurations that are automatically generated by the AutoQoS VoIP feature in case the default AutoQoS configuration is not sufficient.

CiscoWorks QPM can be used in conjunction with the AutoQoS VoIP feature to provide a centralized, web-based tool to cost effectively manage and monitor network-wide QoS policies. The AutoQoS VoIP feature, together with CiscoWorks QPM, eases QoS implementation, provisioning, and management.

Note: Cisco AutoQoS was introduced in the 12.2(15)T Cisco IOS software release.

QoS Implementation Methods Compared

This topic compares the different methods for implementing QoS.

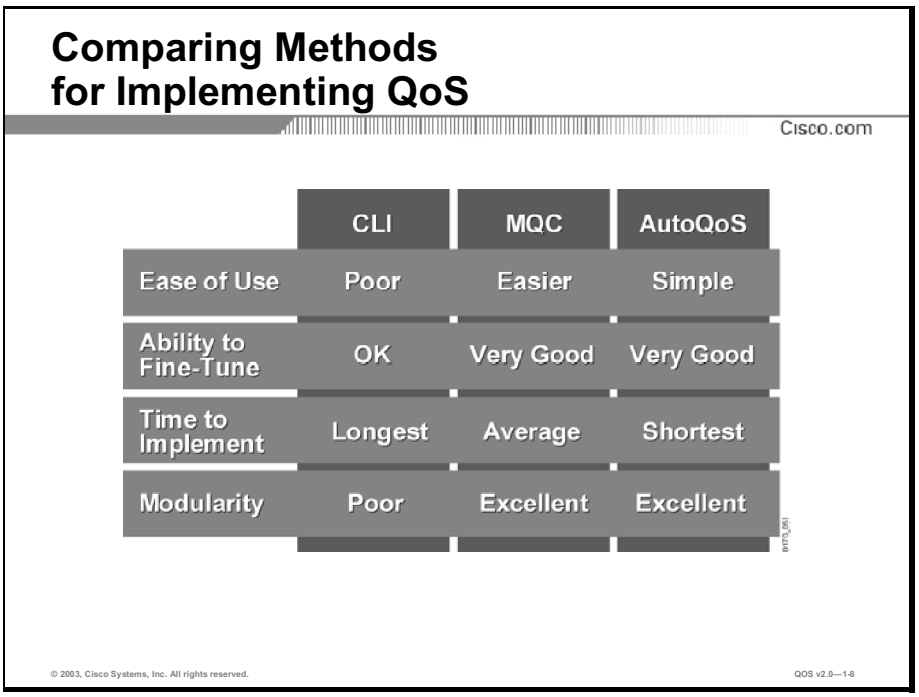

The three methods for configuring QoS on a network are the legacy CLI method, the MQC method, and Cisco AutoQoS.

Cisco recommends the use of the MQC and Cisco AutoQoS for implementing QoS.

While MQC is much easier to use than CLI, AutoQoS can simplify the configuration of QoS. As a result, the fastest implementation possible can usually be accomplished with AutoQoS.

MQC offers excellent modularity and the ability to fine-tune complex networks. AutoQoS offers the fastest way to implement QoS, but has limited fine-tuning capabilities. When an AutoQoS configuration has been generated, it is necessary to use CLI commands to fine-tune an AutoQoS configuration if necessary. (On most networks fine-tuning will not be necessary for AutoQoS.)

QoS Policy Manager

This topic describes the use of the QPM for managing QoS on a network.

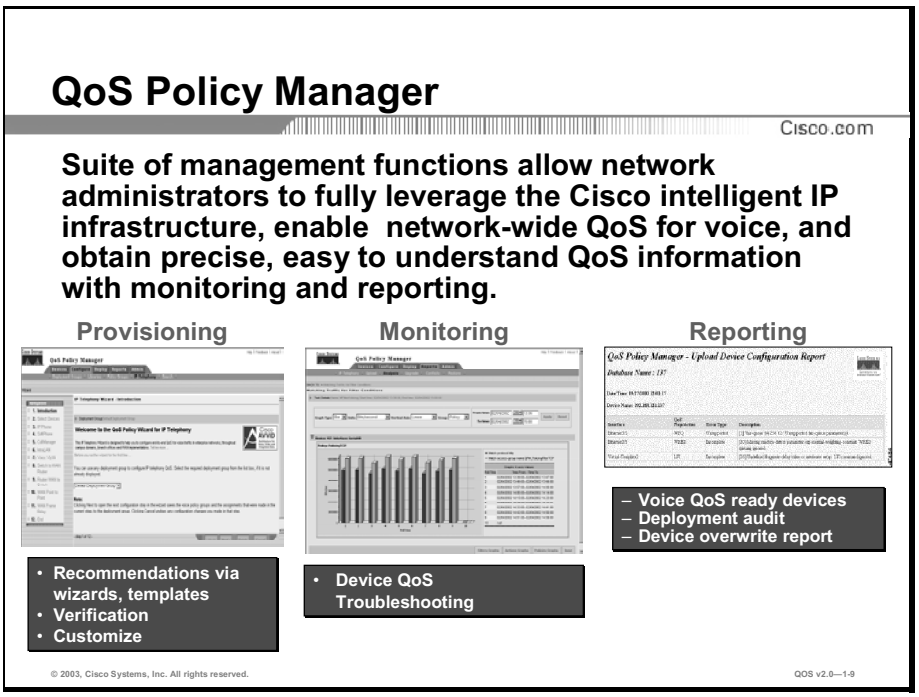

CiscoWorks QPM provides a scalable platform for defining, applying, and monitoring QoS policy on a system-wide basis for Cisco devices, including routers and switches.

QPM enables you to baseline profile network traffic, create QoS policies at an abstract level, control the deployment of policies, and then monitor QoS to verify intended results. As a centralized tool QPM is used to monitor and provision QoS for groups of interfaces and devices.

QPM provides a web-based intuitive user interface to define QoS policies and translates those policies into the device CLI commands.

QPM lets you analyze traffic throughput by application or service class. This analysis leverages that information to configure QoS policies to differentiate traffic and define the QoS functions that are applied to each type of traffic flow.

By simplifying QoS policy definition and deployment, QPM makes it easier for you to create and manage end-to-end differentiated services in your network, thus making more efficient and economical use of your existing network resources. For example, you can deploy policies that ensure that your mission-critical applications always get the bandwidth required to run your business.

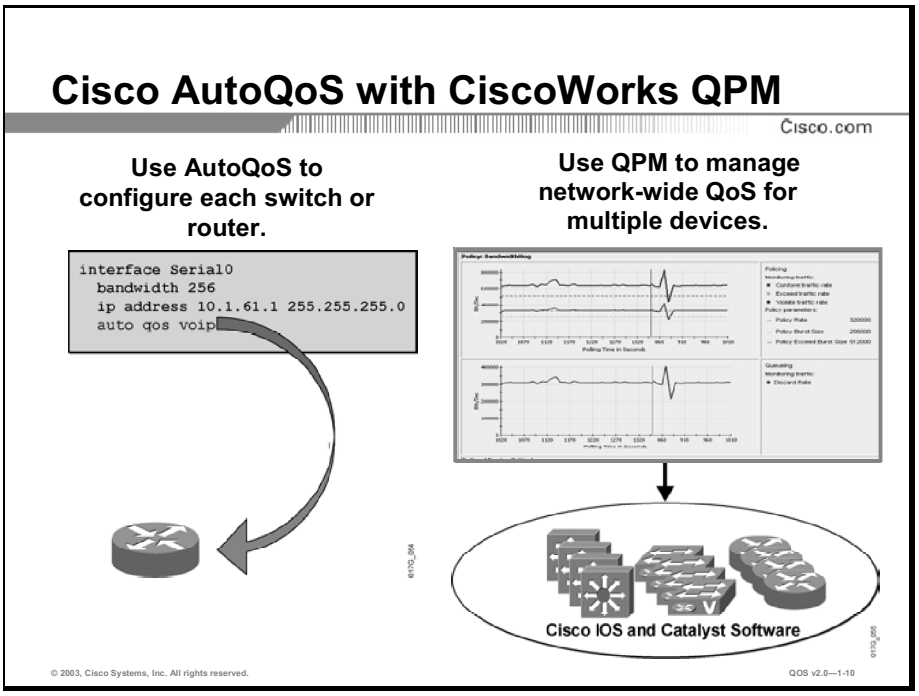

Customers can more easily provision and manage successful QoS deployments using Cisco AutoQoS together with QPM. Cisco AutoQoS provides QoS provisioning for individual routers and switches, simplifying deployment and reducing human error. CiscoWorks QPM provides centralized QoS design, administration, and traffic monitoring that scales to large QoS deployments.

Network Management MIBs for Monitoring QoS

This topic describes the key Management Information Bases (MIBs) that are used in managing QoS implementations.

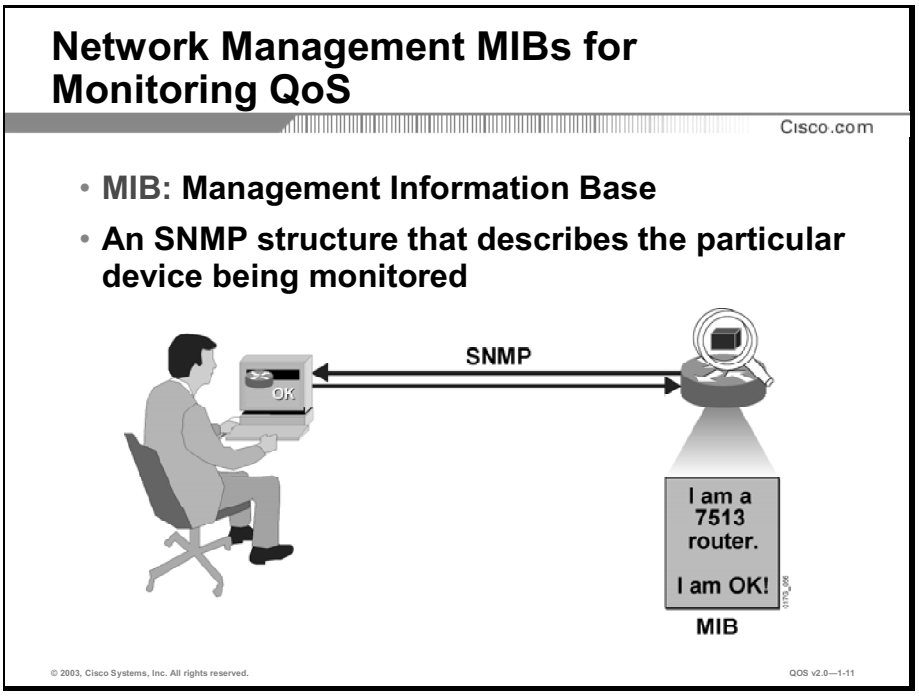

An MIB is a Simple Network Management Protocol (SNMP) structure that describes the particular device being monitored. Cisco provides many standards-based MIBs for use in monitoring the status of devices on a network.

Advanced network management products, such as CiscoWorks QPM, use these MIBs to generate statistics on the performance of the network. Specialized QoS MIBs enable QPM to graphically display key QoS information to aid in the management of QoS policies on the network.

Note: See the “Cisco Network Management Toolkit for MIBs” at http://www.cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml

MIBs for Managing QoS

This topic describes the key MIBs that are used for managing QoS policy on a network.

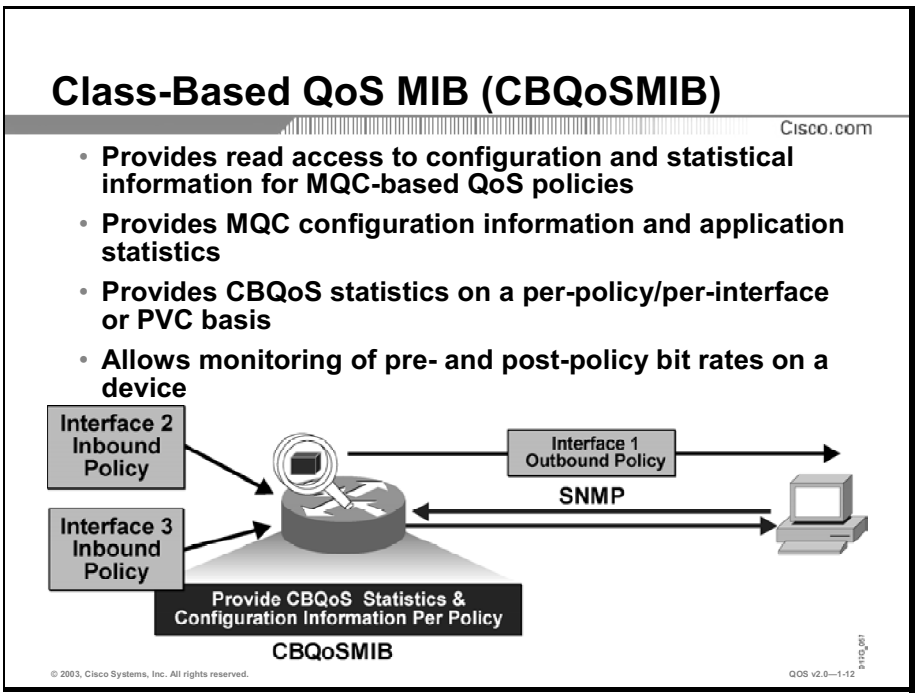

The Class-Based QoS MIB (CBQoSMIB) provides read access to QoS configurations. This MIB also provides QoS statistics information based on the MQC, including information regarding class map and policy map parameters.

This CBQoSMIB actually contains two MIBs: Cisco Class-Based QoS MIB and Cisco Class-Based QoS Capability MIB.

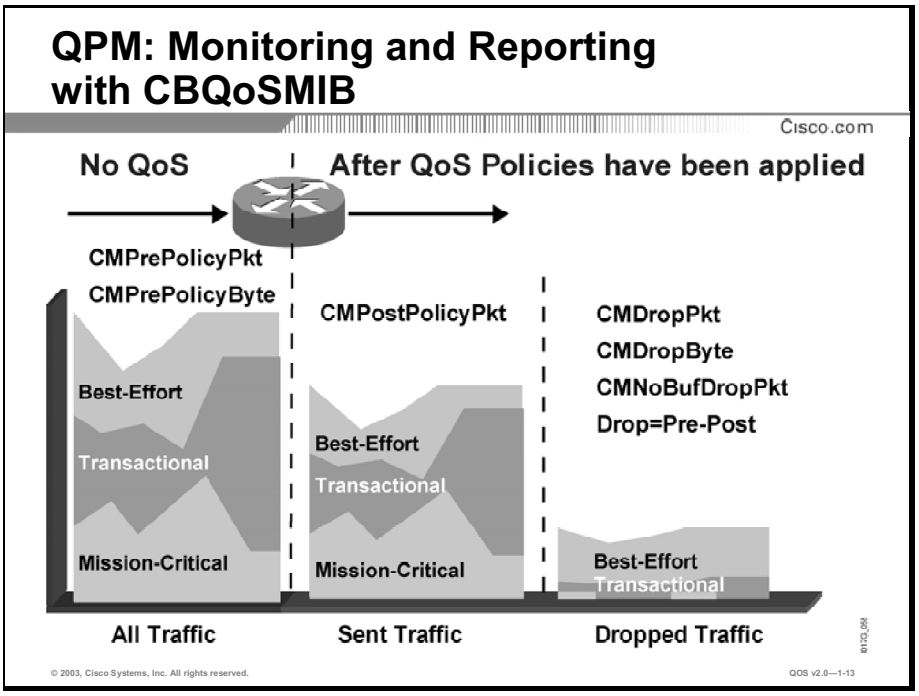

CiscoWorks QPM uses the information collected in the class-based MIB to build a number of reports showing the effect of QoS policies on the network.

These reports can graphically illustrate the overall input traffic flow divided by traffic class, the traffic that was actually sent, and the traffic that was dropped because of QoS policy enforcement.

In the Reports tab for QPM, under Historical Reports, it is possible to create graphs, some of which include:

- "Matching Traffic Per Class *Prior to QoS* Actions" graphs that display the traffic that matched each policy group filters, before any policy actions were performed
- "Matching Traffic Per Class *After QoS* Actions" graphs that display the traffic that matched each policy group filters and was transmitted (not dropped) by the configured QoS policies
- "Matching Traffic Per Class *Discarded* by QoS Drop Actions" graphs that display the traffic that matched each policy group filters and was dropped (not transmitted) by QoS policy drop actions

Cisco NBAR Protocol Discovery MIB

Cisco.com

- **NBAR protocol discovery statistics only available on the configured device**
- **Provides ability to retrieve statistics via SNMP into a central performance monitoring system**

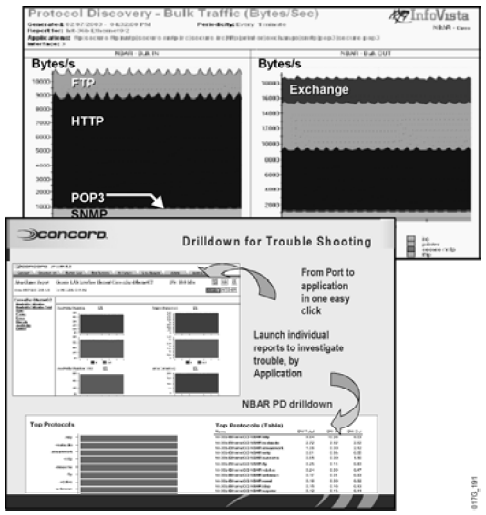

Another important MIB that is used for monitoring QoS is the Cisco NBAR Protocol Discovery MIB (CNPD MIB). Using the information collected by this MIB, it is possible to collect detailed protocol and application-level network utilization statistics.

Note:	CNPD MIB is discussed further in the module “Classification and Marking” when NBAR is explained.

Summary

This topic summarizes the key points discussed in this lesson.

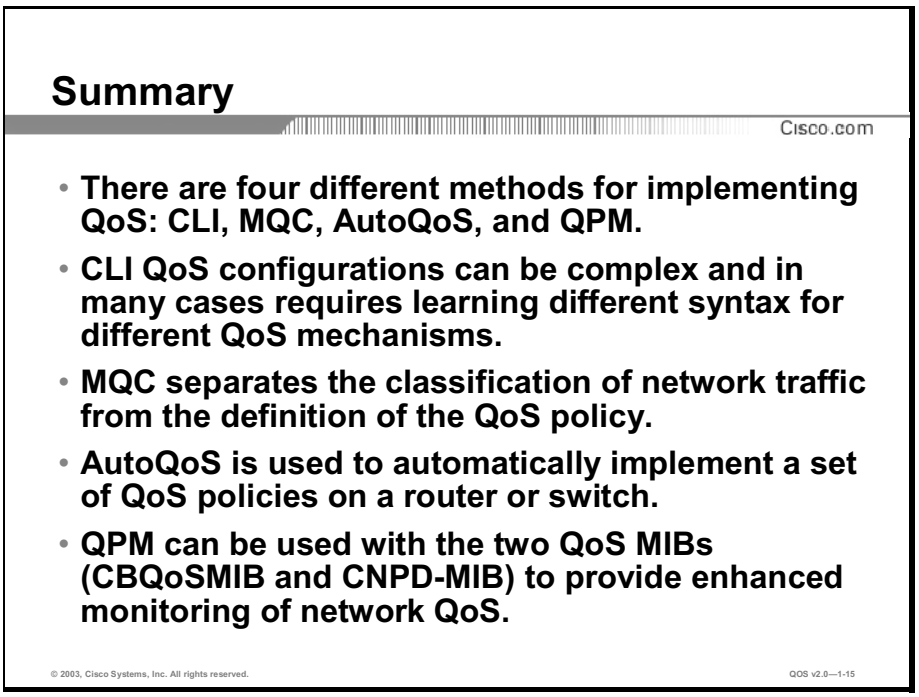

References

For additional information, refer to these resources:

- For more information on QPM, refer to “Introduction to QPM” at the following URL: http://www.cisco.com/en/US/products/sw/cscowork/ps2064/products_user_guide_chapter09186a00800e0a00.html
- For more information on Cisco MIBs, refer to “Cisco Network Management Toolkit for MIBs” at the following URL: http://www.cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) Which three methods are used to implement QoS policy? (Choose three.)

A) AutoQoS

B) QoS CLI Manager

C) QoS Policy Manager

D) Command-Line Interface

Q2) Which of the following is a major advantage of MQC?

A) capability to generate QoS CLI code automatically

B) ability to separate classification from policy definition

C) capability to do fine-tuning with “raw” CLI commands

D) ability to automatically recognize new classes of traffic

Q3) Which three of the following are advantages of Cisco AutoQoS over other methods for implementing QoS? (Choose three.)

A) reduces human error

B) lowers training costs

C) increases consistency

D) works for all situations

Q4) Which three of the following are features of QPM? (Choose three.)

A) baseline profile network traffic

B) control deployment of policies

C) create QoS policies at an abstract level

D) auto alert users of QoS policy violation

Q5) MIB is an acronym for_____.

A) Management Interrupt Block

B) Management Information Base

C) Management Information Block

D) Management Implementation Block

Q6) Using the Class-Based QoS MIB with QPM, it is possible to create which two of the following? (Choose two.)

A) AutoQoS Macro Efficiency Analysis

B) matching Traffic Per Class *After QoS* Actions

C) matching Traffic Per Class *Prior to QoS* Actions

D) matching Traffic by AutoQoS Generated Classes

Quiz Answer Key

Q1) A, C, D

Relates to: Methods for Implementing QoS Policy

Q2) B

Relates to: Modular QoS CLI

Q3) A, B, C

Relates to: AutoQoS

Q4) A, B, C

Relates to: QoS Policy Manager

Q5) B

Relates to: Network Management MIBs for Monitoring QoS

Q6) B, C

Relates to: MIBs for Managing QoS

Module Assessment

Overview

Use this assessment to test what you learned in this module. The correct answers and solutions are found in the Module Assessment Answer Key.

Quiz: Introduction to IP QoS

Complete the Quiz to assess what you have learned in the module.

Objectives

This activity tests your knowledge on how to meet these objectives:

- Identify problems that could lead to poor QoS and explain how the problems might be resolved
- Define the term QoS and identify and explain the key steps to implementing QoS on a converged network
- List and describe methods for implementing QoS

Instructions

Complete these steps:

Step 1 Answer all questions in this quiz by selecting the best answer(s) to each question.

Step 2 Verify your results against the answer key located at the end of this section.

Step 3 Review the topics in this module that relate to the questions that you answered incorrectly.

Q1) Which of the following terms is used to describe the time it takes to actually transmit a packet on a link (put bits on the wire)?

A) encoding delay

B) processing delay

C) serialization delay

D) transmission delay

Q2) What is the “best” solution for reducing delay on a link?

A) compress data and headers

B) drop low priority packets early

C) increase the bandwidth of the link

D) incorporate advanced queuing technologies

Q3) Which three of the following are characteristics of converged network traffic? (Choose three.)

A) constant small packet flow

B) time-sensitive packets

C) brief outages are unacceptable

D) bursty small packet flow

Q4) How much one-way delay can a voice packet tolerate?

A) 15 ms
B) 150 ms
C) 300 ms
D) 200 ms

Q5) Which transport layer protocol is used for voice traffic?

A) UDP
B) TCP
C) XNS
D) HTTP

Q6) Which three of the following represent components in the definition of a QoS policy? (Choose three.)

A) user validated
B) network-wide
C) specific levels of quality of service
D) different classes of network traffic

Q7) How are QoS implementations generated using AutoQoS fine-tuned?

A) command-line interface
B) Modular QoS CLI
C) QoS AutoTune
D) QoS Policy Manager

Q8) Which three of the following are advantages of using MQC? (Choose three.)

A) reduction in time to configure a complex policy
B) ability to apply one policy to multiple interfaces
C) separation of classification from policy definition
D) automatic generation of CLI commands from MQC macros

Q9) Which QoS implementation method has the quickest implementation time for simple networks?

A) CLI

B) MQC

C) AutoQoS

D) AutoTuner

Q10) Which two of the following are MIBS specifically designed for managing QoS in a network? (Choose two.)

A) Modular QoS MIB

B) class-based QoS MIB

C) QoS Policy Manager MIB

D) Cisco NBAR Protocol Discover MIB

Scoring

You have successfully completed the quiz for this lesson when you earn a score of 80 percent or better.

Module Assessment Answer Key

Q1) C

Relates to: The Need for QoS

Q2) C

Relates to: The Need for QoS

Q3) A, B, C

Relates to: The Need for QoS

Q4) B

Relates to: Understanding QoS

Q5) A

Relates to: Understanding QoS

Q6) B, C, D

Relates to: Understanding QoS

Q7) A

Relates to: Implementing IP QoS

Q8) A, B, C

Relates to: Implementing IP QoS

Q9) C

Relates to: Implementing IP QoS

Q10) B, D

Relates to: Implementing IP QoS

Module Summary

This topic summarizes the key points discussed in this module.

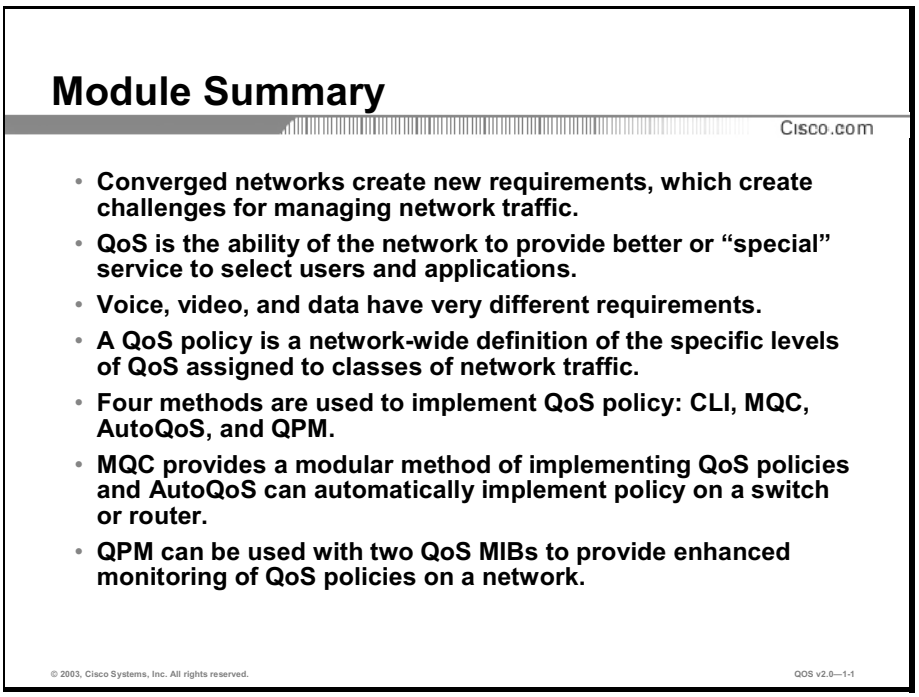

Voice and video traffic present new challenges to networking. QoS is the network glue that makes it possible to incorporate voice and video traffic into a traditional networking environment.

The MQC, Cisco AutoQoS, and QPM offer much more cost-effective and simple ways to configure and manage a QoS-enabled network.

Module 2

The Building Blocks of IP QoS

Overview

Quality of service (QoS) and its implementations are necessarily complex. The complex requirements of different applications in a converged network can create many challenges for network administrators and architects. As technology evolves over time, different approaches to solving the problems of providing service quality to network applications are introduced. Many of these QoS "building blocks" or "features" operate at different parts of a network to create an end-to-end QoS system. Managing how these "building blocks" are assembled and how different QoS "features" are used can be a difficult task. In response to these difficulties, three different implementation models for QoS have been developed.

This module discusses the different implementation models of QoS and describes how the different "building blocks" of QoS integrate into each of them. This module also discusses the different QoS features and where they are typically implemented within a network. Because the end result of QoS implementations is to effect application traffic traversing over a QoS enabled network, this module also describes the effects that different QoS features have on network traffic.

Module Objectives

Upon completing this module, you will be able to identify and describe different models used for ensuring QoS in a network and explain key IP QoS mechanisms that are used to implement the models.

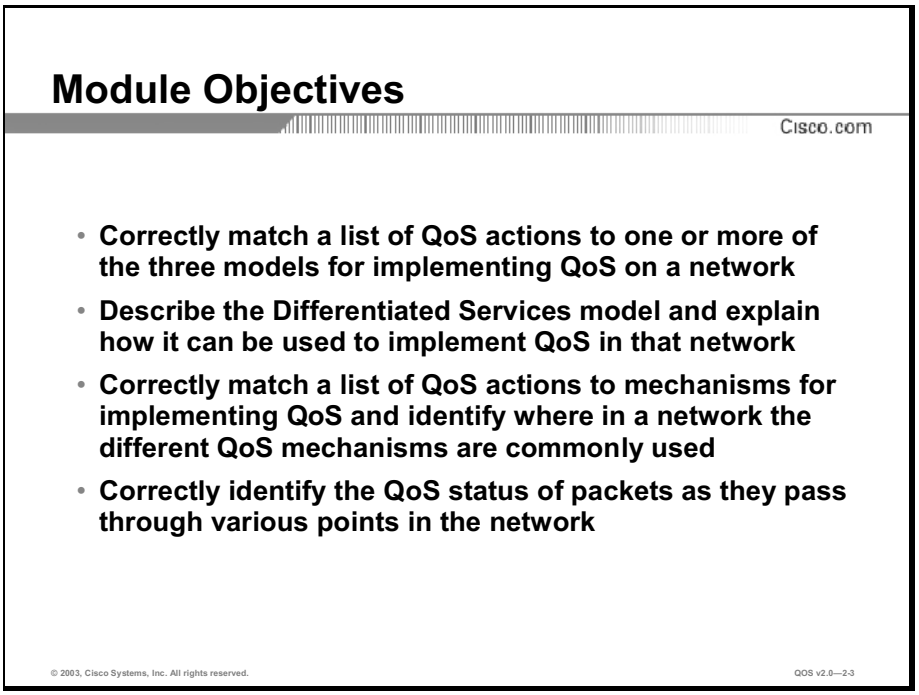

Module Outline

The outline lists the components of this module.

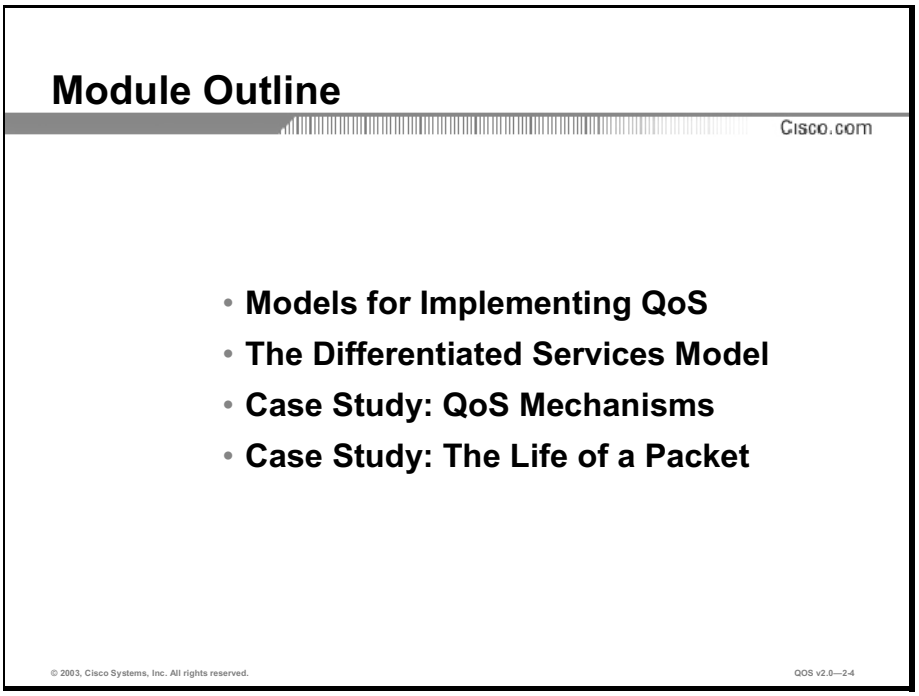

Models for Implementing QoS

Overview

Three different models exist for implementing QoS on a network. The Best-Effort model was designed for best-effort, no-guarantee delivery of packets. This model is still predominant in the Internet today. The Integrated Services (IntServ) model was introduced to supplement the best-effort delivery by setting aside some bandwidth for applications that require bandwidth and delay guarantees. The IntServ model expects applications to signal their requirements to the network. The Differentiated Services (DiffServ) model was added to provide greater scalability in providing QoS to IP packets. The main difference between the IntServ and DiffServ models is that in the DiffServ model the network recognizes packets (no signaling is needed) and provides the appropriate services to them. IP networks today can use all three models at the same time.

Relevance

To select the most appropriate method or methods for implementing QoS on a network, it is vital to understand the primary methods available.

Objectives

Upon completing this lesson, you will be able to correctly match QoS actions to one or more models for implementing QoS on a network. This includes being able to meet these objectives:

- List the models for providing QoS on a network
- Explain the key features of the Best-Effort model for QoS
- Explain the key features of the IntServ model for QoS
- Explain the key features of the DiffServ model for QoS

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Basic knowledge of internetworking with TCP/IP concepts

Outline

The outline lists the topics included in this lesson.

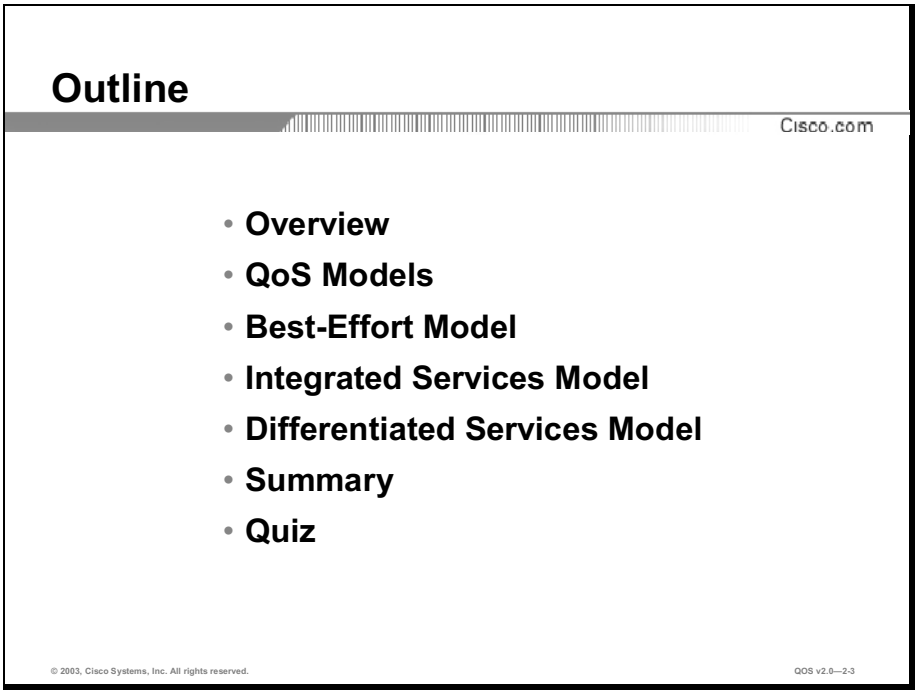

QoS Models

This topic lists the models for providing QoS on a network.

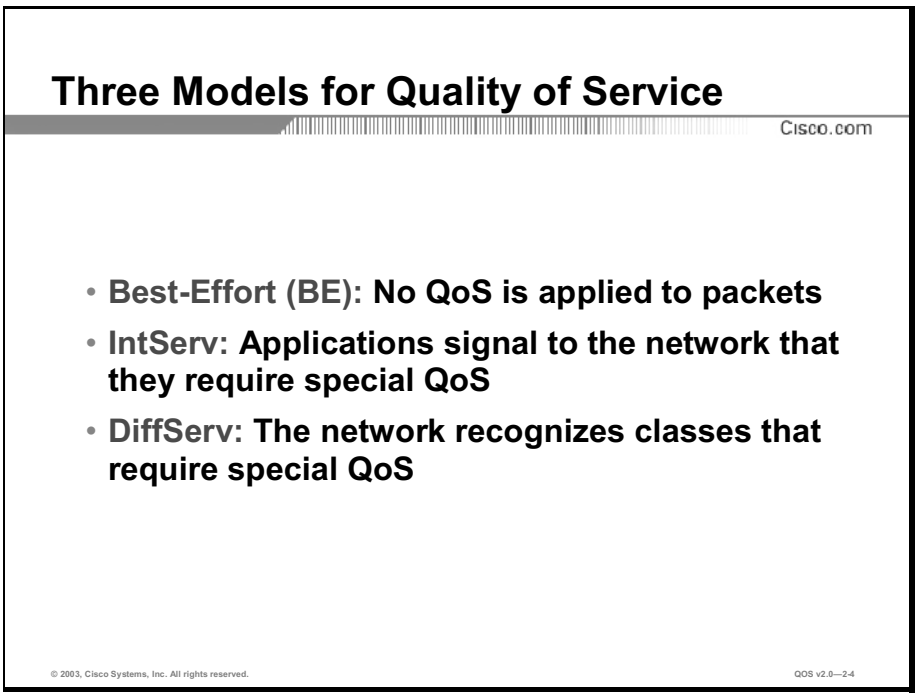

Three different models exist for implementing QoS in a network:

- With the Best-Effort model, QoS is not applied to packets. If it is not important when or how packets arrive, the Best-Effort model is appropriate.
- The IntServ model can provide very high QoS to IP packets. Essentially, applications signal to the network that they will require special QoS for a period of time and bandwidth is reserved. With the IntServ model, packet delivery is guaranteed. However, the use of the IntServ model can severely limit the scalability of a network.
- The DiffServ model provides the greatest scalability and flexibility in implementing QoS in a network. Network devices recognize traffic classes and provide different levels of QoS to different traffic classes.

Best-Effort Model

This topic explains the key features of the Best-Effort model for QoS.

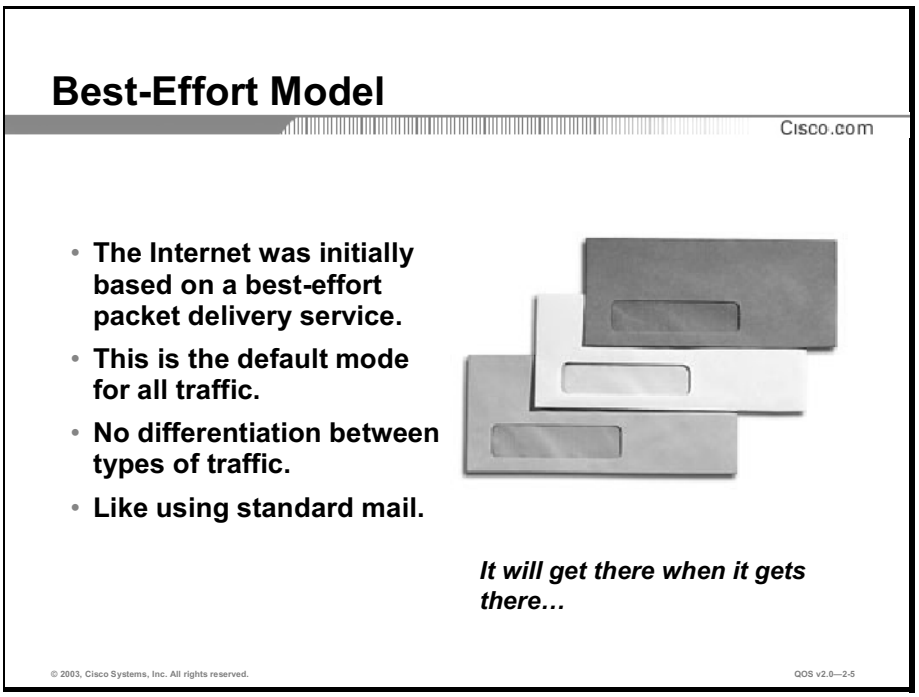

The Internet was designed for best-effort, no-guarantee delivery of packets. This behavior is still predominant on the Internet today.

If QoS policies are not implemented, traffic is forwarded using the Best-Effort model. All network packets are treated exactly the same—an emergency voice message is treated exactly like a digital photograph attached to an e-mail. Without QoS implemented, the network cannot tell the difference and, as a result, cannot treat packets preferentially.

When you drop a letter in standard postal mail, you are using a Best-Effort model. Your letter will be treated exactly the same as every other letter—*"it will get there when it gets there."* With the Best-Effort model the letter may actually never arrive and, unless you have a separate notification arrangement with the letter recipient, you may never know if the letter does not arrive.

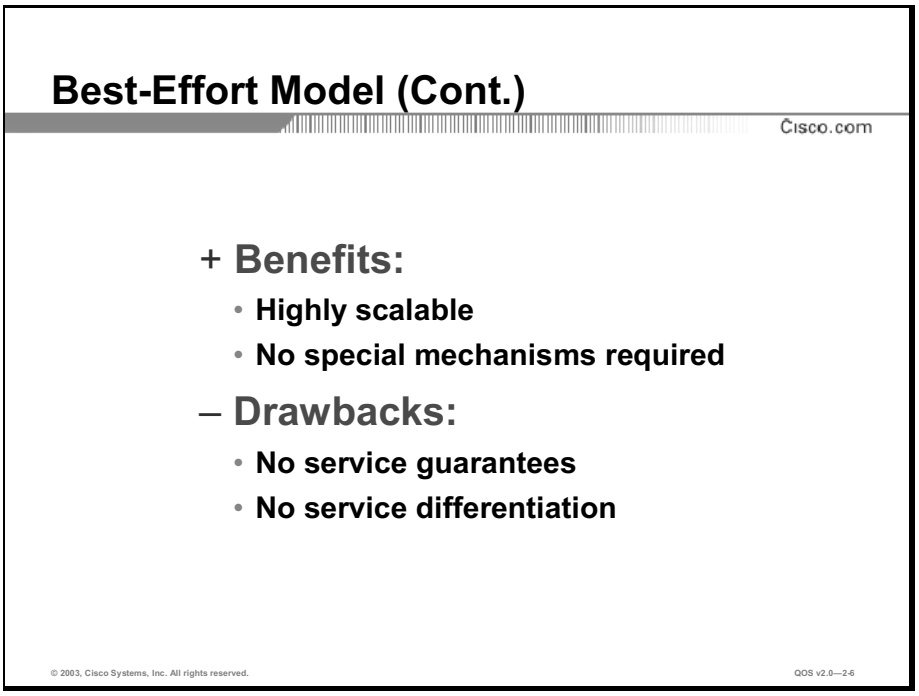

The Best-Effort model does have two significant benefits:

- It has virtually unlimited scalability. The only way to reach scalability limits is to reach bandwidth limits; then, everything becomes equally delayed.
- No special QoS mechanisms need be employed to use the Best-Effort model. It is, as a result, the easiest and quickest to deploy.

The Best-Effort model also has obvious drawbacks:

- Nothing is guaranteed. Packets will arrive whenever they can, in any order possible, if they arrive at all.
- Packets are not given preferential treatment. Critical data is treated the same as casual e-mail.

Integrated Services Model

This topic explains the key features of the IntServ model for QoS.

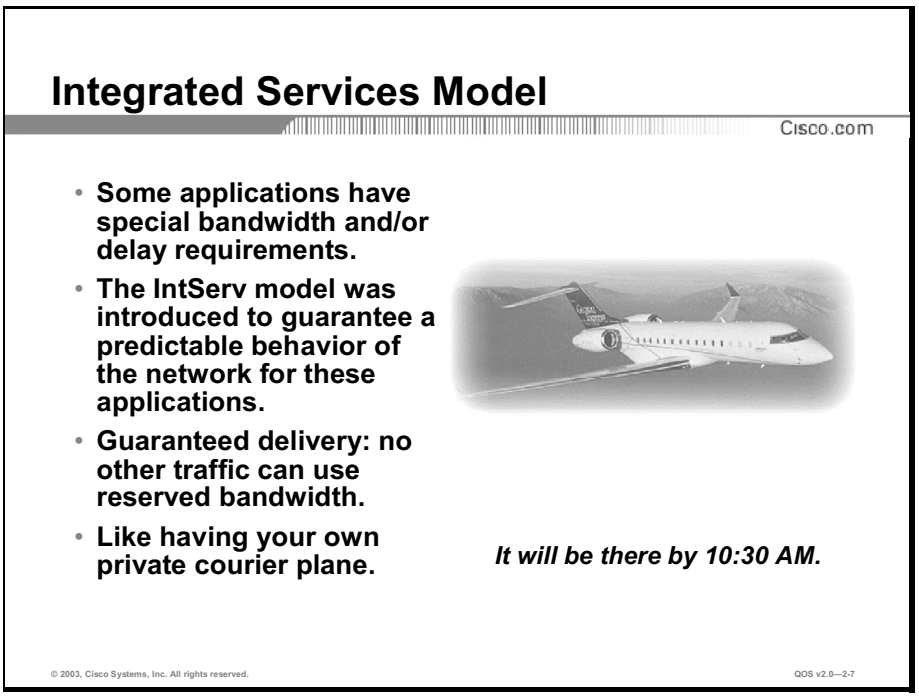

Some applications, such as high-resolution video, require consistent, dedicated bandwidth to provide sufficient quality for viewers. The IntServ model was introduced to guarantee predictable network behavior for these applications.

Because the IntServ model reserves bandwidth throughout a network, no other traffic can use the reserved bandwidth. Bandwidth unused, but reserved, is wasted.

This is similar to a concept known as “Hard QoS.” With Hard QoS traffic characteristics such as bandwidth, delay, and packet-loss rates, are guaranteed end-to-end. This guarantee ensures both predictable and guaranteed service levels for mission-critical applications. Guaranteed traffic cannot be impacted when guarantees are made, regardless of additional network traffic. Hard QoS is accomplished by negotiating specific QoS requirements upon connection establishment and by using Call Admission Controls (CACs) to ensure that no new traffic will violate the guarantee. Such guarantees require an end-to-end QoS approach with both complexity and scalability limitations. Large network environments that contain heavy traffic loads will be extremely challenged to track QoS guarantees for hundreds of thousands of signaled flows.

Using the IntServ model is like having a private courier airplane or truck dedicated to the delivery of your traffic. It ensures quality and delivery, is expensive, and is not scalable.

Integrated Services Model (Cont.)

Cisco.com

- **Provides multiple service levels**
- **Requests specific kind of service from the network before sending data**
- **Uses RSVP to reserve network resources**
- **Uses intelligent queuing mechanisms**
- **End-to-end**

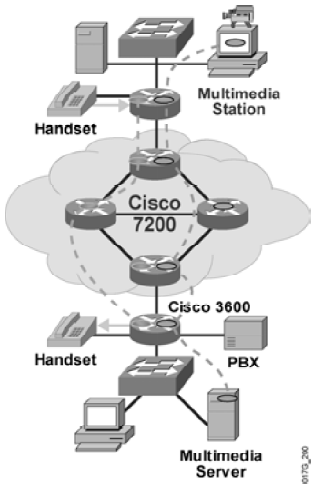

QOS v2.0—2-8

IntServ is a multiple-service model that can accommodate multiple QoS requirements. The IntServ model inherits the connection-oriented approach from telephony network design. Every individual communication must explicitly specify its traffic descriptor as well as requested resources to the network. The edge router performs admission control to ensure that available resources are sufficient in the network. The IntServ standard assumes that routers along a path set and maintain state for each individual communication.

The role of Resource Reservation Protocol (RSVP) in the Cisco QoS architecture is to provide resource admission control for Voice over IP (VoIP) networks. If resources are available, RSVP accepts a reservation and installs a traffic classifier in the QoS forwarding path. The traffic classifier tells the QoS forwarding path how to classify packets from a particular flow and what forwarding treatment to provide.

In this model the application requests a specific kind of service from the network before sending data. The application informs the network of its traffic profile and requests a particular kind of service that can encompass its bandwidth and delay requirements. The application is expected to send data only *after* it gets a confirmation from the network. It is also expected to send data that lies within its described traffic profile.

The network performs admission control based on information from the application and available network resources. It commits to meeting the QoS requirements of the application as long as the traffic remains within the profile specifications. The network fulfills its commitment by maintaining per-flow state, and then performing packet classification, policing, and intelligent queuing based on that state.

The QoS feature set in Cisco IOS software includes these features that provide controlled-load service:

- RSVP can be used by applications to signal their QoS requirements to the router.
- Intelligent queuing mechanisms can be used with RSVP to provide these QoS service levels:
  - Guaranteed-rate: Allows applications to reserve bandwidth to meet their requirements. For example, a VoIP application can reserve 32 Mbps end-to-end using this type of service. Cisco IOS QoS uses low latency queuing (LLQ) with RSVP to provide this type of service.
  - Controlled-load: Allows applications to have low delay and high throughput, even during times of congestion. For example, adaptive real-time applications such as the playback of a recorded conference can use this service. Cisco IOS QoS uses RSVP with weighted random early detection (WRED) to provide this type of service.

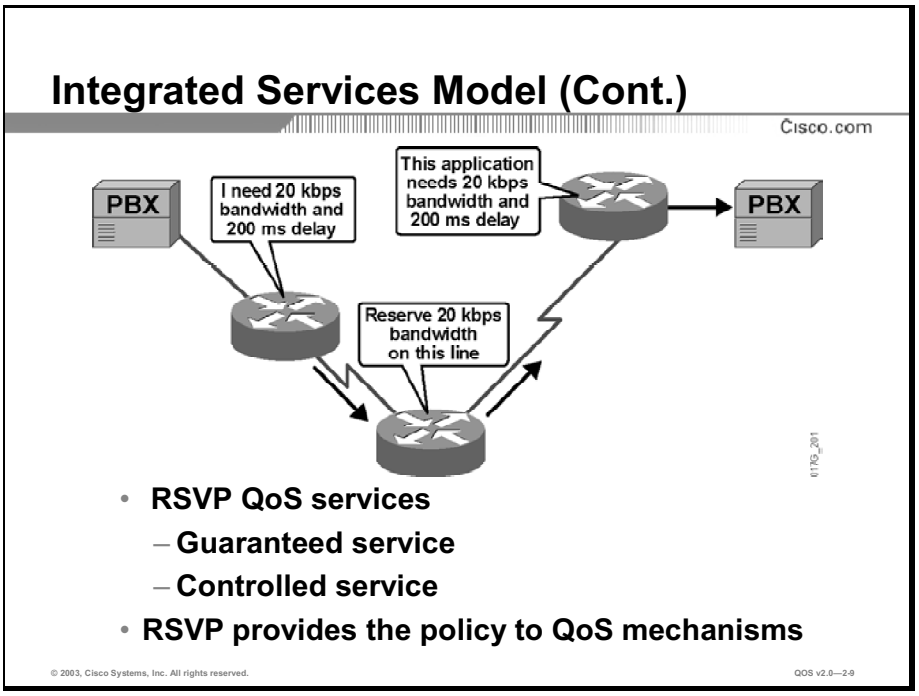

RSVP is an IP service that allows end systems or hosts on either side of a router network to establish a reserved-bandwidth path between them to predetermine and ensure QoS for their data transmission. RSVP is currently the only standard signaling protocol designed to guarantee network bandwidth from end to end for IP networks.

RSVP is an Internet Engineering Task Force (IETF) standard (RFC 2205) protocol for allowing an application to dynamically reserve network bandwidth. RSVP enables applications to request a specific QoS for a data flow (shown in the graphic). Cisco implementation also allows RSVP to be initiated within the network, using configured proxy RSVP. Network managers can take advantage of RSVP benefits in the network, even for non-RSVP-enabled applications and hosts.

Hosts and routers use RSVP to deliver QoS requests to the routers along the paths of the data stream. Hosts and routers also use RSVP to maintain router and host state to provide the requested service, usually bandwidth and latency. RSVP uses a mean data rate; that is, the largest amount of data the router will keep in queue and the minimum QoS used to determine bandwidth reservation.

LLQ or WRED act as the workhorses for RSVP, setting up the packet classification and scheduling that is required for the reserved flows. Using LLQ, RSVP can deliver an IntServ guaranteed service. Using WRED, it can deliver a controlled-load service. RSVP can be deployed in existing networks with a software upgrade.

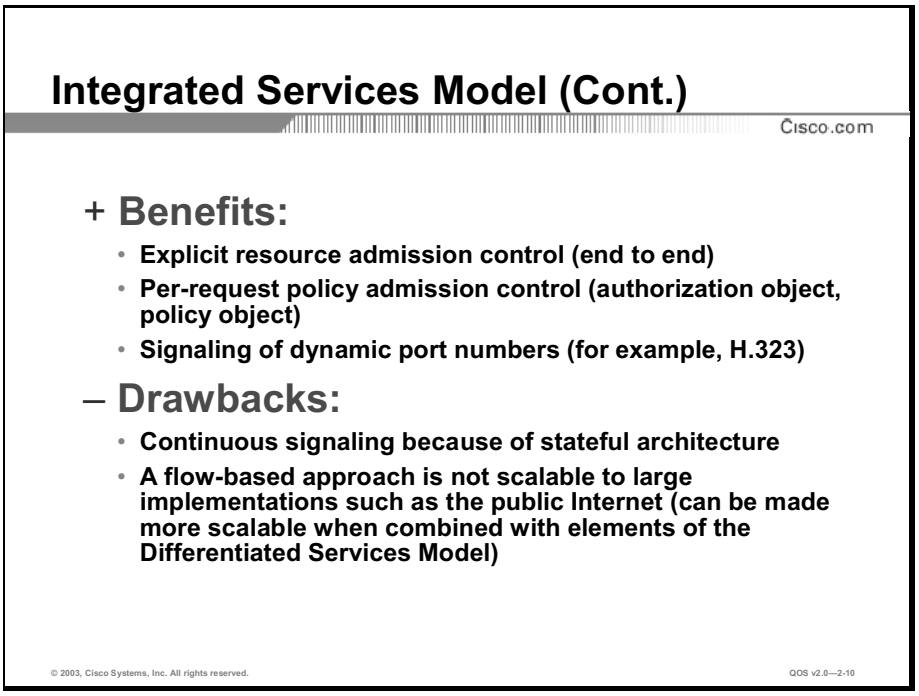

The main benefits of the IntServ model and RSVP are:

- It signals QoS requests per individual flow. The network can then provide guarantees to these individual flows. The problem with this is that it does not scale to large networks because of the large numbers of concurrent RSVP flows.
- It informs network devices of flow parameters (IP addresses and port numbers). Some applications use dynamic port numbers, which can be difficult for network devices to recognize. Network-based application recognition (NBAR) is a mechanism that has been introduced to supplement RSVP for applications that use dynamic port numbers but do not use RSVP.
- It supports admission control that allows a network to reject (or downgrade) new RSVP sessions if one of the interfaces in the path has reached the limit (that is, all reservable bandwidth is booked).

The main drawbacks of the IntServ model and RSVP are:

- Continuous signaling because of the stateful RSVP operation.
- RSVP is not scalable to large networks where per-flow guarantees would have to be made to thousands of flows.

Differentiated Services Model

This topic explains the key features of the DiffServ model for QoS.

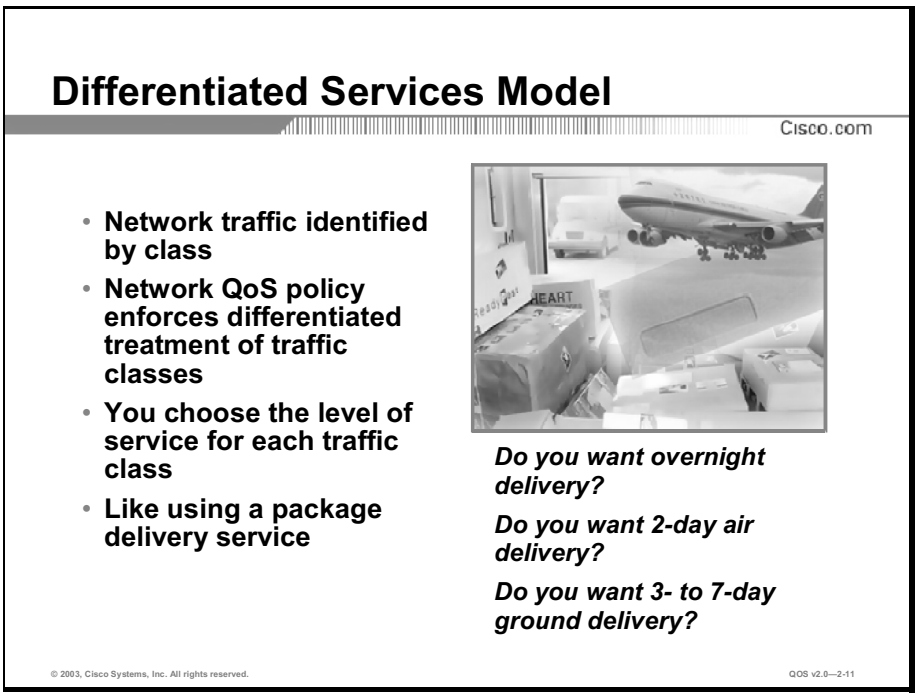

The DiffServ model was designed to overcome the limitations of both the Best-Effort and IntServ models. The DiffServ model can provide an "almost guaranteed" QoS while still being cost-effective and scalable.

This is similar to a concept known as "Soft QoS." With Soft QoS, QoS mechanisms are used without prior signaling. In addition, QoS characteristics (bandwidth and delay, for example), are managed on a hop-by-hop basis by policies that are established independently at each intermediate device in the network. The soft QoS approach is not considered an end-to-end QoS strategy because end-to-end guarantees cannot be enforced. However, soft QoS is a more scalable approach to implementing QoS than hard QoS because many (hundreds or potentially thousands) of applications can be mapped into a small set of classes upon which similar sets of QoS behaviors are implemented. Although QoS mechanisms in this approach are enforced and applied on a hop-by-hop basis, uniformly applying global meaning to each traffic class provides both flexibility and scalability.

With DiffServ, network traffic is divided into classes based on business requirements. Each of the classes can then be assigned a different level of service. As the packets traverse a network, each of the network devices identifies the packet class and services the packet according to that class.

It is possible to choose many levels of service with DiffServ. For example, voice traffic from IP Phones is usually given preferential treatment over all other application traffic. E-mail is generally given “best-effort” service. And nonbusiness traffic can either be given very poor service or blocked entirely.

DiffServ works like a package delivery service. You request (and pay for) a level of service when you send your package. Throughout the package network, the level of service is recognized and your package is given either preferential or normal service, depending on what you requested.

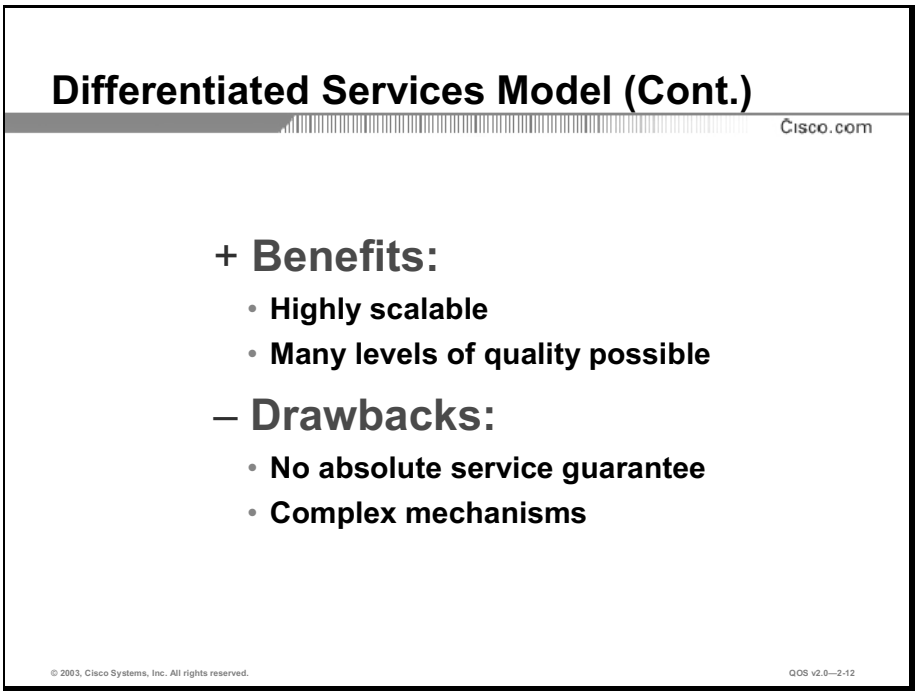

The DiffServ model has two key benefits:

- It is highly scalable.
- It provides many different levels of quality.

The DiffServ model also has drawbacks:

- No absolute guarantee of service quality can be made.
- It requires a set of complex mechanisms to work in concert throughout the network.

Summary

This topic summarizes the key points discussed in this lesson.

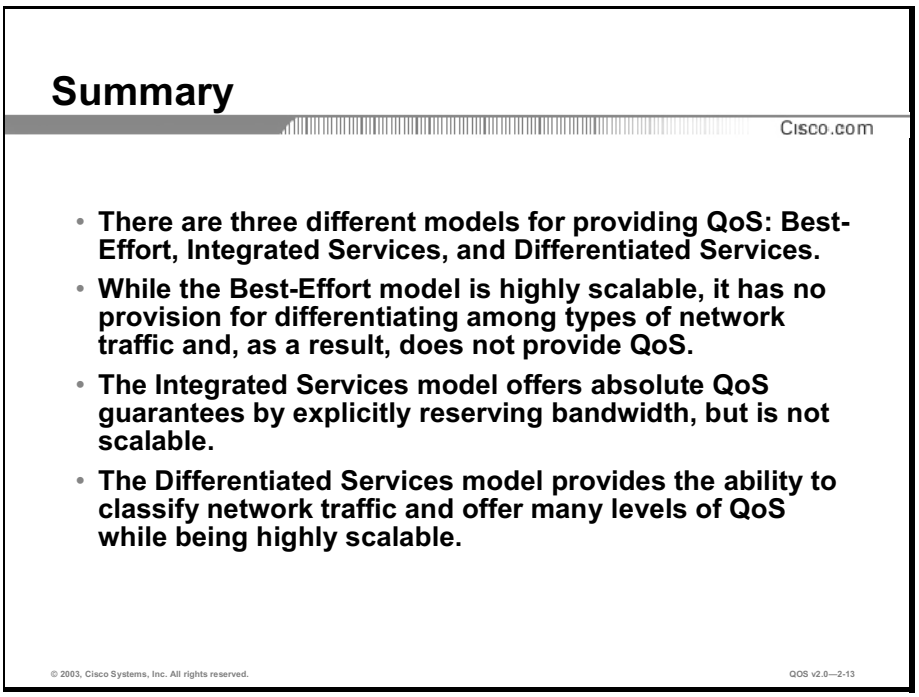

References

For additional information, refer to these resources:

- To learn more about the Integrated Services model, refer to “Integrated Services in the Internet Architecture: an Overview” at the following URL: http://www.ietf.org/rfc/rfc1633.txt
- To learn more about RSVP, refer to RFC 2210, “The Use of RSVP with IETF Integrated Services” at the following URL: http://www.ietf.org/rfc/rfc2210.txt
- To learn more about the Differentiated Services model, refer to RFC 2475 “An Architecture for Differentiated Services” at the following URL: http://www.ietf.org/rfc/rfc2475.txt

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) Which one of the models for implementing QoS requires that applications signal their special QoS requirements thereby reserving bandwidth?

A) Integrated Services

B) Best-Effort

C) Differentiated Services

D) Quantitative Services

Q2) Which one of the models for implementing QoS offers no guarantee of packet delivery?

A) Best-Effort

B) Integrated Services

C) Differentiated Services

D) Quantitative Services

Q3) Which of the models for implementing QoS was introduced because certain applications have special bandwidth and/or delay requirements?

A) Best-Effort

B) Integrated Services

C) Differentiated Services

D) Quantitative Services

Q4) Which QoS mechanism does Cisco IOS software rely upon for providing Integrated Services?

A) Low Latency Queuing (LLQ)

B) Generic Traffic Shaping (GTS)

C) Real Time Protocol (RTP)

D) Resource Reservation Protocol (RSVP)

Q5) Which two of the following represent benefits of the Integrated Services model? (Choose two.)

A) dynamic port number signaling

B) explicit resource admission control

C) a highly scalable QoS implementation

D) continuous signaling because of a stateless architecture

Q6) Which two of the following represent benefits of the Differentiated Services model? (Choose two.)

A) highly scalable

B) simple QoS mechanisms

C) absolute service guarantee

D) many levels of quality possible

Quiz Answer Key

Q1) A

Relates to: QoS Models

Q2) A

Relates to: Best-Effort Model

Q3) B

Relates to: Integrated Services Model

Q4) D

Relates to: Integrated Services Model

Q5) A, B

Relates to: Integrated Services Model

Q6) A, D

Relates to: Differentiated Services Model

The Differentiated Services Model

Overview

DiffServ is a multiple-service model designed to satisfy various QoS requirements. With DiffServ, the network tries to deliver a particular kind of service that is based on the QoS specified by each packet. This specification can occur in different ways; for example, using the DiffServ code point (DSCP) in IP packets or source and destination addresses. The network uses the QoS specification of each packet to classify, shape, and police traffic and to perform intelligent queuing.

Relevance

The DiffServ model is the primary model used to implement QoS in IP networks.

Objectives

Upon completing this lesson, you will be able to describe the Differentiated Services model and explain how it can be used to implement QoS in that network. This includes being able to meet these objectives:

- Explain the purpose and function of the DiffServ model
- Describe the basic format of and explain the purpose of the DSCP field in the IP header
- Define and explain the different per-hop behaviors used in DSCP
- Explain the interoperability between DSCP-based and IP precedence-based devices in a network

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Basic knowledge of internetworking with TCP/IP concepts

Outline

The outline lists the topics included in this lesson.

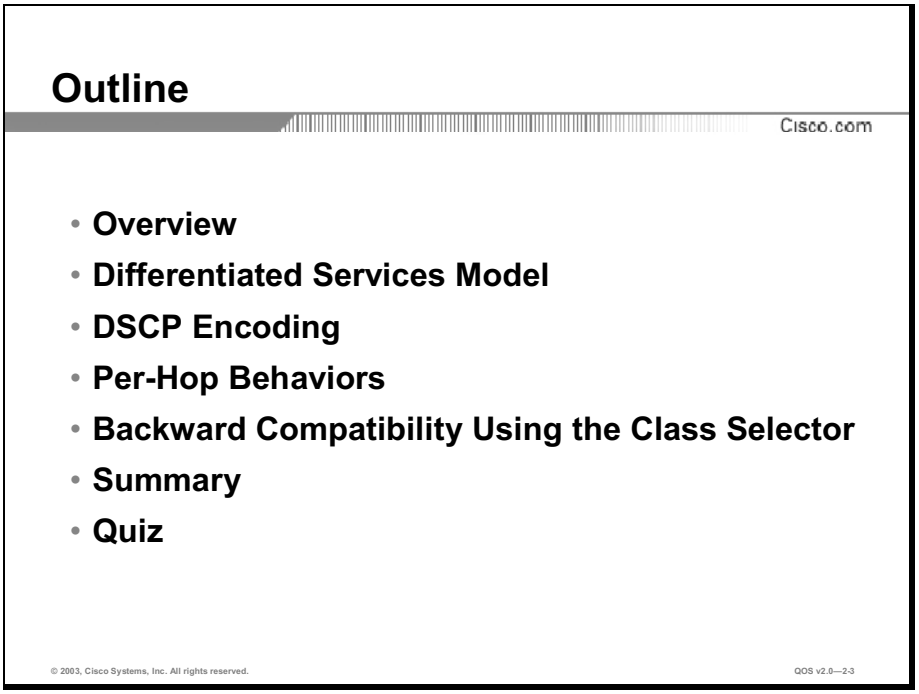

Differentiated Services Model

This topic explains the purpose and function of the DiffServ model.

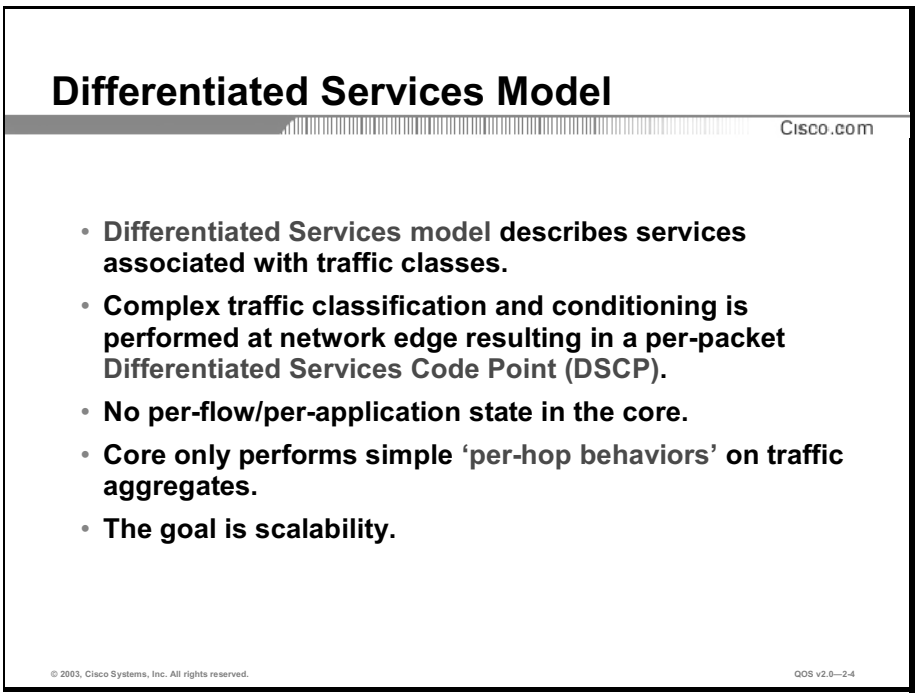

The DiffServ architecture is based on a simple model where traffic entering a network is classified and possibly conditioned at the boundaries of the network. The traffic class is then identified with a DSCP or bit marking in the IP header.

DSCP values are used to mark packets to select a per-hop behavior. Within the core of the network, packets are forwarded according to the per-hop behavior that is associated with the DSCP. The per-hop behavior is defined as an externally observable forwarding behavior applied at a DS-compliant node to a collection of packets with the same DSCP value.

One of the primary principles of the DiffServ model is that you should mark packets as close to the edge of the network as possible. It is often a difficult and time-consuming task to understand to which traffic class a data packet belongs. Therefore, you want to classify the data as few times as possible. By marking the traffic at the network edge, core network devices and other devices along the forwarding path will be able to quickly determine the proper class of service (CoS) to apply to a given traffic flow.

The primary advantage of the DiffServ model is scalability.

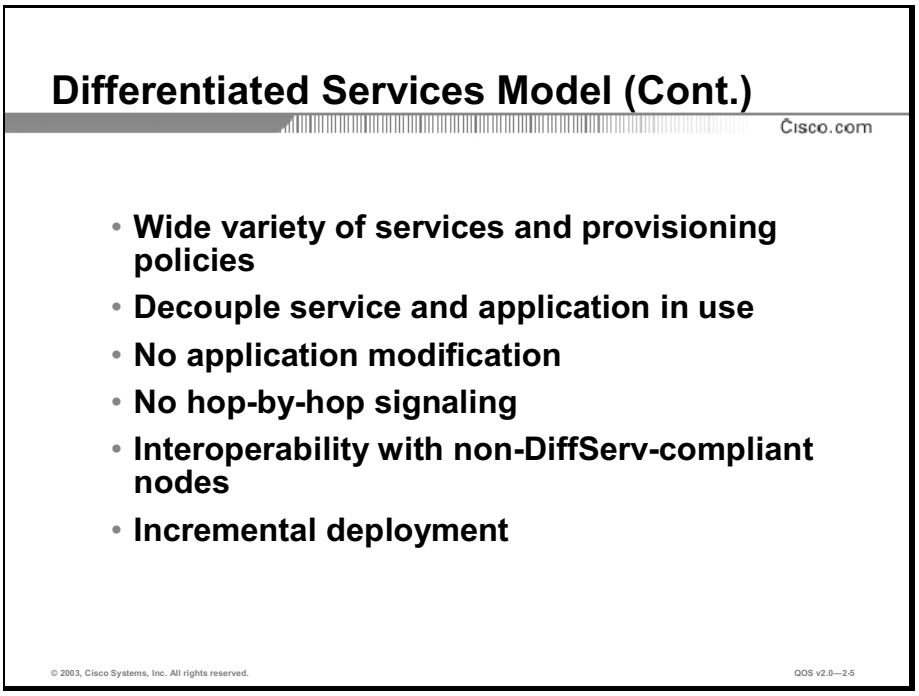

DiffServ is used for mission-critical applications and for providing end-to-end QoS. Typically, DiffServ is appropriate for aggregate flow because it performs a relatively coarse level of traffic classification.

The DiffServ model describes services and allows many user-defined services to be used in a DiffServ-enabled network.

Services are defined as QoS requirements and guarantees that are provided to a collection of packets with the same DSCP value. Services are provided to classes. A class can be identified as a single application, multiple applications with like service needs, or, based on source or destination IP addresses.

Provisioning is used to allocate resources to defined traffic classes. An example of provisioning would be the set of methods that are used to set up the network configurations on devices that would correctly enable the devices to provide the correct set of capabilities for a particular traffic class.

The idea is for the network to recognize a class without having to receive a request from applications. This allows the QoS mechanisms to be applied to other applications that do not have the RSVP functionality, which is the case in 99 percent of applications that use IP.

The introduction of DSCPs replaces IP precedence, a three-bit field in the type of service (ToS) byte of the IP header originally used to classify and prioritize types of traffic. However, DiffServ maintains interoperability with non-DiffServ-compliant devices (those that still use IP precedence). Because of this backward compatibility, DiffServ can be deployed gradually in large networks.

DSCP Encoding

This topic describes the basic format of DSCP and explains the purpose of the DSCP field in the IP header.

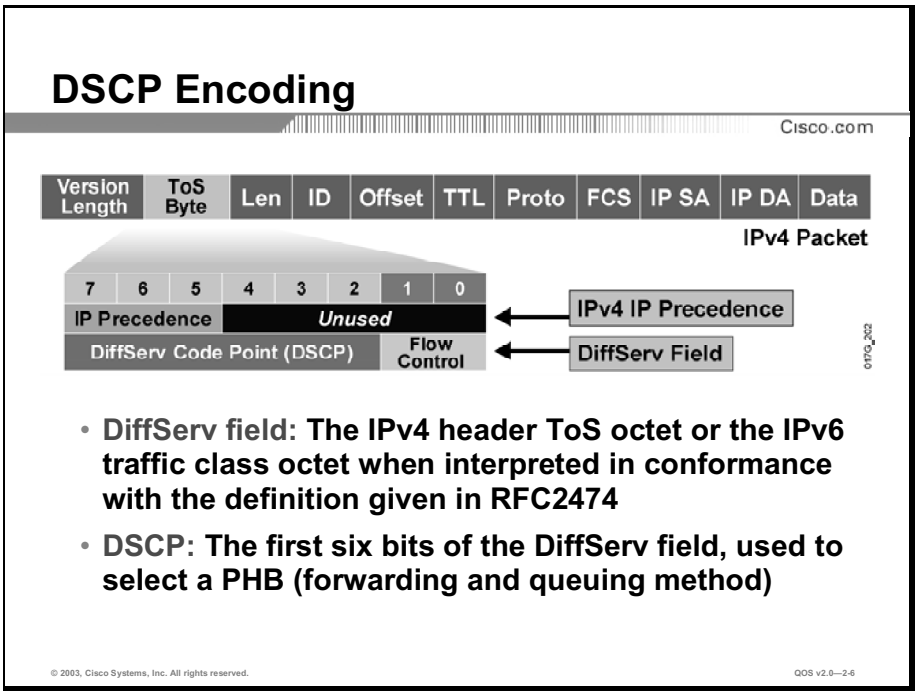

The DiffServ model uses the DiffServ field in the IP header to mark packets according to their classification into behavior aggregates (BAs). The DiffServ field occupies the same eight bits of the IP header that were previously used for the ToS byte.

There are three IETF standards describing the purpose of those eight bits:

- RFC 791 includes specification of the ToS field where the high-order three bits are used for IP precedence. The other bits are used for delay, throughput, reliability, and cost.
- RFC 1812 modifies the meaning of the ToS field by removing meaning from the five low-order bits (those bits should all be zero). This gained widespread use and became known as the original IP precedence.
- RFC 2474 replaces the ToS field with the DiffServ field where the six high-order bits are used for the DSCP. The remaining two bits are used for explicit congestion notification.

Each DSCP value identifies a BA. Each BA is assigned a per-hop behavior (PHB). Each PHB is implemented using the appropriate QoS mechanism or a set of QoS mechanisms.

Per-Hop Behaviors

This topic defines and explains the different per-hop behaviors that are used in DSCP.

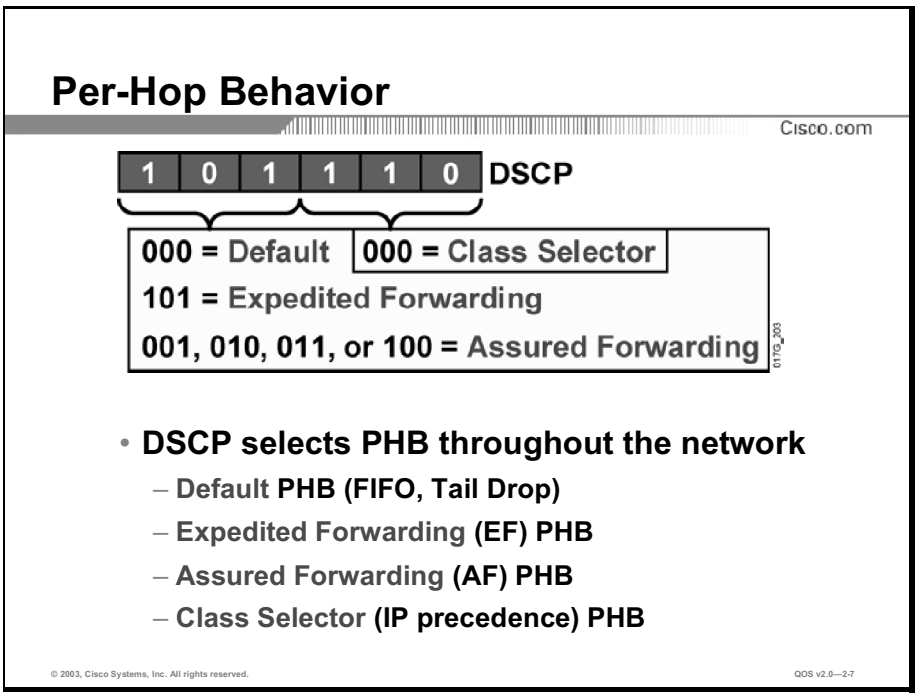

The following PHBs are defined by IETF standards:

- **Default PHB:** Used for Best-Effort service (bits 5 to 7 of DSCP = “000”)
- **Expedited Forwarding PHB:** Used for low-delay service (bits 5 to 7 of DSCP = “101”)
- **Assured Forwarding PHB:** Used for guaranteed bandwidth service (bits 5 to 7 of DSCP = “001”, “010”, “011”, or “100”)
- **Class Selector PHB:** Used for backward compatibility with non-DS-compliant devices (RFC 1812 compliant devices [bits 2 to 4 of DSCP = “000”])

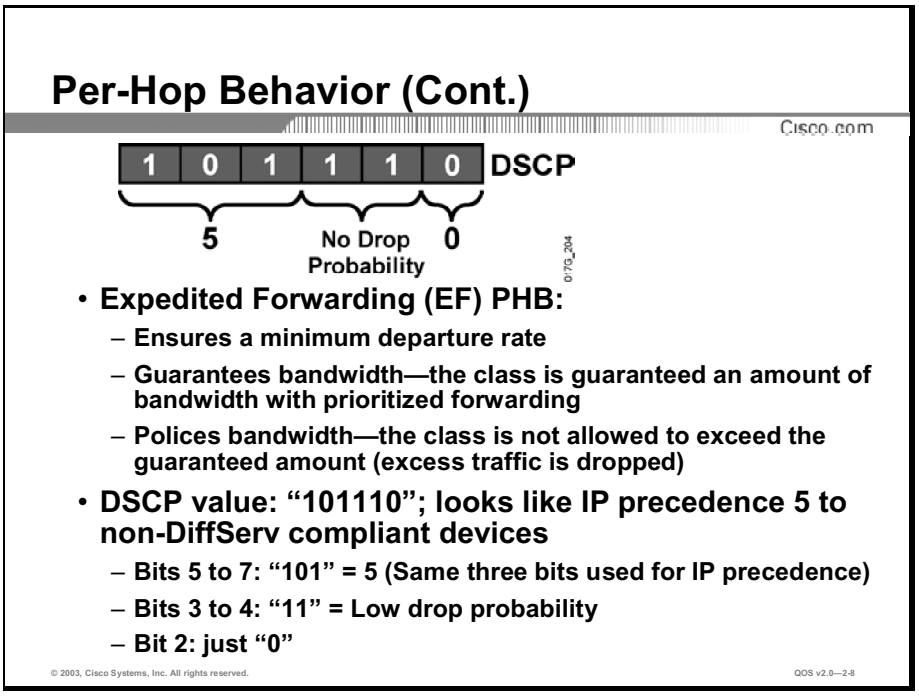

The EF PHB is identified, based on the following parameters:

- Ensures a minimum departure rate. Provides the lowest possible delay to delay-sensitive applications.
- Guarantees bandwidth. Prevents starvation of the application if there are multiple applications using EF PHB.
- Polices bandwidth. Prevents starvation of other applications or classes that are *not* using this PHB.
- Packets requiring Expedited Forwarding should be marked with DSCP binary value “101110” (46 or 0x2E).

Non-DiffServ-compliant devices will regard EF DSCP value “101110” as IP precedence 5 (101).This precedence is the highest user-definable IP precedence and is typically used for delay-sensitive traffic (such as VoIP). Bits 5 to 7 of the EF DSCP value are “101,” which matches IP precedence 5 and allows backward compatibility.

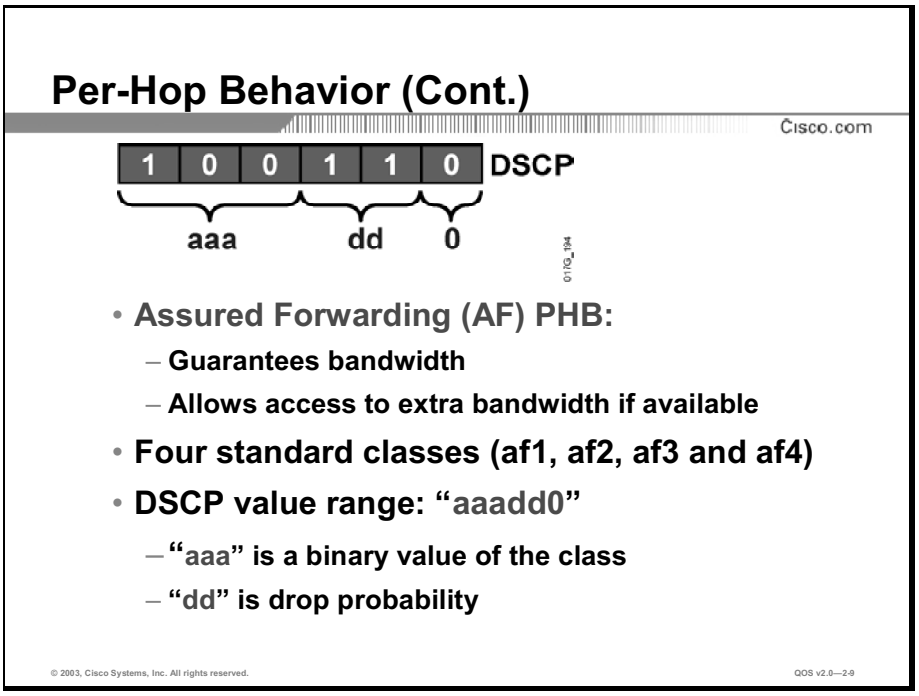

The Assured Forwarding (AF) PHB is identified, based on the following parameters:

- Guarantees a certain amount of bandwidth to an AF class.
- Allows access to extra bandwidth, if available.
- Packets requiring AF PHB should be marked with DSCP value "aaadd0" where "aaa" is the number of the class and "dd" is the drop probability.

There are four standard-defined AF classes. Each class should be treated independently and have allocated bandwidth that is based on the QoS policy.

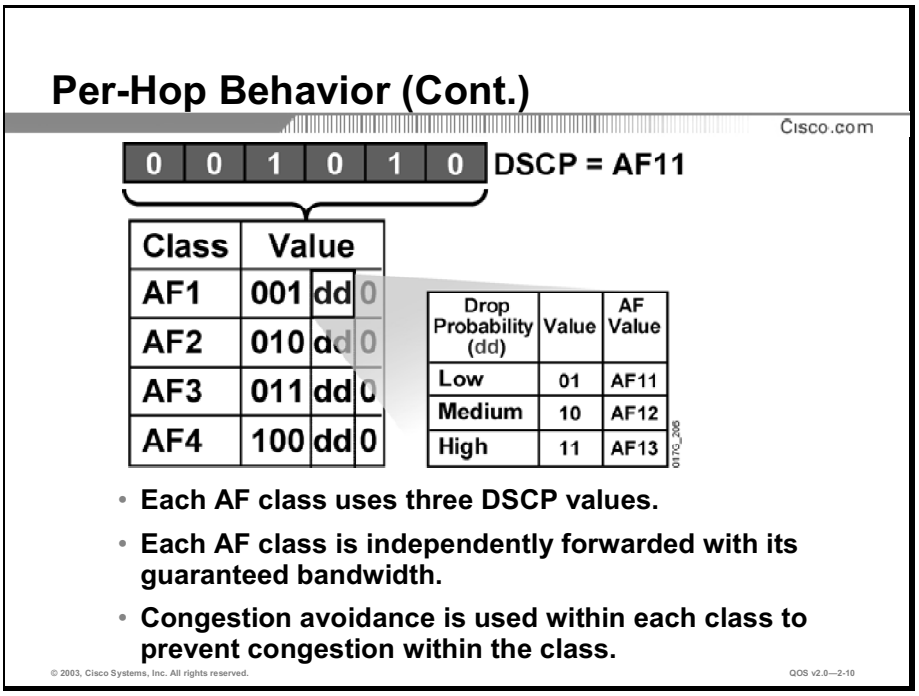

As illustrated in the figure and the table, there are three DSCP values assigned to each of the four AF classes.

AF Class

AF Class	Drop Probability	DSCP Value
AF Class 1	Low	001 01 0
	Medium	001 10 0
	High	001 11 0
AF Class 2	Low	010 01 0
	Medium	010 10 0
	High	010 11 0
AF Class 3	Low	011 01 0
	Medium	011 10 0
	High	011 11 0
AF Class 4	Low	100 01 0
	Medium	100 10 0
	High	100 11 0

Per-Hop Behavior (Cont.)

Cisco.com

- **A DiffServ node must allocate a configurable, minimum amount of forwarding resources (buffer space and bandwidth) per AF class.**
- **Excess resources may be allocated between non-idle classes. The manner must be specified.**
- **Reordering of IP packets of the same flow is not allowed if they belong to the same AF class.**

An AF implementation must attempt to minimize long-term congestion within each class, while allowing short-term congestion resulting from bursts. This requires an active queue management algorithm. An example of such an algorithm is WRED that is described in detail in the “Congestion Avoidance” module in this course.

The AF specification does not define the use of a particular algorithm, but does require that several properties hold.

An AF implementation must detect and respond to long-term congestion within each class by dropping packets, while handling short-term congestion (packet bursts) by queuing packets. This implies the presence of a smoothing or filtering function that monitors the instantaneous congestion level and computes a smoothed congestion level. The dropping algorithm uses this smoothed congestion level to determine when packets should be discarded.

The dropping algorithm must treat all packets within a single class and precedence level identically. Therefore, within a single traffic class, the discard rate of a particular packet flow will be proportional to the percentage of the total amount of traffic.

Backward Compatibility Using the Class Selector

This topic explains the interoperability between DSCP-based and IP-precedence-based devices in a network.

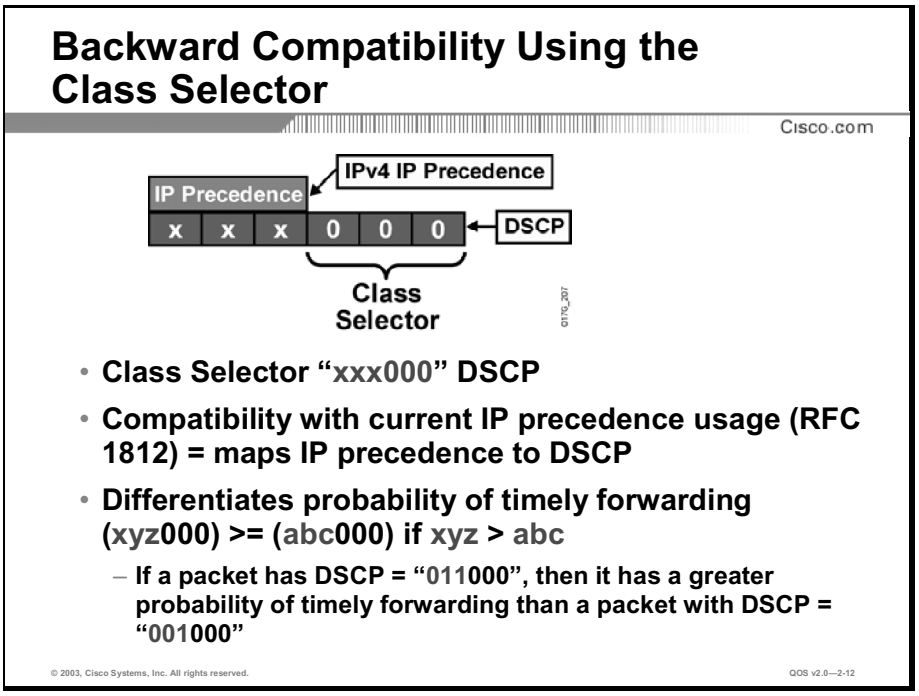

The meaning of the 8 bits in the DiffServ field of the IP packet has changed over time to meet the expanding requirements of IP networks.

Originally, the field was referred to as the ToS field and the first 3 bits of the field (bits 7 to 5) defined a packet **IP precedence** value. A packet could be assigned one of six priorities based on the value of the IP precedence value (8 total values minus 2 reserved values). IP precedence 5 ("101") was the highest priority that could be assigned. (RFC 791)

RFC 2474 replaced the ToS field with the DiffServ field where a range of eight values (class selector) is used for backward compatibility with IP precedence. There is no compatibility with other bits used by the ToS field.

The class selector PHB was defined to provide backwards compatibility for DSCP with ToS-based IP precedence. RFC 1812 simply prioritizes packets according to the precedence value. The PHB is defined as the probability of timely forwarding. Packets with higher IP precedence should be (on average) forwarded in less time than packets with lower IP precedence.

The last three bits of the DSCP (2 to 4) set to zero identify a class selector PHB.

Summary

This topic summarizes the key points discussed in this lesson.

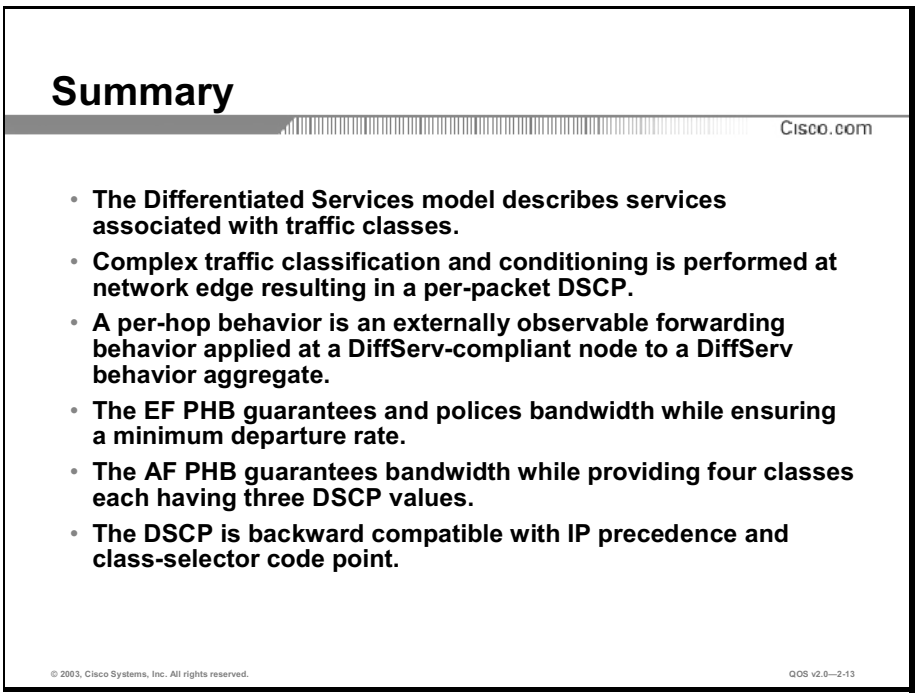

References

For additional information, refer to these resources:

- To learn more about Differentiated Services, refer to “An Architecture for Differentiated Services” at the following URL: http://www.ietf.org/rfc/rfc2475.txt
- To learn more about the DS field, refer to “Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers” at the following URL: http://www.ietf.org/rfc/rfc2474.txt
- To learn more about Assured Forwarding, refer to “Assured Forwarding per-hop behavior (PHB) Group” at the following URL: http://www.ietf.org/rfc/rfc2597.txt
- To learn more about Expedited Forwarding, refer to “An Expedited Forwarding per-hop behavior (PHB)” at the following URL: http://www.ietf.org/rfc/rfc2598.txt
- To learn more about the Class Selector PHB, refer to RFC 2474 at the following URL: http://www.ietf.org/rfc/rfc2474.txt

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) In the differentiated service model, where should packets be marked?

A) just as they hit the first congested link

B) as they leave the service provider core

C) as soon as they hit the service provider core

D) as close to the edge of the network as possible

Q2) Using DiffServ, how aware of the actual application is the core of the network?

A) no awareness at all

B) awareness at layer 2 through MAC address

C) awareness at layer 5 through HTTP connections

D) awareness at layer 4 through TCP/UDP port numbers

Q3) What is the term used to describe the forwarding behavior that is applied at a DiffServ-compliant node to a traffic class?

A) Per-Hop Behavior (PHB)

B) Behavior Aggregate (BA)

C) Behavior Modification (BM)

D) Class Behavior Mechanism (CBM)

Q4) The DSCP field makes up which bits of the entire 8-bit DiffServ field?

A) the entire 8 bits of the DiffServ field

B) the last 6 least significant bits (0 to 5)

C) the first 6 most significant bits (2 to 7)

D) maps over the IP precedence portion (bits 5 to 7)

Q5) Which three of the following Per-Hop Behaviors (PHBs) is determined by content of the first 3 bits (5 to 3) of the DSCP field? (Choose three.)

A) default PHB

B) class selector PHB

C) Assured Forwarding PHB

D) Expedited Forwarding PHB

Q6) Which DSCP value appears to be IP precedence 5 to a non-DiffServ-compliant device?

A) 5 “101”

B) 45 “101101”

C) 46 “101110”

D) 47 “101111”

Quiz Answer Key

Q1) D

Relates to: Differentiated Services Model

Q2) A

Relates to: Differentiated Services Model

Q3) A

Relates to: Differentiated Services Model

Q4) C

Relates to: DSCP Encoding

Q5) A, C, D

Relates to: Per-Hop Behaviors

Q6) C

Relates to: Backward Compatibility Using the Class Selector

IP QoS Mechanisms

Overview

IP QoS mechanisms are used to implement a coordinated QoS policy in devices throughout the network. The moment an IP packet enters the network, it is classified and usually marked with its class identification. From that point on, the packet is treated by a variety of IP QoS mechanisms according to the packet classification. Depending upon the mechanisms it encounters, the packet could be expedited, delayed, compressed, fragmented, or even dropped.

Relevance

The IP QoS mechanisms described in this lesson form the base technologies that are used to implement QoS in any IP network.

Objectives

Upon completing this lesson, you will be able to correctly match a list of QoS actions to mechanisms for implementing QoS and identify where in a network the different QoS mechanisms are commonly used. This includes being able to meet these objectives:

- List the key mechanisms used to implement QoS in an IP network
- Define classification and identify where classification is commonly implemented in a network
- Define marking and identify where marking is commonly implemented in a network
- Define congestion management and identify where congestion management is commonly implemented in a network
- Define congestion avoidance and identify where congestion avoidance is commonly implemented in a network
- Define policing and shaping and identify where policing and shaping are commonly implemented in a network
- Explain the functions of compression and identify where compression is commonly implemented in the network
- Explain the functions of LFI and identify where LFI is commonly implemented in the network
- Identify whether QoS mechanisms are used for input or output or both

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Basic knowledge of internetworking with TCP/IP concepts

Outline

The outline lists the topics included in this lesson.

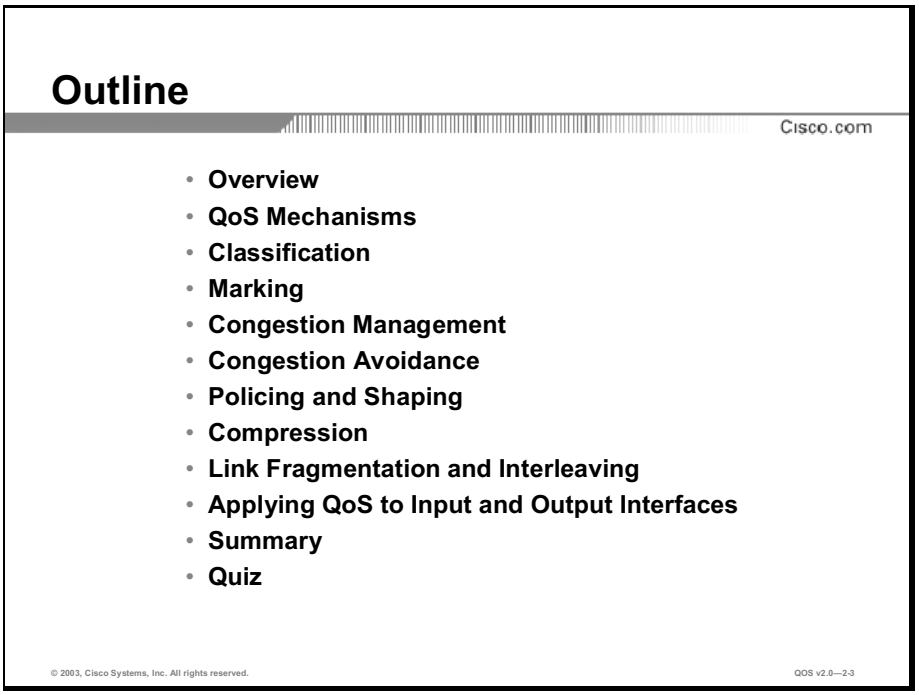

QoS Mechanisms

This topic lists the key mechanisms use to implement QoS in an IP network.

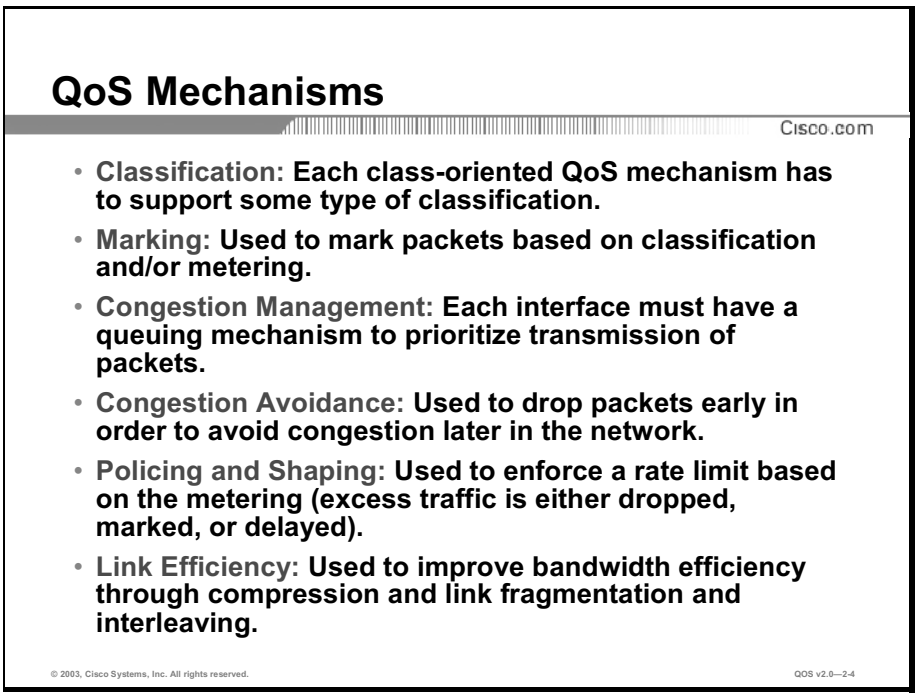

This slide shows the main categories of tools and describes in lay terms how they contribute to QoS.

Classification and marking is the identifying and splitting of traffic into different classes and the marking of traffic according to behavior and business policies.

Congestion management is the prioritizing, protection, and isolation of traffic based on markings.

Congestion avoidance discards specific packets based on markings, to avoid network congestion.

Traffic conditioning mechanisms police traffic by dropping misbehaving traffic to maintain network integrity. They also shape traffic to control bursts by queuing traffic.

One type of link efficiency technology is packet header compression that improves the bandwidth efficiency of a link. Another technology is link fragmentation and interleaving (LFI) that can decrease the “jitter” of voice transmission by reducing voice packet delay.

Classification

This topic defines classification and identifies where classification is commonly implemented in a network.

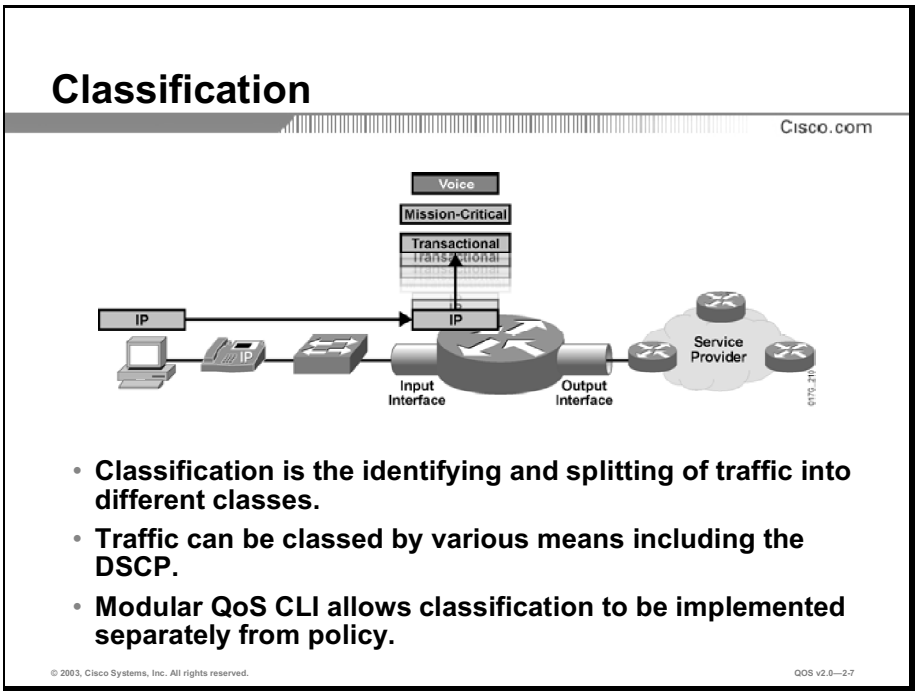

Classification is the identifying and splitting of traffic into different classes. In a QoS-enabled network, all traffic is classified at the input interface of every QoS-aware device. Packet classification can be recognized based on many factors including:

- DSCP
- IP precedence
- Source address
- Destination address

The concept of "trust" is key for deploying QoS. When an end device (such as a workstation or an IP Phone) marks a packet with CoS or DSCP, a switch or router has the option of accepting or not accepting values from the end device. If the switch or router chooses to accept the values, the switch or router "trusts" the end device. If the switch or router trusts the end device, it does not need to do any reclassification of packets coming from that interface. If the switch or router does not trust the interface, then it must perform a reclassification to determine the appropriate QoS value for the packet coming from that interface. Switches and routers are generally set to "not trust" end devices and must specifically be configured to "trust" packets coming from an interface.

Classification tools include NBAR, policy-based routing (PBR), and classification and marking using modular QoS command-line interface (CLI [MQC]).

Note: The tools for classification are covered in detail in the “Classification and Marking” module in this course.

Marking

This topic defines marking and identifies where marking is commonly implemented in a network.

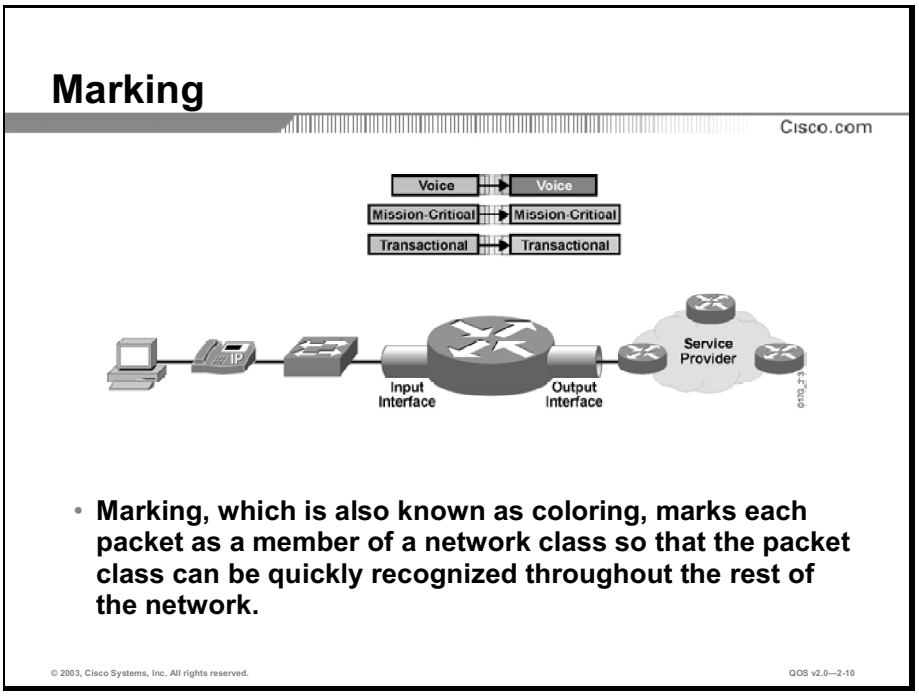

Marking, which is also known as coloring, involves marking each packet as a member of a network class so that devices throughout the rest of the network can quickly recognize the packet class. Marking is performed as close to the network edge as possible, and is typically done using the MQC.

QoS mechanisms set bits in the DSCP or IP precedence fields of each IP packet according to the class that the packet is in. The settings for the DSCP field and their relationship to the IP precedence fields were discussed in the previous lesson. Other fields can also be marked to aid in the identification of a packet class.

Other QoS mechanisms use these bits to determine how to treat the packets when they arrive. If they are marked as high-priority voice packets, the packets will generally never be dropped by congestion avoidance mechanisms and be given immediate preference by congestion management queuing mechanisms. On the other hand, if the packets are marked as low-priority file transfer packets, they will be dropped when congestion is occurring and generally move to the end of the congestion management queues.

Note: The tools for marking are covered in detail in the “Classification and Marking” module in this course.

Congestion Management

This topic defines congestion management and identifies where congestion management is commonly implemented in a network.

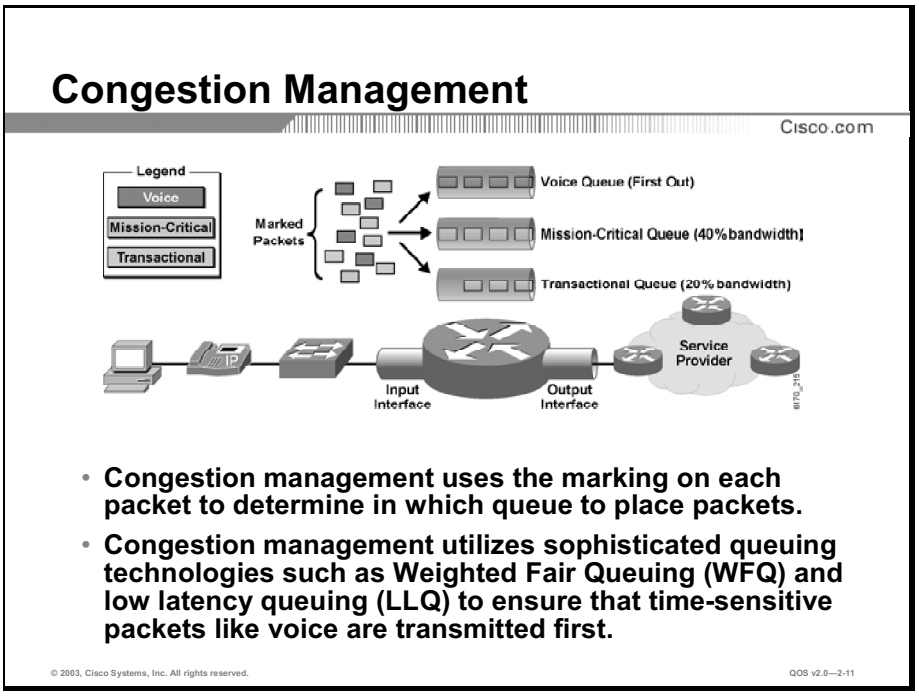

Congestion management mechanisms (queuing algorithms) use the marking on each packet to determine in which queue to place packets. Different queues are given different treatment by the queuing algorithm based on the class of packets in the queue. Generally, queues with higher priority packets receive preferential treatment.

All output interfaces in a QoS-enabled network use some kind of congestion management (queuing) mechanism to manage the outflow of traffic. Each queuing algorithm was designed to solve a specific network traffic problem and has a particular effect on network performance.

The Cisco IOS software features for congestion management or queuing, include:

- FIFO
- Priority queuing (PQ)
- Custom queuing (CQ)
- Weighted fair queuing (WFQ)
- Class-based weighted fair queuing (CBWFQ)
- LLQ

LLQ is currently the preferred queuing method. It is a hybrid (PQ and CBWFQ) queuing method that was developed to specifically meet the requirements of real-time traffic, such as voice.

Note: All of the queuing technologies discussed are described further in the “Congestion Management” module in this course.

Congestion Avoidance

This topic defines congestion avoidance and identifies where congestion avoidance is commonly implemented in a network.

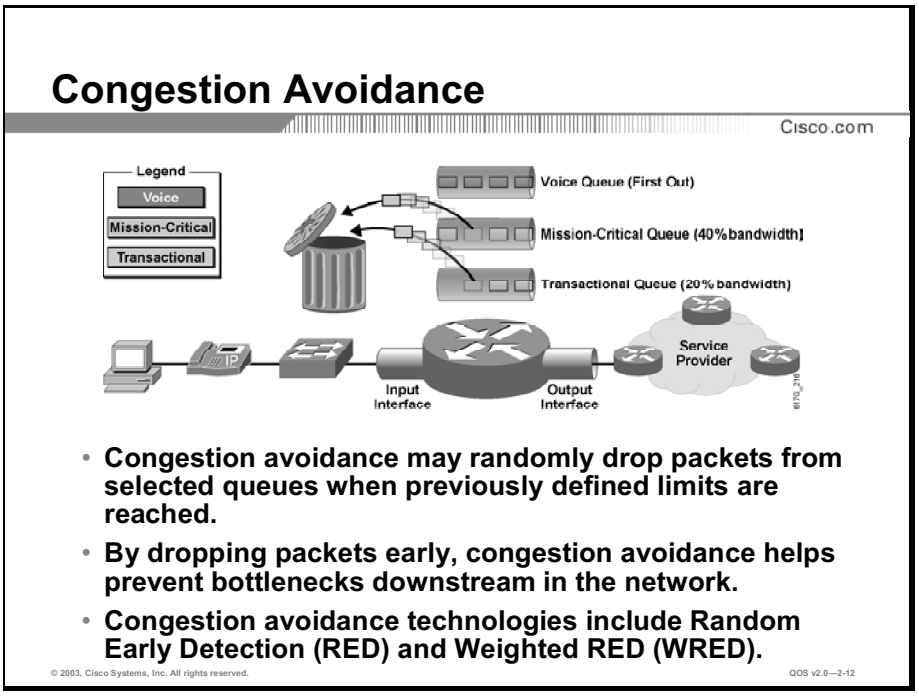

Congestion-avoidance mechanisms monitor network traffic loads in an effort to anticipate and avoid congestion at common network bottlenecks. Congestion avoidance is achieved through packet dropping.

Congestion avoidance mechanisms are typically employed on output interfaces wherever a high-speed link or set of links feed into a lower-speed link (that is, a LAN feeding into a slower WAN link.) This ensures that the WAN is not instantly congested by LAN traffic.

WRED is a Cisco primary congestion-avoidance technique.

WRED increases the probability that congestion is avoided by dropping low-priority packets rather than high-priority packets.

Note:	WRED is not recommended for voice queues. A network should not be designed to drop voice packets.

Note:	The tools for congestion avoidance are covered in detail in the "Congestion Avoidance" module in this course.

Policing and Shaping

This topic defines policing and shaping and identifies where policing and shaping are commonly implemented in a network.

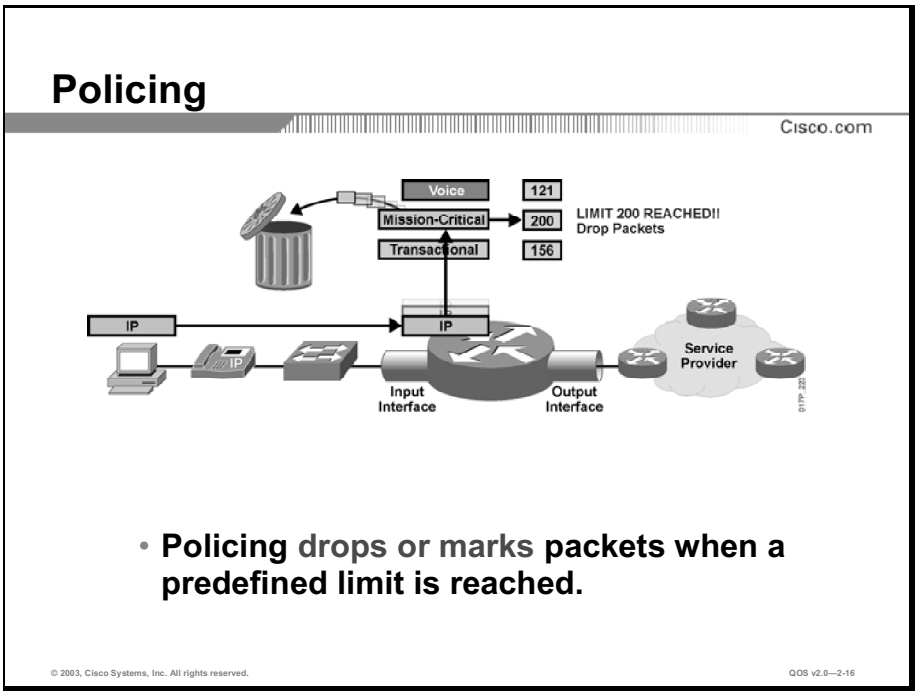

Policing or shaping mechanisms are often used to condition traffic before transmitting traffic to a network or receiving traffic from a network.

Policing is the ability to control bursts and conform traffic to ensure that certain types of traffic get certain types of bandwidth.

Policing drops or marks packets when predefined limits are reached.

Policing mechanisms can be set to first drop traffic classes that have lower QoS priority markings.

Policing mechanisms can be used at either input or output interfaces. They are typically used to control the flow into a network device from a high-speed link by dropping excess low-priority packets. A good example would be the use of policing by a service provider to throttle a high-speed inflow from a customer that was in excess of the service agreement. In a TCP environment, this will cause the sender to slow their packet transmission.

Tools include class-based policing and committed access rate (CAR).

Note:	The tools for policing are covered in detail in the “Traffic Policing and Shaping” module in this course.

Shaping

Cisco.com

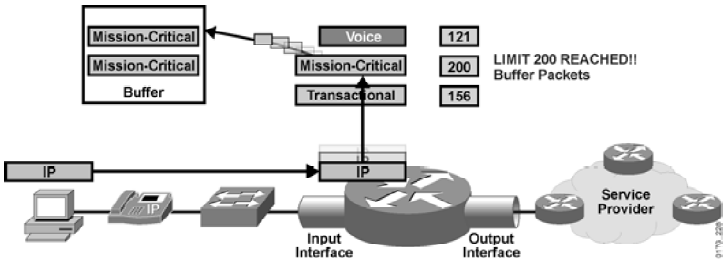

- **Shaping queues packets when a predefined limit is reached.**

Shaping helps smooth out speed mismatches in the network and limits transmission rates.

Shaping mechanisms are used on output interfaces. They are typically used to limit the flow from a high-speed link to a lower-speed link to ensure that the lower-speed link does not become overrun with traffic. Shaping could also be used to manage the flow of traffic at a point in the network where multiple flows are aggregated. Service providers use it to manage the flow of traffic to and from customers to ensure that the flows conform to service agreements between the customer and provider.

Cisco QoS software solutions include two traffic-shaping tools to manage traffic and congestion on the network: Generic Traffic Shaping (GTS) and Frame Relay traffic shaping (FRTS).

Note: The tools for shaping are covered in detail in the “Traffic Policing and Shaping” module in this course.

Compression

This topic explains the functions of compression and identifies where compression is commonly implemented in the network.

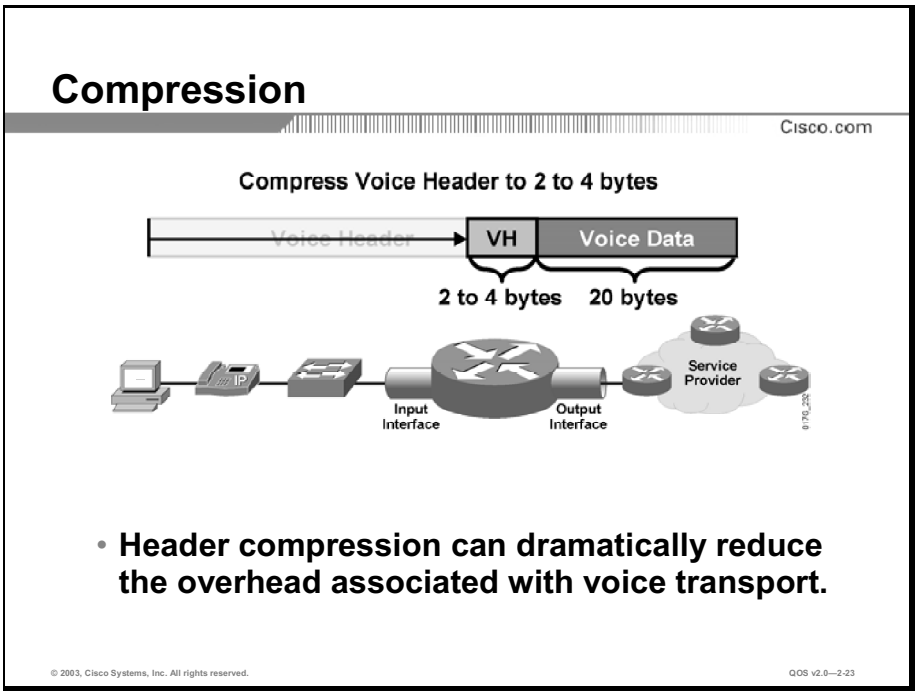

Cisco IOS QoS software offers link-efficiency mechanisms that work in conjunction with queuing and traffic shaping to manage existing bandwidth more efficiently and predictably. One of these is compressed Real-Time Transport Protocol (cRTP).

Real-Time Transport Protocol (RTP) is a host-to-host protocol that is used for carrying converged traffic (including packetized audio and video) over an IP network. RTP provides end-to-end network transport functions intended for applications that transmit real-time requirements such as audio, video, simulation data multicast, or unicast network services.

A voice packet carrying a 20-byte voice payload, for example, typically carries a 20-byte IP header, an 8-byte User Datagram Protocol (UDP) header, and a 12-byte RTP header. By using cRTP, as shown in the illustration, the three headers of a combined 40 bytes are compressed down to 2 or 4 bytes, depending on whether or not the CRC is transmitted. This compression can dramatically improve the performance of a link.

Compression would typically be used on WAN links between sites to improve bandwidth efficiency.

Note:	Compression technology is discussed in the “Link Efficiency Mechanisms” module in this course.

Link Fragmentation and Interleaving

This topic explains the functions of LFI and identifies where LFI is commonly implemented in the network.

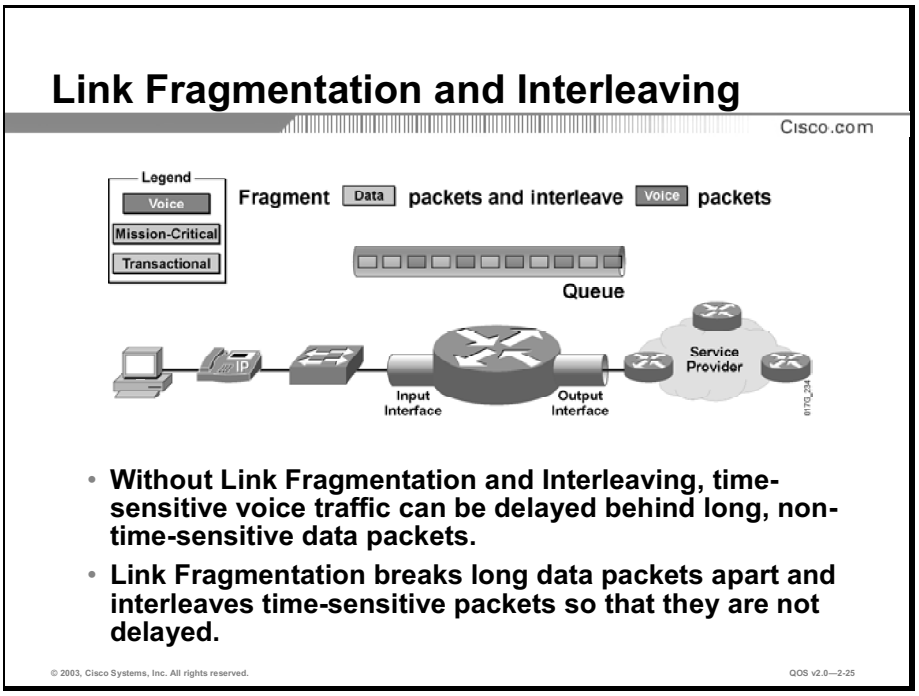

Interactive traffic, such as Telnet and VoIP, is susceptible to increased latency and jitter when the network processes large packets, such as LAN-to-LAN FTP Telnet transfers, traversing a WAN link. This susceptibility increases as the traffic is queued on slower links.

LFI can reduce delay and jitter on slower-speed links by breaking up large datagrams and interleaving low-delay traffic packets with the resulting smaller packets.

LFI would typically be used on WAN links between sites to ensure minimal delay for voice and video traffic.

Note: LFI technology is covered in detail in the “Link Efficiency Mechanisms” module in this course.

Applying QoS to Input and Output Interfaces

This topic identifies whether QoS mechanisms are used for input or output or both.

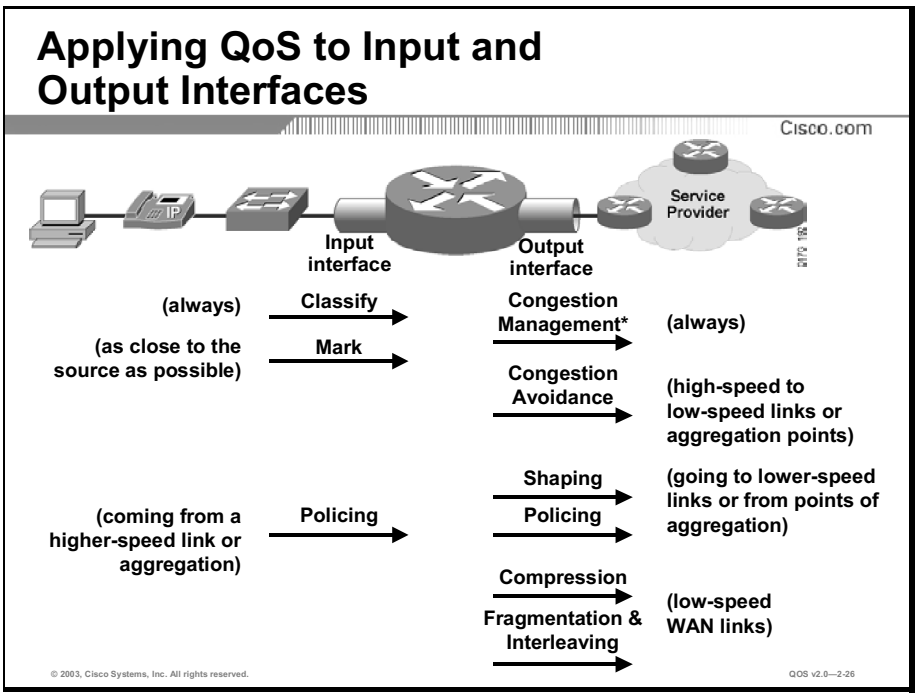

In a QoS-enabled network, classification is performed on every input interface.

Marking should be performed as close to the network edge as possible—in the originating network device, if possible. Devices farther from the edge of the network, such as routers and switches, can be configured to "trust" or "untrust" the markings made by devices on the edge of the network. An IP Phone, for example, will not "trust" the markings of an attached PC while a switch will generally be configured to "trust" the markings of an attached IP Phone.

Congestion management, congestion avoidance, and traffic shaping mechanisms only make sense to use on output interfaces as they help maintain smooth operation of links by controlling how much and which type of traffic is allowed on a link. On some router and switch platforms, congestion management mechanisms such as weighted round robin (WRR) and modified deficit round robin (MDRR) can be applied on the input interface.

Congestion avoidance is typically employed on an output interface wherever there is a chance that a high-speed link or aggregation of links feeds into a slower link (a LAN feeding into a WAN).

Policing and shaping are typically employed on output interfaces to control the flow of traffic from a high-speed link to lower-speed links. Policing is also employed on input interfaces to control the flow into a network device from a high-speed link by dropping excess low-priority packets.

Both compression and LFI are typically used on slower-speed WAN links between sites to improve bandwidth efficiency.

Summary

This topic summarizes the key points discussed in this lesson.

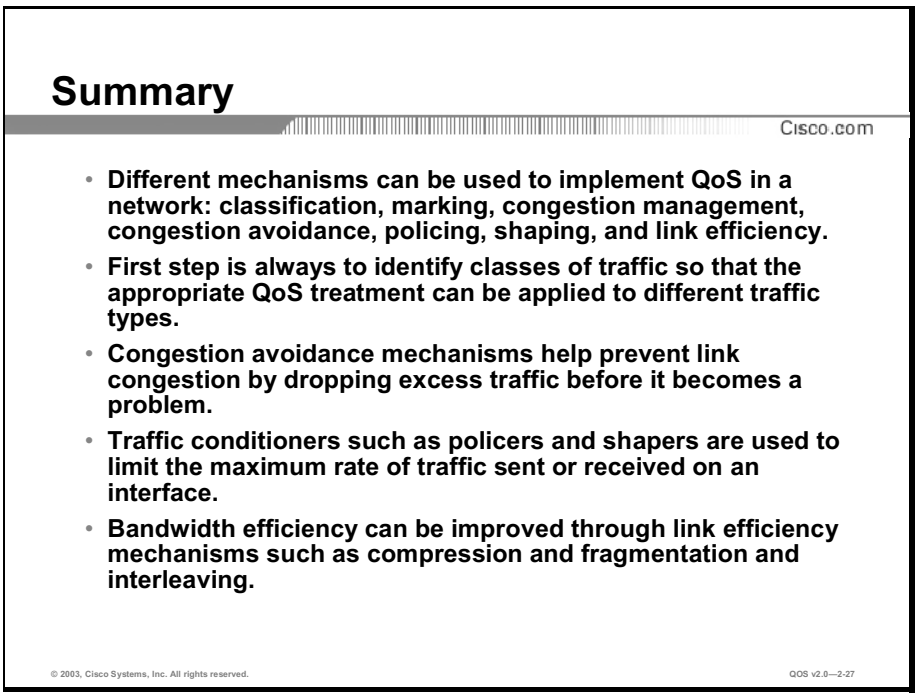

References

For additional information, refer to this resource:

- To see more information on Cisco IP QoS mechanisms, refer to “Quality of Service (QoS)” at the following URL: http://www.cisco.com/en/US/tech/tk543/tech_topology_and_network_serv_and_protocol_suite_home.html

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) Which of the following IP QoS mechanisms queues the transmission of packets?

A) metering

B) traffic shaping

C) traffic policing

D) congestion avoidance

Q2) By which four fields are IP packets often classified? (Choose four.)

A) TTL

B) DSCP

C) IP precedence

D) source address

E) destination address

Q3) Which one of the following congestion management technologies was developed specifically to meet the requirements of real-time traffic such as voice?

A) low latency queuing (LLQ)

B) weighted fair queuing (WFQ)

C) class-based WFQ (CBWFQ)

D) priority-voice queuing (PVQ)

Q4) The acronym RED stands for _____.

A) random early detection

B) regular expedited dropping

C) regular early dropping

D) random early dropping

Q5) Which two IP QoS mechanisms manage traffic by queuing packets? (Choose two.)

A) traffic shaping

B) traffic policing

C) congestion avoidance

D) congestion management

Q6) How many header bytes does the normal voice packet carry for IP, UDP, and RTP?

A) 32

B) 40

C) 64

D) 80

Quiz Answer Key

Q1) B

Relates to: QoS Mechanisms

Q2) B, C, D, E

Relates to: Classification

Q3) A

Relates to: Congestion Management

Q4) A

Relates to: Congestion Avoidance

Q5) A, D

Relates to: Policing and Shaping

Q6) B

Relates to: Compression

Case Study: QoS Mechanisms

Overview

This case study activity provides information regarding the QoS administrative policy requirements of a large, multisite network. Your task is to work with a partner to evaluate the QoS requirements, and based on these requirements, identify where QoS mechanisms should be applied. You will discuss your solution with the instructor and other classmates, and the instructor will present a solution for the case study to the class.

Relevance

The ability to properly sort traffic into service classes and correctly position QoS mechanisms are important steps in correctly implementing an administrative QoS policy.

Objectives

In this activity, you will correctly identify which QoS mechanisms can be used, and where QoS mechanisms should be applied to the network to implement an administrative QoS policy. Upon completing this case study, you will be able to meet these objectives:

- Review customer QoS requirements
- Identify QoS service class requirements
- Identify where QoS mechanisms should be applied to the network to meet customer requirements
- Present a solution to the case study

Learner Skills and Knowledge

To benefit fully from this activity, you must have these prerequisite skills and knowledge:

- Basic knowledge of internetworking with TCP/IP concepts

Outline

The outline lists the topics included in this activity.

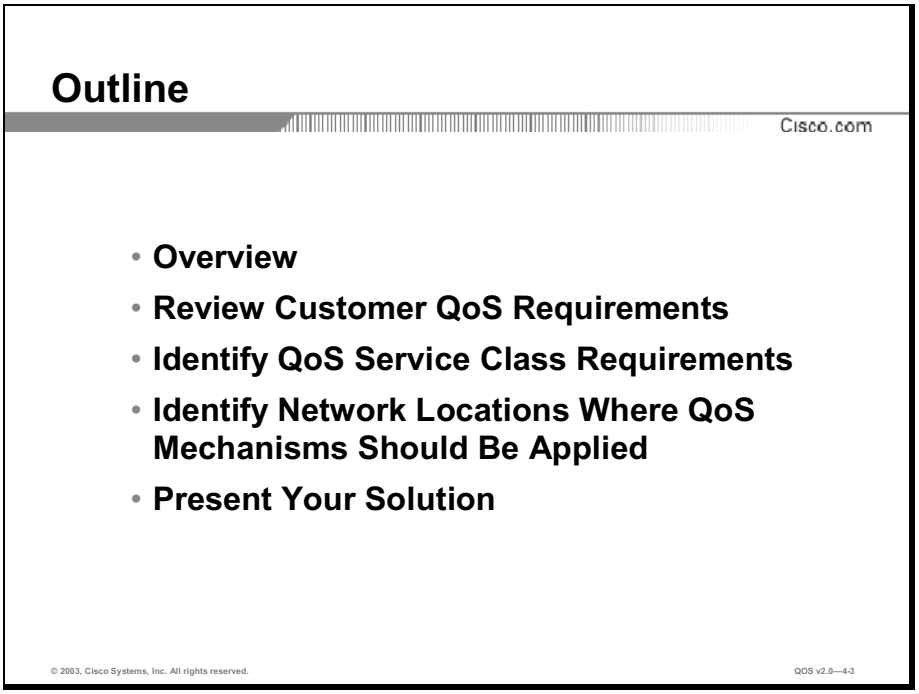

Required Resources

These are the resources required to complete this exercise:

- Case Study Activity: QoS Mechanisms
- A workgroup consisting of two learners

Job Aids

No job aids are required to complete this case study.

Case Study Tasks

The activity includes these tasks:

Step 1 **Review customer QoS requirements:** Completely read the customer requirements provided.

Step 2 **Identify QoS service class requirements:** With the aid of your partner, identify the service classes required to implement the administrative QoS policy based on customer requirements.

Step 3 **Identify network locations where QoS mechanisms should be applied:** Identify locations in the network where the QoS mechanisms should be applied to most effectively implement QoS policy.

Step 4 **Present your solution:** After the instructor presents a solution to the case study, present your solution to the class with your partner.

Case Study Verification

You have completed this activity when your case study solution has been presented to the class and you have justified any major deviations from the case study solution supplied by the instructor.

Review Customer QoS Requirements

Company Background

Nuevo Health Care Systems (NHCS) provides health care information to health care professionals in ten major regions of the country.

Customer Situation

NHCS network currently has limited bandwidth capacity in their WAN links and they do not envision being able to increase bandwidth in the near future. All ten remote sites (two are pictured in the network illustration) connect to the central site through a service provider through a Frame Relay, Layer 2, 1-Mbps link service. The NHCS headquarters site also connects to the service provider via a Frame Relay, Layer 2, 1-Mbps link. NHCS LAN bandwidth is 10 Mbps. NHCS connects to the Internet through its headquarters site.

Since the installation of a new IP telephony system, NHCS has been encountering increasingly serious problems with their network:

- Users of the enterprise resource planning (ERP) applications have been complaining of unacceptable response times. Their sub-second response time has now stretched to multiple seconds in many cases and up to a minute in some cases.
- Key patient information files that used to arrive almost instantly are now taking 10 to 15 minutes to be transferred from headquarters to users at the remote sites. (These are moderate sized, mostly text files.)
- Patient graphics files (x-rays, MRIs) that used to take 20 to 30 minutes to transfer between the remote sites and headquarters now often have to be transferred overnight. (This is acceptable as they are usually not needed immediately and they tend to be extremely large graphics files.)
- Users of the new IP telephony devices are the most upset. The quality of their calls is very poor and their calls often just drop.

The key applications running on NHCS network are:

Applications Running on NHCS Network

Application	Application Importance	Response Time Requirements	Use of Bandwidth (Daytime)
Enterprise Resource Planning	**Critical**	**Immediate**	**Moderate**
Patient Information Files	**Important**	**Immediate**	**Moderate**
Patient Graphics Files	**Important**	**Minimal**	**Heavy**
IP Telephony	**Important**	**No delay**	**Moderate**
Browser Traffic	**Not important**	**Minimal**	**Heavy**

Nuevo Health Care Systems Network

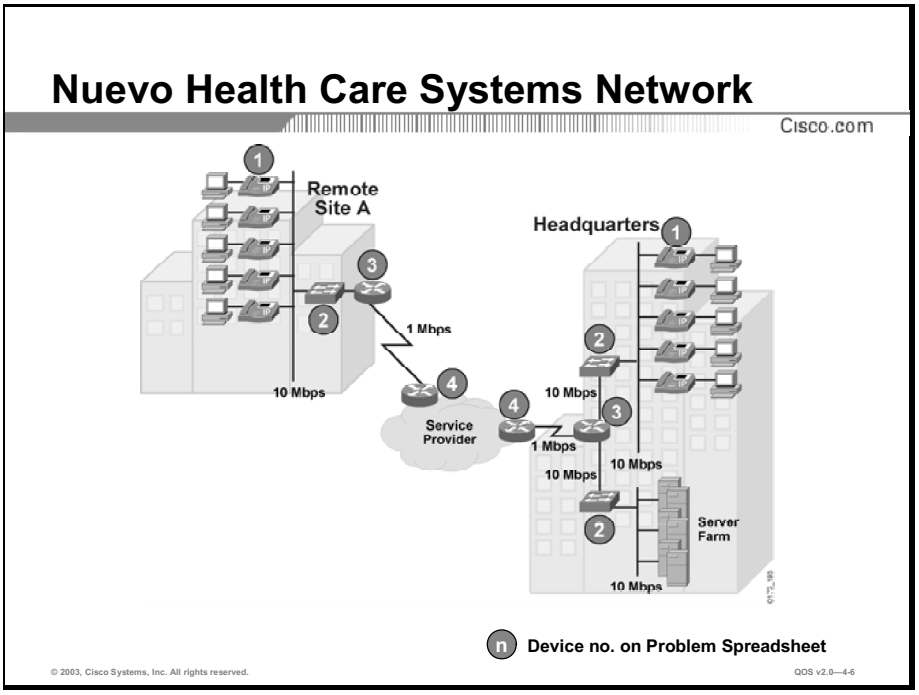

Device Number	Device Type
1	IP Phone
2	LAN Switch
3	Customer Edge Router
4	Service Provider Router

Identify QoS Service Class Requirements

Given the NHCS network as described, how would you recommend classifying network traffic?

Traffic Classification and Prioritization

Type of Traffic (Application)	Traffic Priority (Rank from 1 to 5)

Identify Network Locations Where QoS Mechanisms Should Be Applied

Given NHCS network as described, how would you recommend deploying QoS mechanisms? Check each box (X) where you believe that QoS mechanisms could be applied to effectively resolve QoS problems at NHCS.

Where to Apply QoS Mechanisms: Classification and Marking

Device #	Network Device Interface	Classification On Input	Classification On Output	Marking On Input	Marking On Output
1	IP Phone—Interface to Workstation				
1	IP Phone—Interface to Switch				
2	Switch—Interface to IP Phone				
2	Switch—Interface to Customer Edge Router				
3	Customer Edge Router—Interface to Switch				
3	Customer Edge Router—Interface to WAN (Service Provider Router)				
4	Service Provider Router—Interface to Customer Edge Router				

Where to Apply QoS Mechanisms: Congestion Management and Avoidance

Device #	Network Device Interface	Congestion Management On Input	Congestion Management On Output	Congestion Avoidance On Input	Congestion Avoidance On Output
2	Switch—Interface to IP Phone				
2	Switch—Interface to Customer Edge Router				
3	Customer Edge Router—Interface to Switch				
3	Customer Edge Router—Interface to WAN (Service Provider Router)				
4	Service Provider Router—Interface to Customer Edge Router				

Where to Apply QoS Mechanisms: Traffic Policing and Traffic Shaping

Device #	Network Device Interface	Traffic Policing On Input	Traffic Policing On Output	Traffic Shaping On Input	Traffic Shaping On Output
2	Switch—Interface to IP Phone				
2	Switch—Interface to Customer Edge Router				
3	Customer Edge Router—Interface to Switch				
3	Customer Edge Router—Interface to WAN (Service Provider Router)				
4	Service Provider Router—Interface to Customer Edge Router				

Where to Apply QoS Mechanisms: Link Efficiency

Device #	Network Device Interface	Compression On Input	Compression On Output	LFI On Input	LFI On Output
2	Switch—Interface to IP Phone				
2	Switch—Interface to Customer Edge Router				
3	Customer Edge Router—Interface to Switch				
3	Customer Edge Router—Interface to WAN (Service Provider Router)				
4	Service Provider Router—Interface to Customer Edge Router				

Present Your Solution

Together with your partner, present your solution to the class. Include the following information:

- Customer service class requirements
- Network diagrams indicating where classification and marking should be applied
- Justification for differences from the solution presented by the instructor

Case Study Answer Key

Traffic Classification and Prioritization

Type of Traffic (Application)	Traffic Priority
IP Telephony	Highest - 1
Enterprise Resource Planning	High - 2
Patient Information Files	Moderate - 3
Patient Graphics Files	Low - 4
Browser Traffic	Low - 4

Where to Apply QoS Mechanisms: Classification and Marking

Device #	Network Device Interface	Classification On Input	Classification On Output	Marking On Input	Marking On Output
1	**IP Phone—Link to Workstation**	X		X*	
1	**IP Phone—Link to Switch**	X			
2	**Switch—Link to IP Phone**	X		No, trusted*	
2	**Switch—Link to Customer Edge Router**	X			
3	**Customer Edge Router—Link to Switch**	X			
3	**Customer Edge Router—Link to WAN (Service Provider Router)**	X			
4	**Service Provider Router—Link to Customer Edge Router**	X			

Note: *The IP Phone will normally be set to remark any traffic coming from its downstream workstation (the IP Phone connection to the workstation is “untrusted”). The switch will not remark traffic coming from the IP Phone (traffic from the IP Phone is “trusted”). Further explanation of “trusted” and “untrusted” interfaces is provided in the “Classification and Marking” module of this course.

Where to Apply QoS Mechanisms: Congestion Management and Avoidance

Device #	Network Device Interface	Congestion Management On Input	Congestion Management On Output	Congestion Avoidance On Input	Congestion Avoidance On Output
2	Switch—Link to IP Phone		X		
2	Switch—Link to Customer Edge Router		X		Possible
3	Customer Edge Router—Link to Switch		X		
3	Customer Edge Router—Link to WAN (Service Provider Router)		X		Possible
4	Service Provider Router—Link to Customer Edge Router		X		Possible

Where to Apply QoS Mechanisms: Traffic Policing and Traffic Shaping

Device #	Network Device Interface	Traffic Policing On Input	Traffic Policing On Output	Traffic Shaping On Input	Traffic Shaping On Output
2	Switch—Link to IP Phone	X			
2	Switch—Link to Customer Edge Router				
3	Customer Edge Router—Link to Switch	X			
3	Customer Edge Router—Link to WAN (Service Provider Router)				Possible
4	Service Provider Router—Link to Customer Edge Router	X			Possible

Where to Apply QoS Mechanisms: Link Efficiency

Device #	Network Device Interface	Compression On Input	Compression On Output	LFI On Input	LFI On Output
2	Switch—Link to IP Phone				
2	Switch—Link to Customer Edge Router				
3	Customer Edge Router—Link to Switch				
3	Customer Edge Router—Link to WAN (Service Provider Router)		X		X
4	Service Provider Router—Link to Customer Edge Router		X		X

Note: Because this is a Frame Relay network the service provider will pass frames through transparently without compressing or fragmenting the frames.

Case Study: The Life of a Packet

Overview

This case study activity provides information regarding the application of QoS mechanisms throughout a simple network. The case study follows two packets—a high-priority voice packet and a low-priority file transfer packet —as they traverse a QoS-enabled network.

Relevance

The ability to recognize the exact impact of QoS mechanisms on packets as they traverse a network is vitally important for correctly implementing QoS in a network.

Objectives

In this activity, you will learn how IP QoS mechanisms impact IP packets. Upon completing this case study, you will be able to meet these objectives:

- On a network diagram, identify key points where the QoS status of a high-priority (VoIP) packet can be altered as QoS policies are applied to the IP packet
- On a network diagram, identify key points where the QoS status of a low-priority (FTP) packet can be altered as QoS policies are applied to the IP packet

Learner Skills and Knowledge

To benefit fully from this activity, you must have these prerequisite skills and knowledge:

- Basic knowledge of internetworking with TCP/IP concepts

Outline

The outline lists the topics included in this case study.

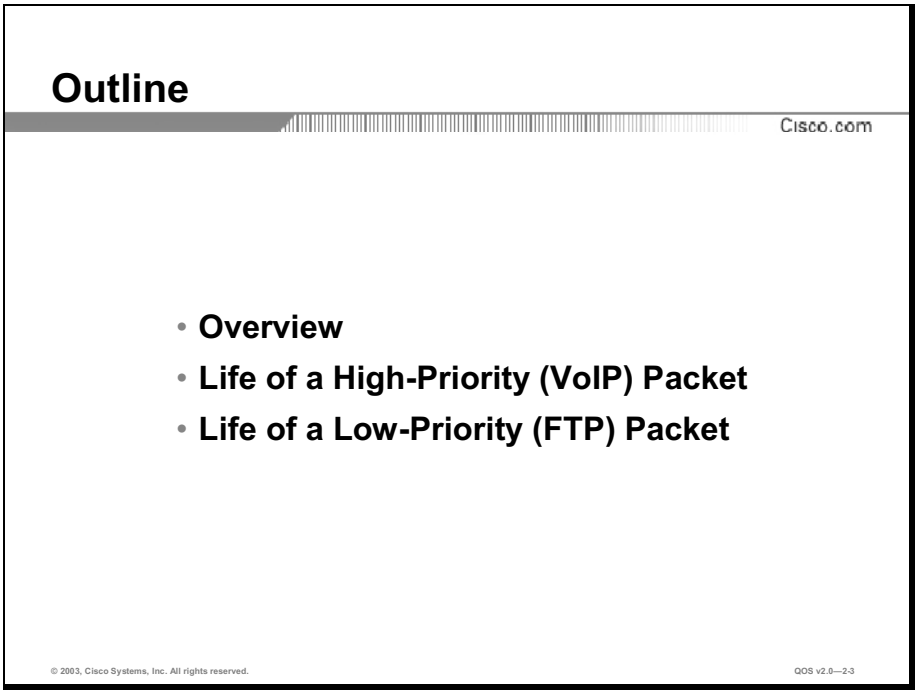

Required Resources

No resources are required to complete this exercise.

Job Aids

No job aids are required to complete this case study.

Case Study Steps

No steps are required for this case study.

Case Study Verification

You have completed this activity when the instructor has completed the presentation on Life of a Packet.

Overview

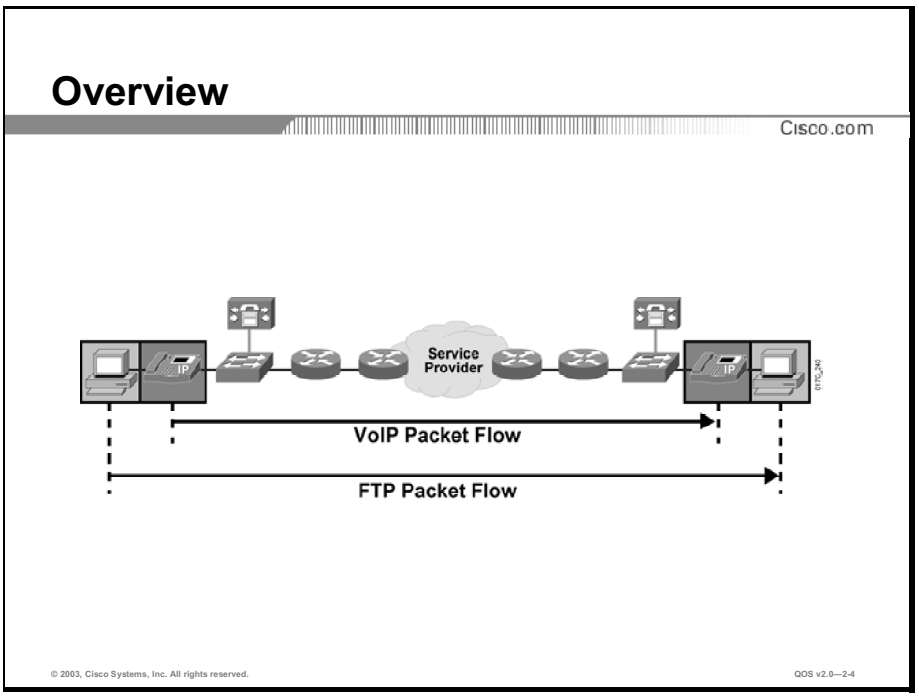

The case study follows two packets—one at a time—as they traverse an IP QoS-enabled network.

The first packet, Packet 1, is a high-priority VoIP packet that will receive preferential treatment as it moves through the network.

The second packet, Packet 2, is a low-priority FTP packet that will receive deferential treatment as it moves through the network.

In this case study, a QoS peering relation between the enterprise and the service provider is assumed. The service provider, in this case study, will recognize and act upon QoS classifications made by the enterprise customer. The relationship shows how QoS can be effectively honored across an enterprise and service provider boundary.

Life of a High-Priority (VoIP) Packet

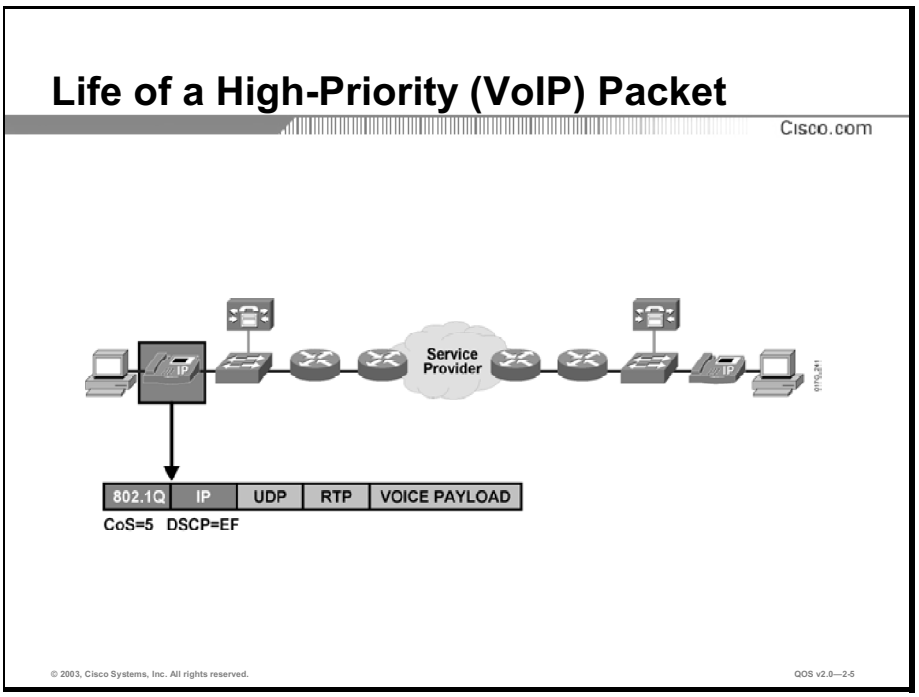

As it begins its life in the IP Phone, the VoIP packet is immediately marked with both:

- Layer 2 – 802.1Q CoS = 5 (highest priority on an Ethernet LAN)
- Layer 3 – DSCP = EF (highest priority in an IP network)

Note: The 802.1Q standard is an IEEE specification for implementing virtual LANs (VLANs) in Layer 2 switched networks. 802.1Q and its use in QoS will be discussed further in the "Classification and Marking" module in this course.

With the frame marked at CoS = 5 and DSCP = EF, this frame should receive priority treatment every time it encounters any QoS mechanism in the network.

Note: A UDP header is used for voice packets rather than TCP.

Notice that an RTP header has been added because this is a voice packet. RTP helps synchronize real-time transmissions such as voice by time-stamping packets so that they can be resynchronized at the receiving end. This helps minimize jitter.

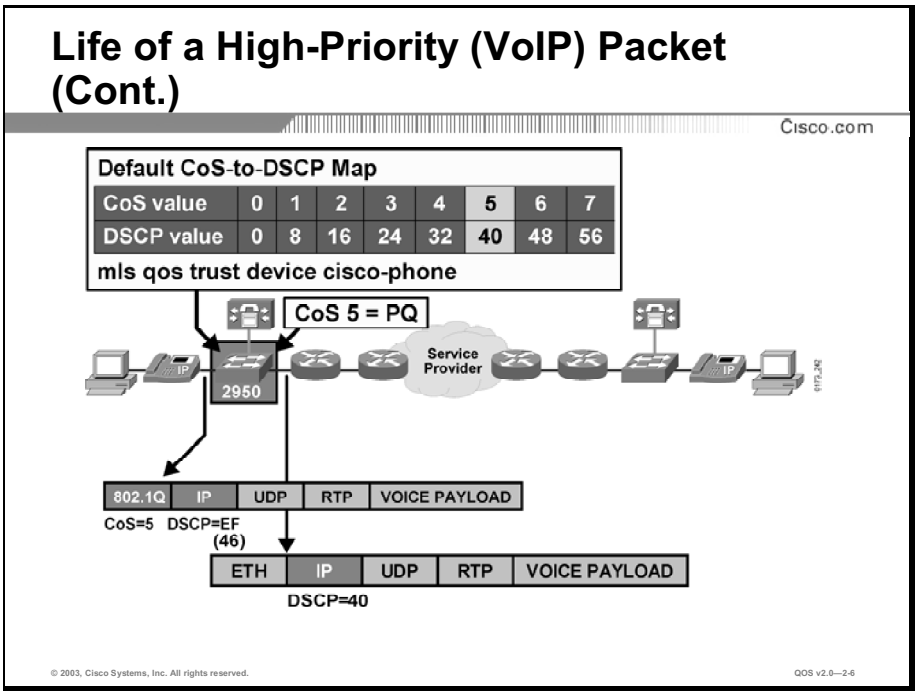

In the Cisco 2950 edge switch, the CoS = 5 means to treat the frame with PQ. This means that you should move the frame before any other frames with lower CoS.

The default CoS-to-DSCP mapping is set to recognize the CoS = 5 as DSCP = 40. To a Cisco 2950 switch, the means EF. The EF value is 46 on input to the switch as set by the IP Phone. Because the default CoS to DSCP marking is CoS 5 = DSCP 40 in the Cisco 2950 switch (not 46), DSCP is set to 40 on output.

When the frame arrives at the 2950 it is instantly recognized as a high-priority frame because of the CoS = 5 and is immediately enqueued in the high-priority, no-delay queue. Because the switch recognizes the frame as a CoS = 5, it re-marks the DSCP field to 40.

Default CoS-to-DSCP Mapping in Cisco 2950 Switch

COS Value	DSCP Value
0	0
1	8
2	16
3	24
4	32
5	**40**
6	48
7	56

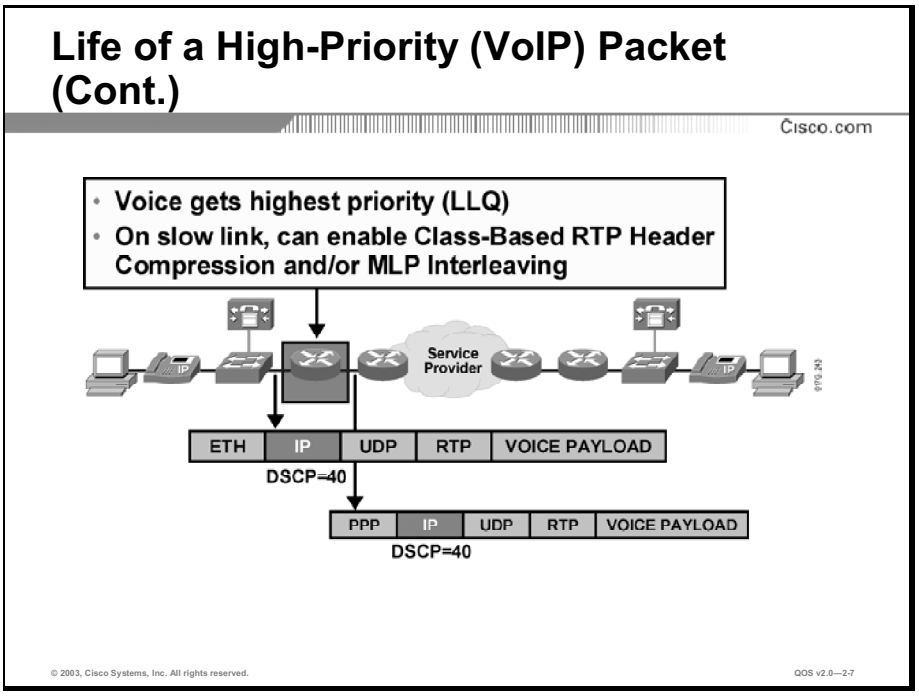

When the packet hits the edge router, the router recognizes the packet as a voice packet due to the DSCP = 40 setting (as was set by the Cisco 2950 switch).

The packet is immediately dispatched ahead of any non-voice packets using LLQ. LLQ is designed to provide instant dispatch of voice packets ahead of data while carefully managing the dispatch of data.

If the link to the service provider is a relatively slow link, then both header compression (in this case, class-based RTP header compression) and LFI would be employed to improve the bandwidth efficiency of the link.

If the WAN link were a Frame Relay link, the packet would use Frame Relay Traffic Shaping (FRTS) and FRF.12.

Note: Both of these technologies are explained further in the “Traffic Shaping and Policing” module in this course.

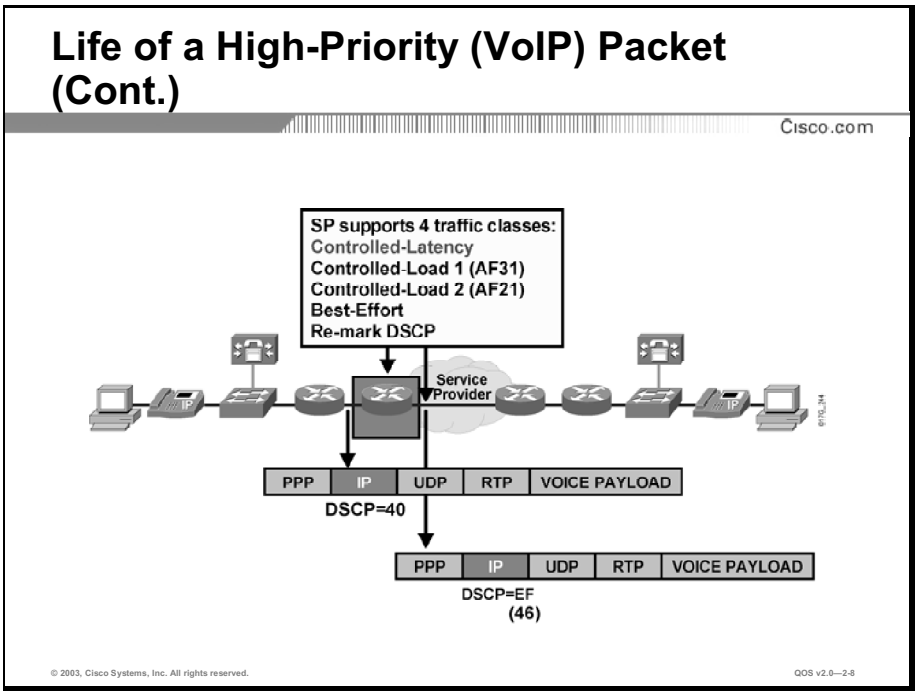

When the packet arrived at the service provider, the service provider would reclassify the packet to fit within the service provider QoS classification policy.

In this case, the service provider has defined four traffic classes:

- Real-Time (EF)
- Gold (CS 4)
- Silver (CS3)
- Best-Effort

In this case study, the service provider is providing IP QoS service level agreement (SLA) for the Real-Time, Gold, Silver, and Best-Effort traffic class. The service provider is mapping the enterprise customer QoS classifications into the service provider four defined traffic classes.

The service provider router recognizes that the packet as a high-priority voice packet and assigns the packet to the Real-Time EF class. The packet is re-marked to DSCP = 46 to fit the service provider classification conventions and sent on its way as a member of the Real-Time class.

Note: The FIFO represents the queuing mechanism that is being used on the output interface of the service provider router.

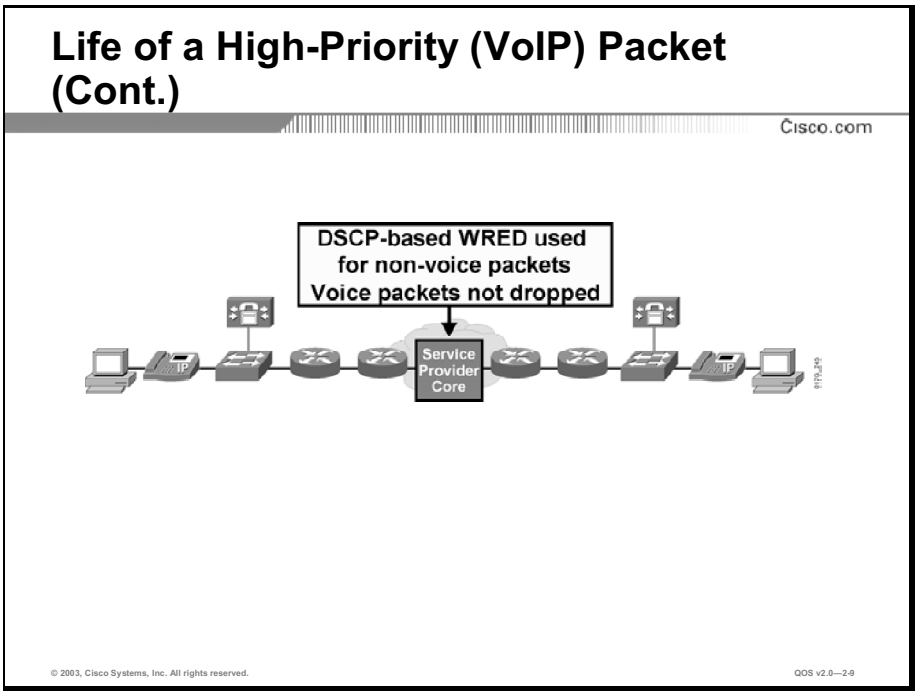

In the service provider core network, the packet will move along with minimal delay using EF.

The key congestion avoidance technology—WRED—is used in the service provider network. WRED will ensure that lower-priority packets are dropped to ensure that priority packets make their way quickly through the network.

Because the voice packet is marked as EF—the service provider Real-Time class—WRED should have no impact on the packet. QoS policy for the service provider should be not to drop voice packets so that WRED would not be applied to packets identified as Real-Time.

The packet will almost certainly not be dropped and will encounter absolute minimal delay.

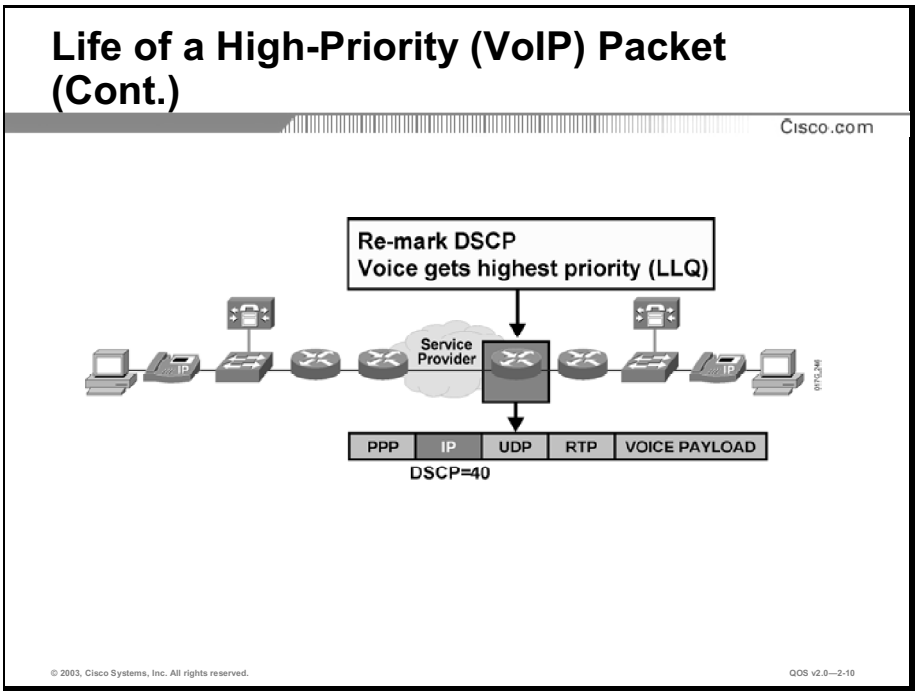

While the packet traversed the service provider network, the DSCP was marked as 46 so that the packet would be immediately dispatched as a member of the service provider Real-time class.

But the packet was marked DSCP = 40 by the customer before it entered the service provider network. At the edge of the service provider network, the DSCP is re-marked to “40” to match the classification scheme being used by the enterprise customer.

The packet is dispatched immediately using the LLQ method that always provides absolute priority to voice packets.

Life of a High-Priority (VoIP) Packet (Cont.)

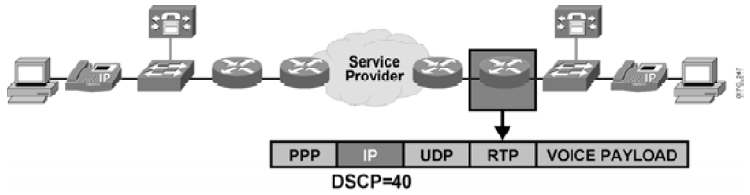

Upon arriving at the enterprise network router, the voice packet is sent out the LAN interface toward the switch using FIFO queuing on the LAN interface.

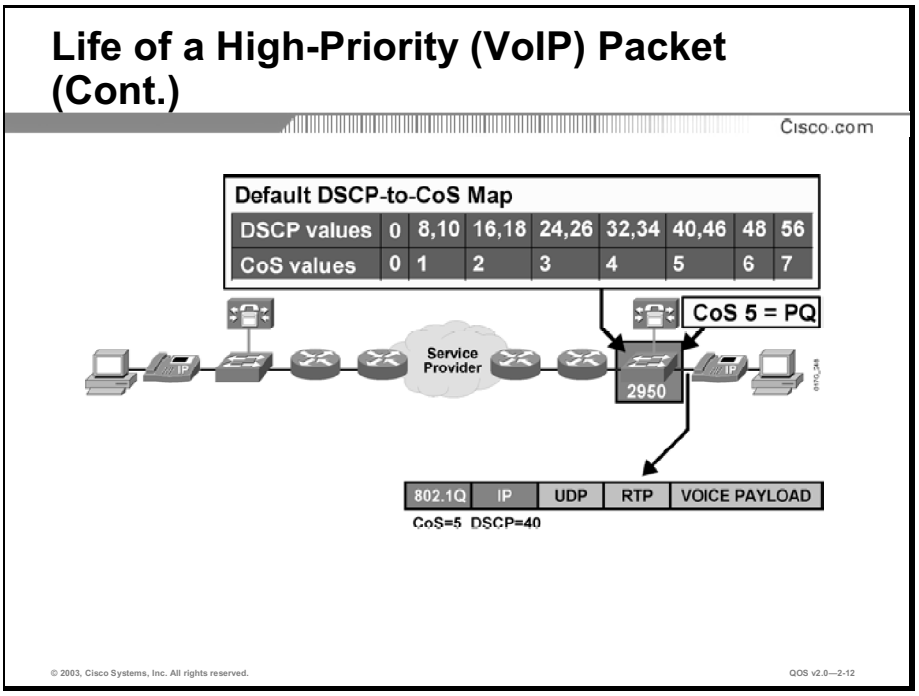

Upon arriving at the Cisco 2950 switch, the DSCP-to-CoS mappings are used to recognize the frame as a high-priority voice frame and the Layer 2 priority is set to CoS = 5. The frame jumps ahead of any non-voice frame and is immediately dispatched to the PQ.

Life of a High-Priority (VoIP) Packet (Cont.)

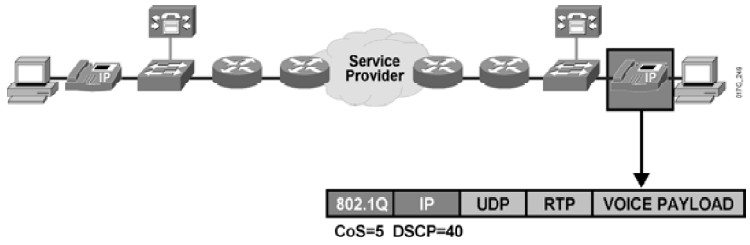

The packet finally arrives at the receiving IP Phone.

The RTP header is used to ensure that the packet is synchronized correctly with other packets from the same voice flow and the voice payload is delivered.

Life of a Low-Priority (FTP) Packet

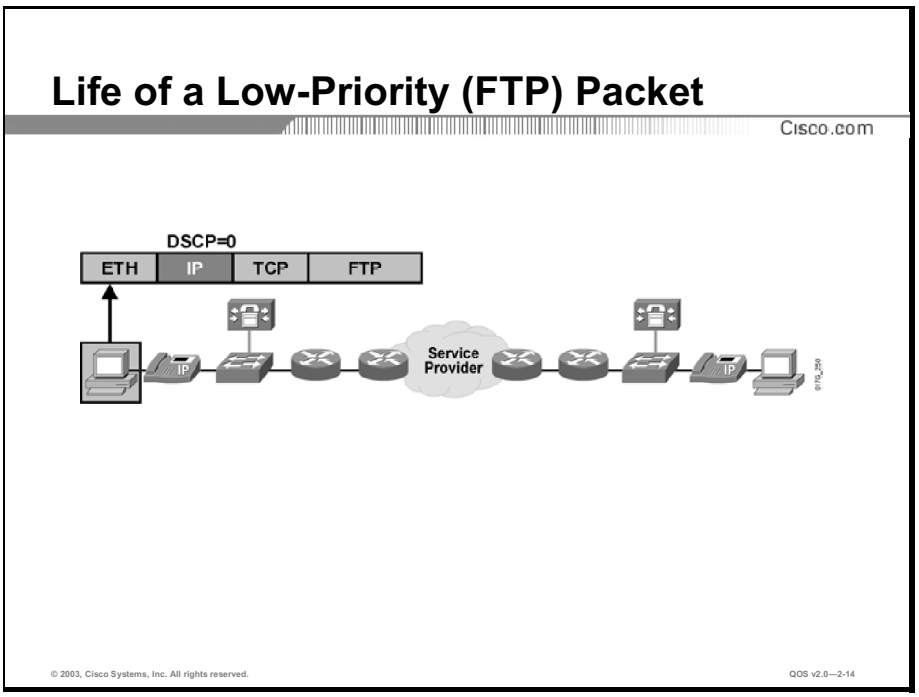

The low-priority FTP packet begins life as a very low-priority DSCP = 0.

Should the user application in the host attempt to mark the packet as a high-priority—DSCP = 46 (EF)—packet, the IP Phone would recognize that the packet was not voice and overwrite the packet DSCP priority with a lower priority.

Note:	The FTP packet is using TCP rather than UDP (which was used by the voice packet).

Life of a Low-Priority (FTP) Packet (Cont.)

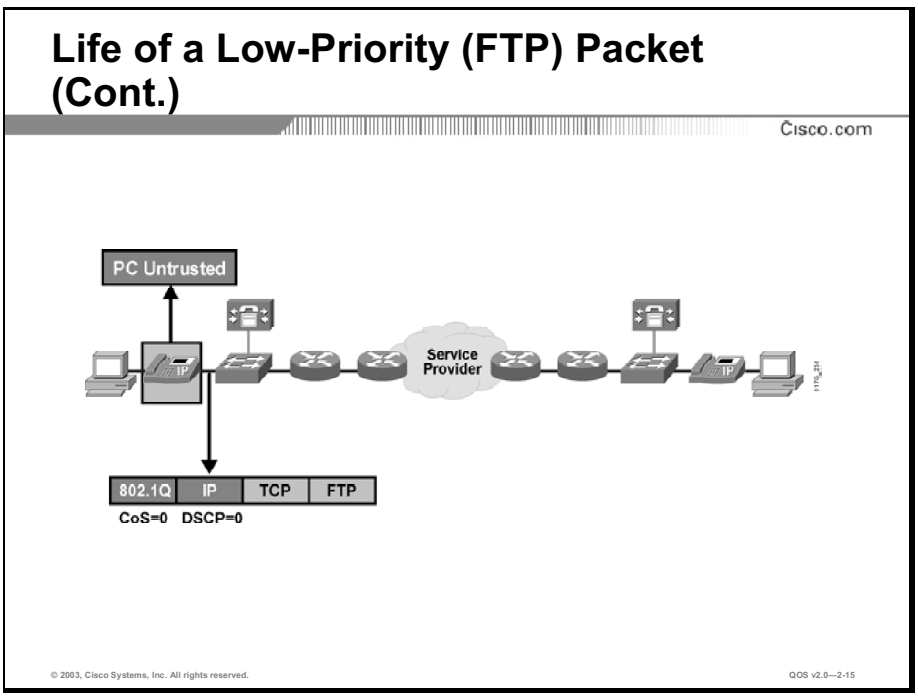

All traffic arriving from the workstation attached to the IP Phone is set to be “untrusted.”

As a result, the IP Phone will not accept any marking done by the workstation and will re-mark all DSCP values from the workstation to DSCP = 0 and set the CoS = 0. This ensures that the voice traffic generated by the IP phone will always receive priority treatment over any traffic generated by the workstation.

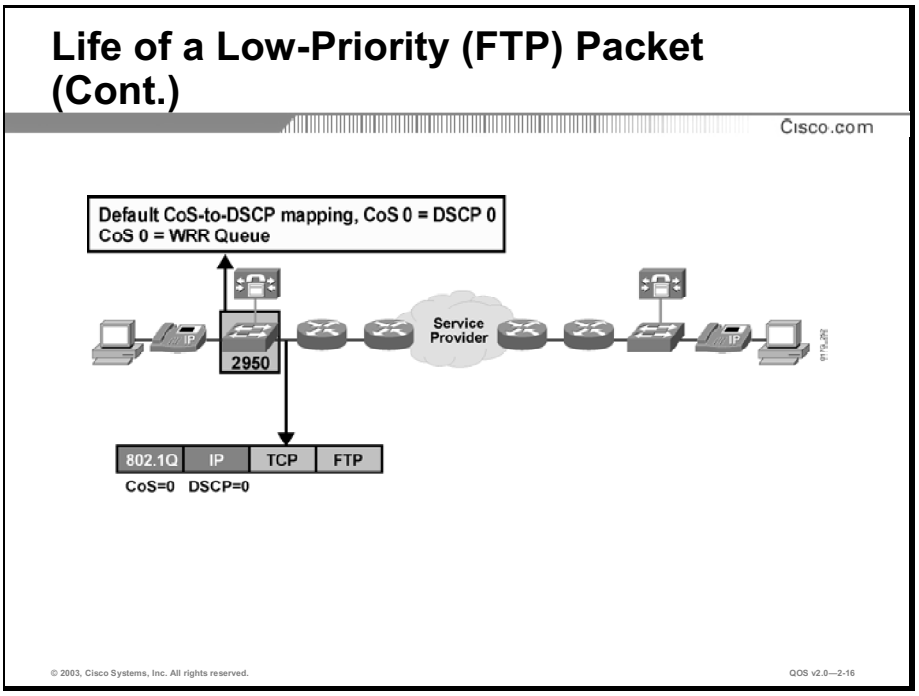

In the Cisco 2950 switch, the CoS-to-DSCP mapping would be used to map the CoS value of the packet to the switches DSCP equivalent. In the case of the FTP frame, the DSCP = 0 matches the CoS = 0, so the frame DSCP value would not change.

The switch congestion management technology—WRR—would dispatch the frame, but not until all high-priority voice frames had been dispatched. (WRR is explained further in the "Congestion Management" module in this course.)

Default CoS-to-DSCP Mapping in Cisco 2950 Switch

COS Value	DSCP Value
0	**0**
1	8
2	16
3	24
4	32
5	40
6	48
7	56

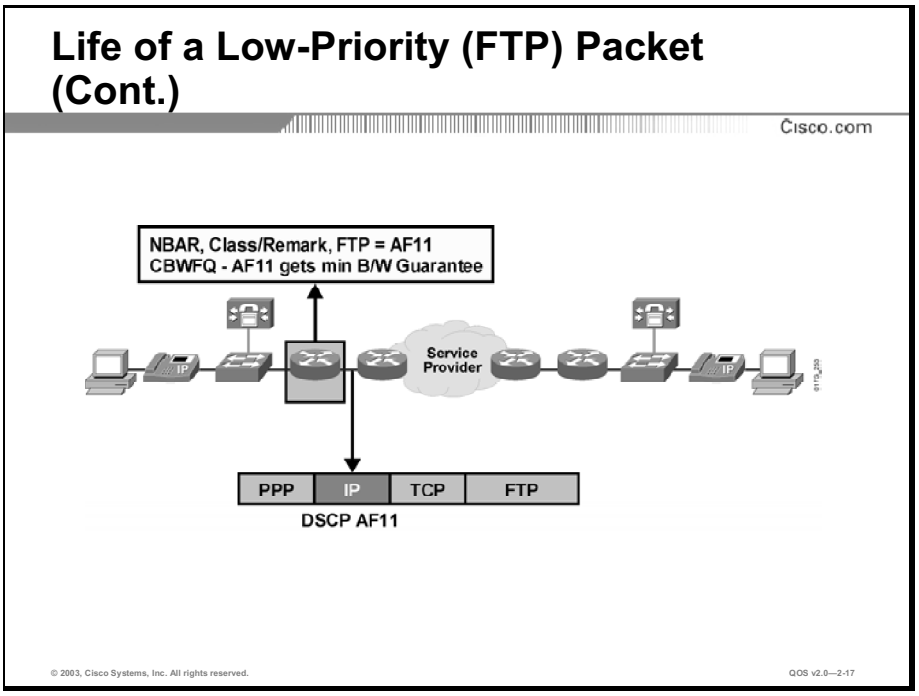

In the enterprise router, a classification technology, NBAR, would recognize the packet as an FTP packet and assign the packet a DSCP = “001010” = AF11.

With a DSCP = AF11, the packet would then be dispatched as a low-priority class packet by CBWFQ. CBWFQ is the component of LLQ that carefully manages the dispatch of data traffic.

The AF11 class is given a minimum guarantee of bandwidth. If the link to the service provider were congested, the packet would have a good probability of being dropped to ensure that higher-priority packets are not delayed.

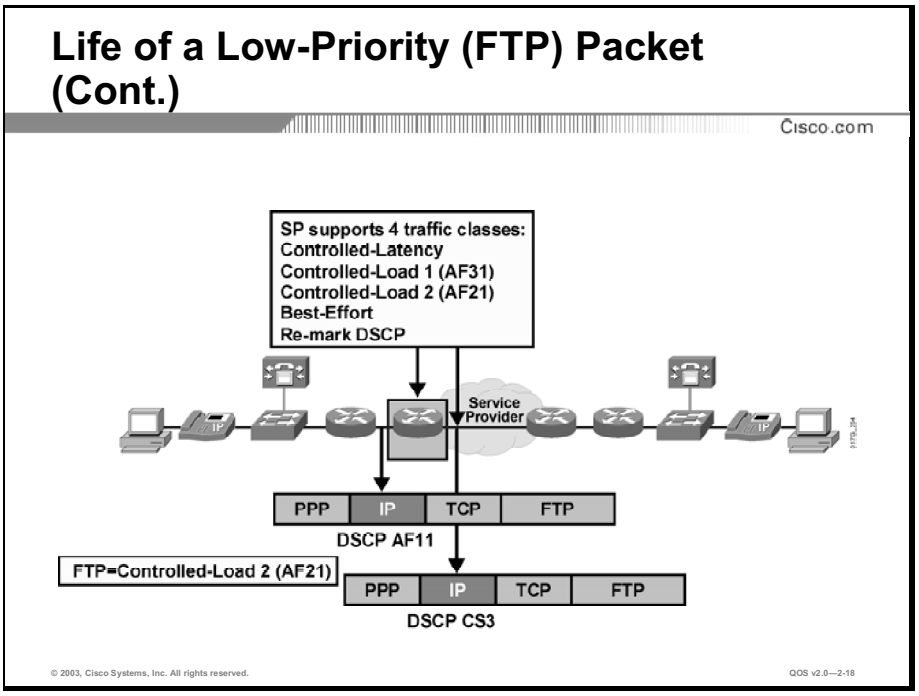

In this case study, the service provider is providing IP QoS SLA for the Real-Time, Gold, Silver, and Best-Effort traffic classes. The service provider is mapping the enterprise customer QoS classifications into the service provider four defined traffic classes.

Upon arriving at the service provider network, the packet would be identified as an FTP packet and assigned to the Silver class (CS3).

Life of a Low-Priority (FTP) Packet (Cont.)

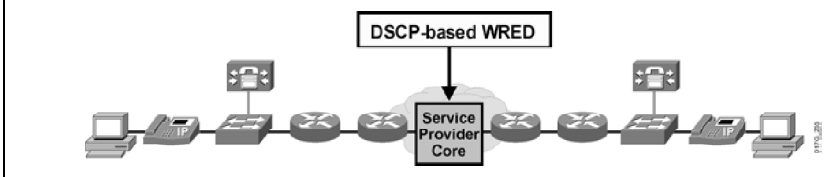

The packet traverses the service provider core marked as a Silver class (CS3) packet.

While in the service provider core network, the FTP packet would have a much better probability of being dropped by WRED than the voice packet.

Life of a Low-Priority (FTP) Packet (Cont.)

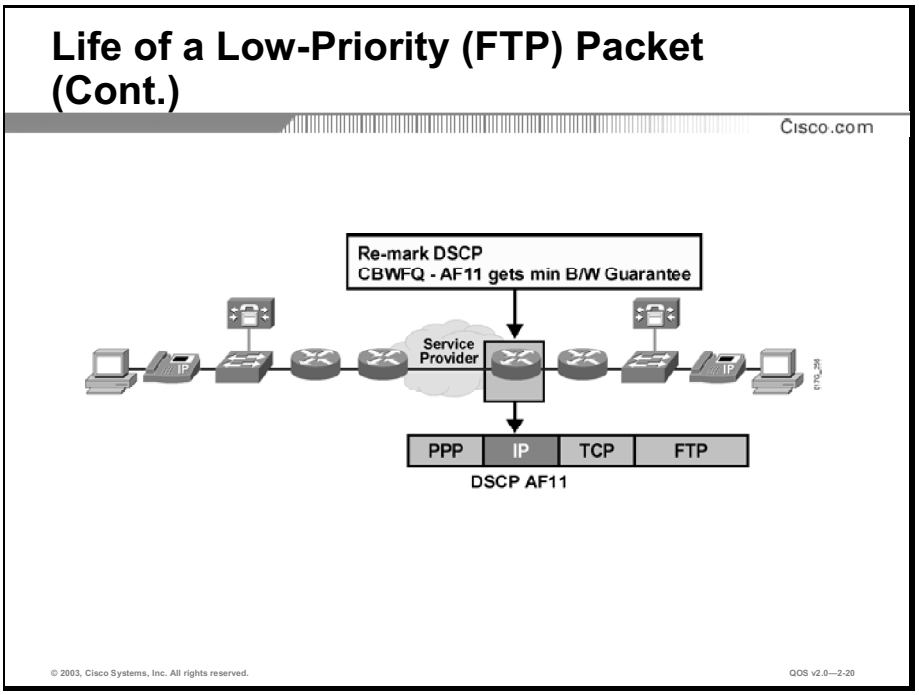

Before the packet entered the service provider network, it was marked DSCP = AF11, which fit the classification scheme used by the enterprise customer. As the packet leaves the service provider network, the packet is re-marked to DSCP = AF11 for the enterprise customer.

The AF11 class is given a minimum guarantee of bandwidth.

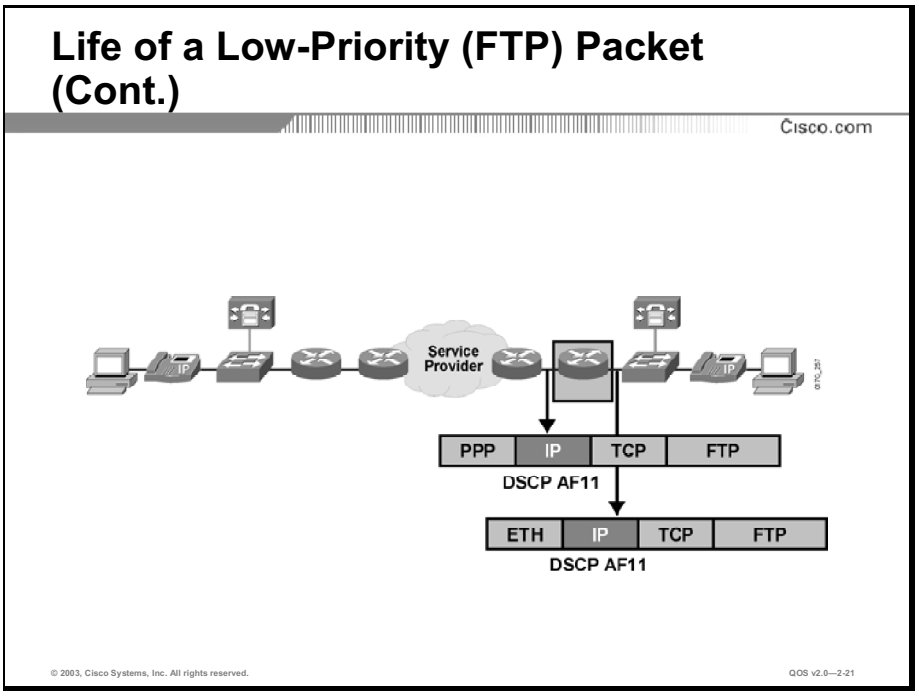

As the packet reenters the enterprise network, it is recognized as an AF11 class packet and is passed through the enterprise router without being re-marked.

The FTP packet is sent out the LAN interface toward the switch using FIFO queuing on the LAN interface.

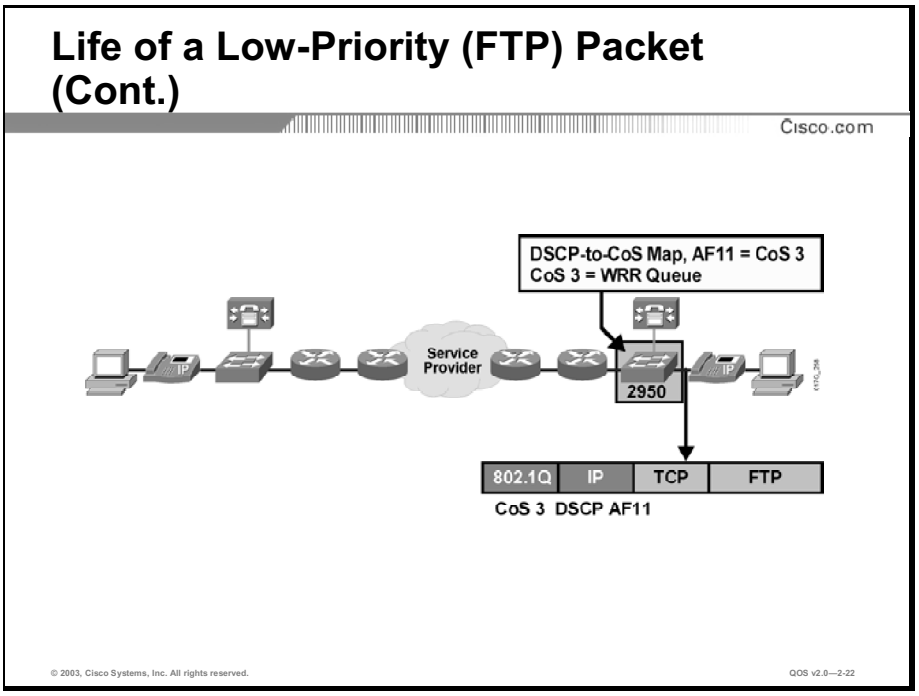

Using its DSCP-to-CoS mapping, the Cisco 2950 switch recognizes the DSCP = AF11 packet (Layer 3) as a CoS = 3 priority frame (Layer 2).

The FTP frame is treated by WRR with the CoS 3 (which can be configured to have *less weight* than CoS 4 or 5 frames, but *more weight* than CoS 1 or 2 frames).

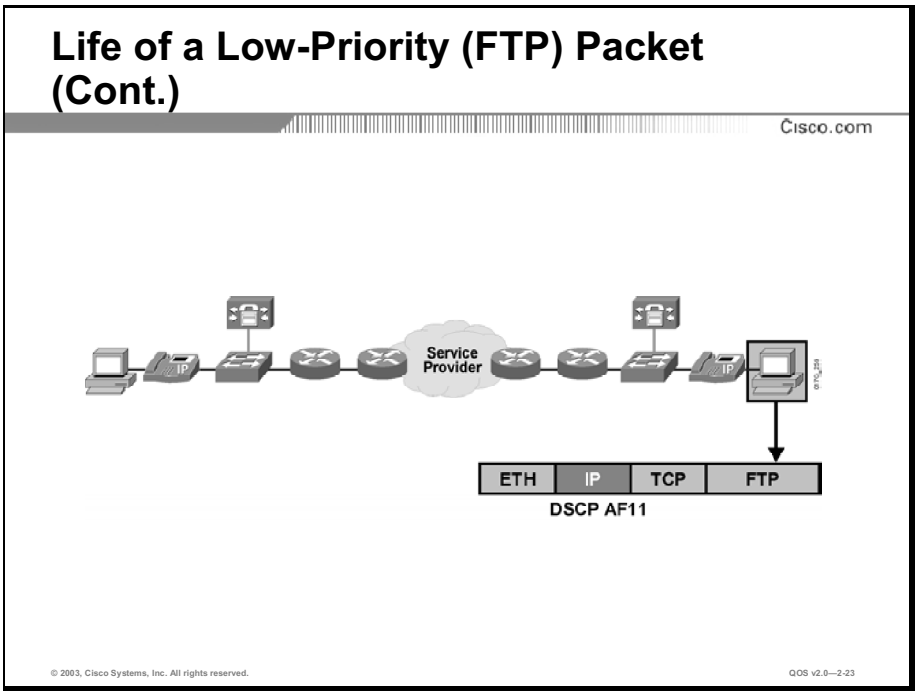

The FTP packet finally arrives at the destination host and the payload is delivered.

If the packet had been dropped at any point along the way, TCP would recognize that fact, and request retransmission of the packet.

Summary

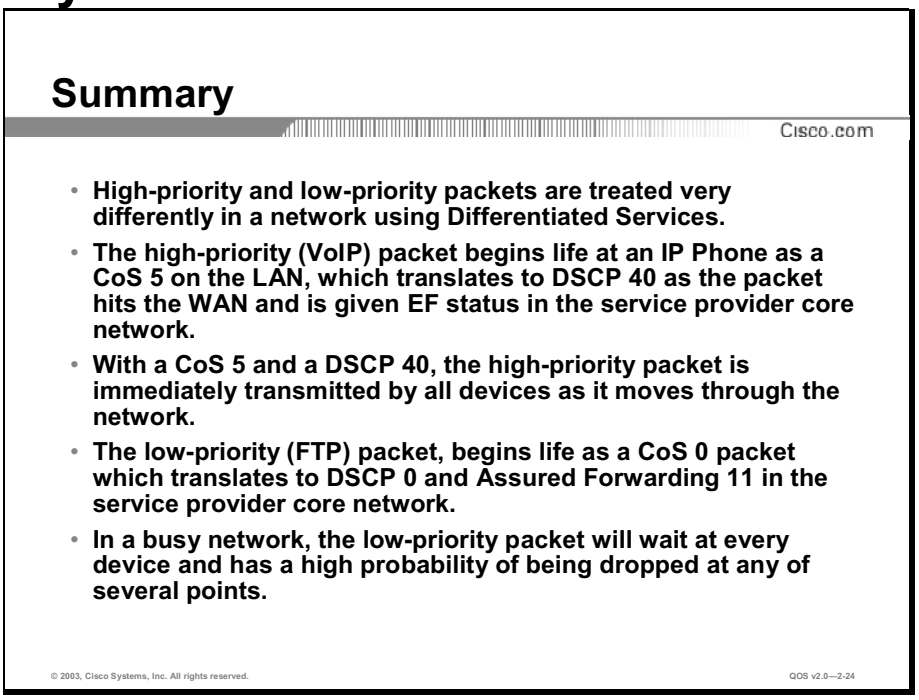

Next Steps

For the associated lab exercise, refer to the following section of the course Lab Guide:

- Lab Exercise 2-1: QoS Lab Setup and Initialization
- Lab Exercise 2-2: Baseline QoS Measurement

Module Assessment

Overview

Use this assessment to test what you learned in this module. The correct answers and solutions are found in the Module Assessment Answer Key.

Quiz: The Building Blocks of IP QoS

Complete the Quiz to assess what you have learned in the module.

Objectives

This activity tests your knowledge on how to meet these objectives:

- Correctly match a list of QoS actions to one or more of the three models for implementing QoS on a network
- Describe the Differentiated Services model and explain how it can be used to implement QoS in that network
- Correctly match a list of QoS actions to mechanisms for implementing QoS and identify where in a network the different QoS mechanisms are commonly used
- Correctly identify the QoS status of packets as they pass through various points in the network

Instructions

Complete these steps:

Step 1 Answer all questions in this quiz by selecting the best answer(s) to each question.

Step 2 Verify your results against the answer key located at the end of this section.

Step 3 Review the topics in this module that relate to the questions that you answered incorrectly.

Q1) Which of the models for implementing QoS is least scalable?

A) Best-Effort
B) Integrated Services
C) Differentiated Services
D) Quantitative Services

Q2) Which three IP QoS mechanisms work together to provide a set of complete integrated services on a network? (Choose three.)

A) Weighted RED (WRED)
B) Weighted Fair Queuing (WFQ)
C) Generic Traffic Shaping (GTS)
D) Resource Reservation Protocol (RSVP)

Q3) What is the most important advantage of Differentiated Services over other QoS models?

A) high scalability
B) many service levels
C) guaranteed service
D) deterministic delays
E) advanced queuing mechanisms

Q4) Services are provided to which entities in the Differentiated Services model?

A) frames
B) packets
C) applications
D) classes of traffic

Q5) How many bits is the DSCP field of the IP header?

A) 3
B) 4
C) 6
D) 8

Q6) What PHB would be indicated if the DSCP was equal to 46 (101110)?

A) default PHB
B) selector PHB
C) Assured Forwarding PHB
D) Expedited Forwarding PHB

Q7) Which Assured Forwarding Class and what drop probability would be indicated if the DSCP was equal to “100100?”

A) AF Class 1 and medium
B) AF Class 4 and medium
C) AF Class 1 and high
D) AF Class 4 and high

Q8) If DSCP for packets A, B, C, and D was respectively set to “101000”, “011000”, “111000”, and “001000”, which packet would have the greatest probability of timely forwarding?

A) A
B) B
C) C
D) D

Q9) Match the following IP QoS mechanisms to their function.

A) congestion avoidance
B) congestion management
C) classification
D) traffic policing
E) traffic shaping
F) packet header compression

_____ 1. Drops misbehaving traffic to maintain network integrity.

_____ 2. Improves the bandwidth efficiency of a link.

_____ 3. Controls traffic by delaying bursts.

_____ 4. Discards specific packets based on markings.

_____ 5. Identifying and splitting of traffic.

_____ 6. Prioritizing, protection, and isolation of traffic based on markings.

Q10) Which of the following IP QoS mechanisms is used on both input and output interfaces?

A) classification
B) traffic policing
C) traffic shaping
D) congestion management

Scoring

You have successfully completed the quiz for this lesson when you earn a score of 80 percent or better.

Module Assessment Answer Key

Q1) B

Relates to: Models for Implementing QoS

Q2) A, B, D

Relates to: Models for Implementing QoS

Q3) A

Relates to: The Differentiated Services Model

Q4) D

Relates to: The Differentiated Services Model

Q5) C

Relates to: The Differentiated Services Model

Q6) D

Relates to: The Differentiated Services Model

Q7) B

Relates to: The Differentiated Services Model

Q8) C

Relates to: The Differentiated Services Model

Q9) 1 = D

2 = F

3 = E

4 = A

5 = C

6 = B

Relates to: IP QoS Mechanisms

Q10) B

Relates to: IP QoS Mechanisms

Module Summary

This topic summarizes the key points discussed in this module.

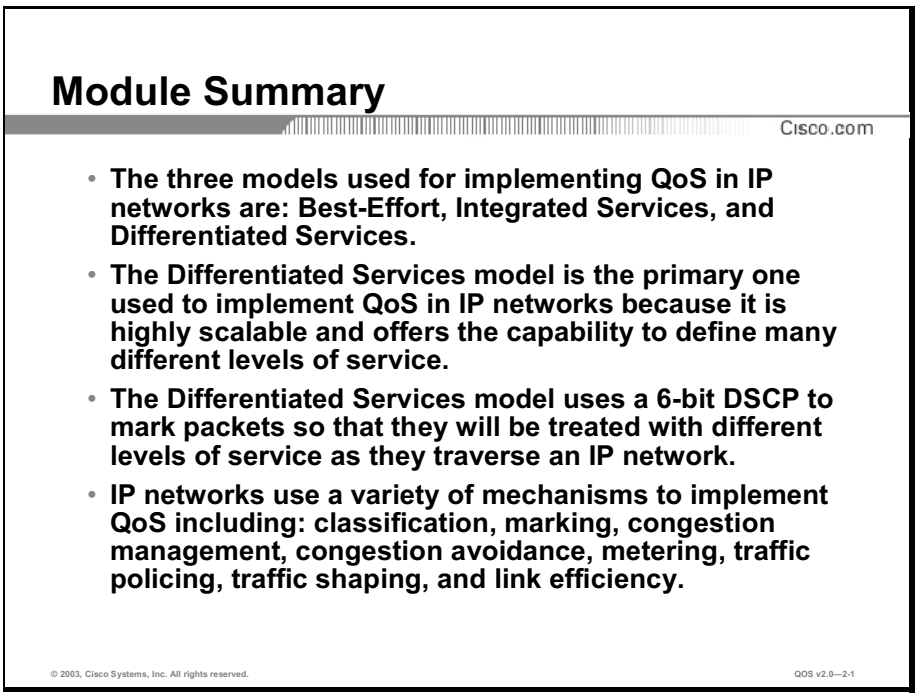

Three different models exist for implementing QoS on a network. The Best-Effort model was designed for best-effort, no-guarantee delivery of packets. This model is still predominant on the Internet today. The Integrated Services model was introduced to supplement the best-effort delivery by setting aside some bandwidth for applications that require bandwidth and delay guarantees. The Integrated Services model expects applications to signal their requirements to the network. The Differentiated Services model was added to provide greater scalability in providing QoS to IP packets. The main difference between the Integrated Services and Differentiated Services models is that with the Differentiated Services model, the network recognizes packets (no signaling is needed) and provides the appropriate services to them. IP networks of today can use all three models at the same time.

Differentiated Services is a multiple-service model that is designed to satisfy various QoS requirements. With Differentiated Services, the network tries to deliver a particular kind of service based on the QoS specified by each packet. This specification can occur in different ways, for example, using the Differentiated Services Code Point in IP packets or source and destination addresses. The network uses the QoS specification of each packet to classify, shape, and police traffic and to perform intelligent queuing.

IP networks use a variety of mechanisms to implement QoS including: classification, marking, congestion management, congestion avoidance, metering, traffic policing, traffic shaping, and link efficiency. IP QoS mechanisms are used to implement a coordinated QoS policy in devices throughout the network. The moment an IP packet enters the network, it is classified and usually marked with its class identification. From that point on, the packet is treated by a variety of IP QoS mechanisms according to the packet classification. Depending upon the mechanisms it encounters, the packet could be expedited, delayed, compressed, fragmented, or even dropped.

Module 3

Introduction to Modular QoS CLI and AutoQoS

Overview

Quality of Service (QoS) configurations can be complex. In Cisco IOS software configurations, there are many different QoS mechanisms, many of which have similar features. Because there are also many options available for providing QoS to different traffic types, it can easily become an overwhelming effort to deploy QoS end-to-end in a network infrastructure. Fortunately, Cisco Systems has unified QoS configuration by separating the different components of a QoS policy into different configuration modules. It is these modules that comprise the Cisco Modular QoS command-line interface (CLI [MQC]) that allow network administrators and network implementers to more easily deploy QoS. MQC configurations are consistent for different QoS mechanisms and are therefore easier to learn, deploy and troubleshoot.

There are cases, however, when some customers do not want to be concerned with the specifics of QoS configuration. These customers would prefer to enable QoS in a global fashion with a single command and allow the Cisco IOS router and switch to automate the required complex QoS configuration. For those customers, Cisco has developed AutoQoS. This module introduces MQC and AutoQoS as configuration methods for implementing QoS. This module will also serve as the foundation for more advanced MQC configurations that include additional QoS features and techniques.

Module Objectives

Upon completing this module, you will be able to explain the use of MQC and AutoQoS, and to implement QoS on the network.

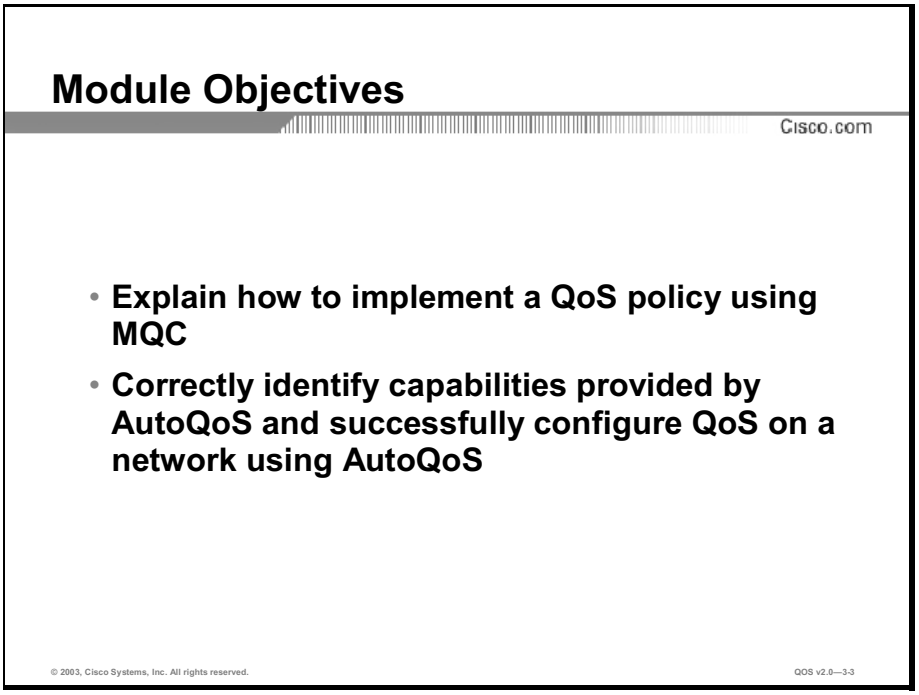

Module Outline

The outline lists the components of this module.

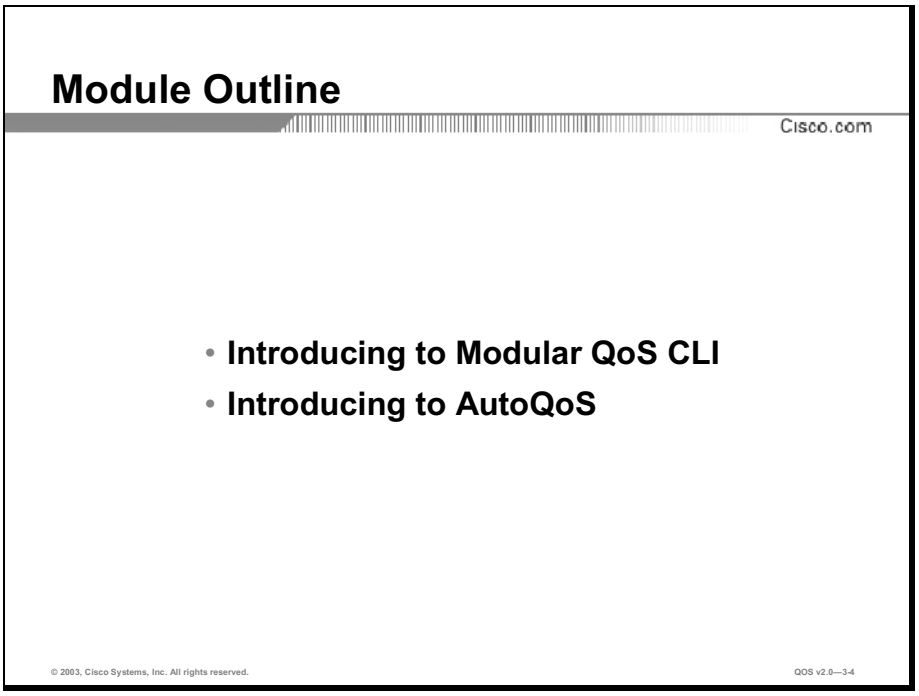

Introducing Modular QoS CLI

Overview

This chapter explains how to implement QoS policies using the MQC.

Relevance

MQC is one of the two key methods recommended for implementing QoS on a network. MQC may be the best choice for implementing a large, finely tuned network incorporating voice and video applications.

Objectives

Upon completing this lesson, you will be able to explain how to implement a QoS policy using MQC. This includes being able to meet these objectives:

- Explain how to implement a given “QoS policy” using MQC
- Differentiate between class maps, policy maps, and service policies
- Describe how a class map is used to define a class of traffic
- Describe the Cisco IOS MQC commands required to configure and monitor a class map
- Describe how a policy map is used to assign a QoS policy to a class of traffic
- Describe the Cisco IOS MQC commands required to configure and monitor a policy map
- Explain how a service policy is assigned to an interface
- Describe the MQC commands used to attach a service policy to an interface

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Basic understanding of the Cisco IOS command-line interface

Outline

The outline lists the topics included in this lesson.

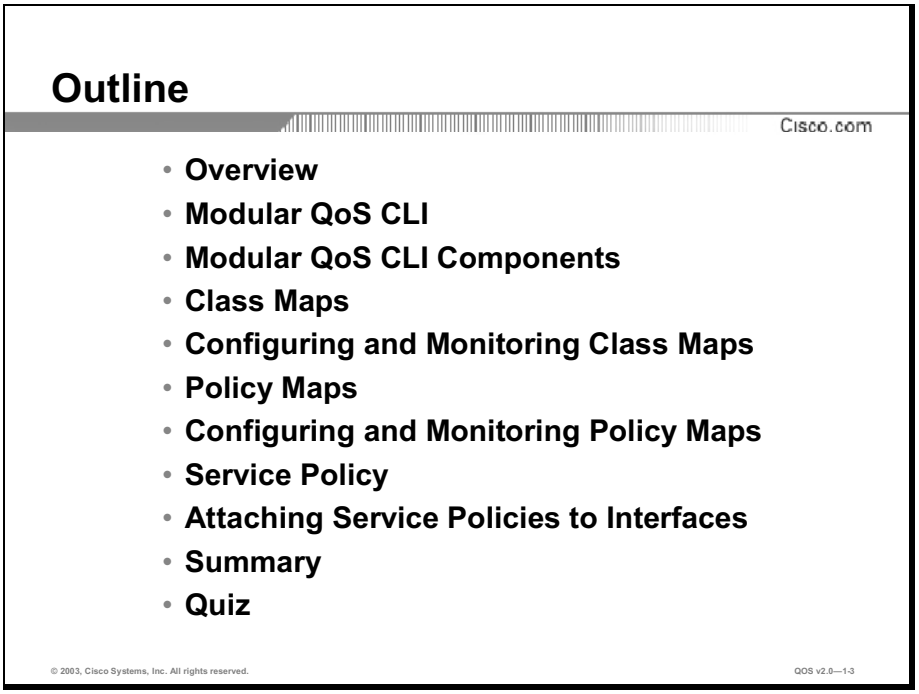

Modular QoS CLI

This topic describes the MQC method for implementing QoS on a network.

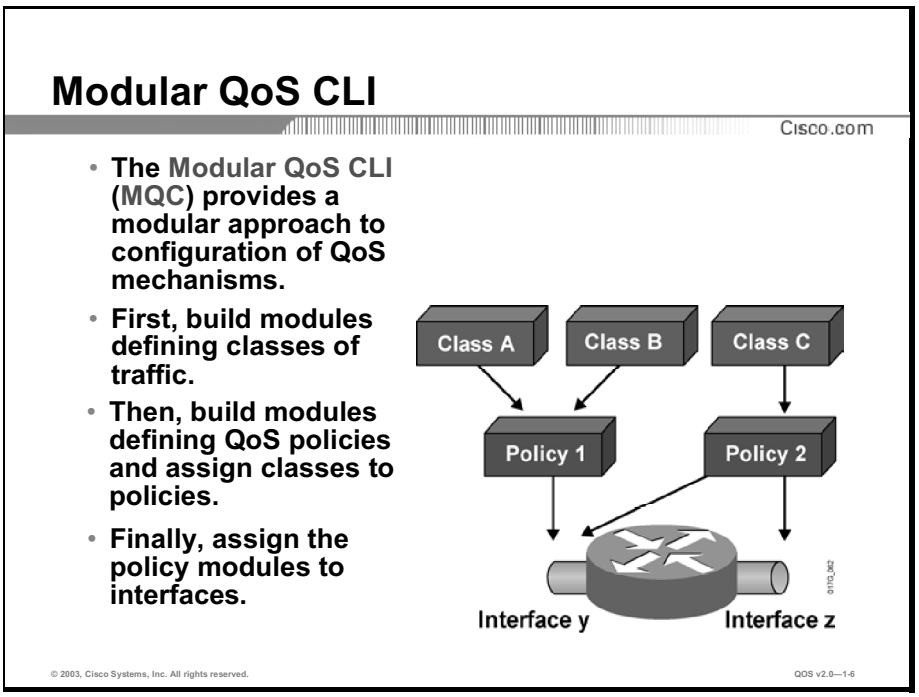

The MQC was introduced to allow any supported classification to be used with any QoS mechanism.

The separation of classification from the QoS mechanism allows new Cisco IOS versions to introduce new QoS mechanisms and reuse all available classification options. On the other hand, old QoS mechanisms can benefit from new classification options.

Another important benefit of the MQC is the reusability of configuration. MQC allows the same QoS policy to be applied to multiple interfaces. The MQC, therefore, is a consolidation of all the QoS mechanisms that have so far only been available as standalone mechanisms.

Example: Advantages of Using MQC

Configuring committed access rate (CAR), for example, required entire configurations to be repeated between interfaces and time-consuming configuration modifications. MQC allows the same QoS policy to be applied to multiple interfaces.

Modular QoS CLI Components

This topic describes the three steps involved in implementing a QoS policy using MQC.

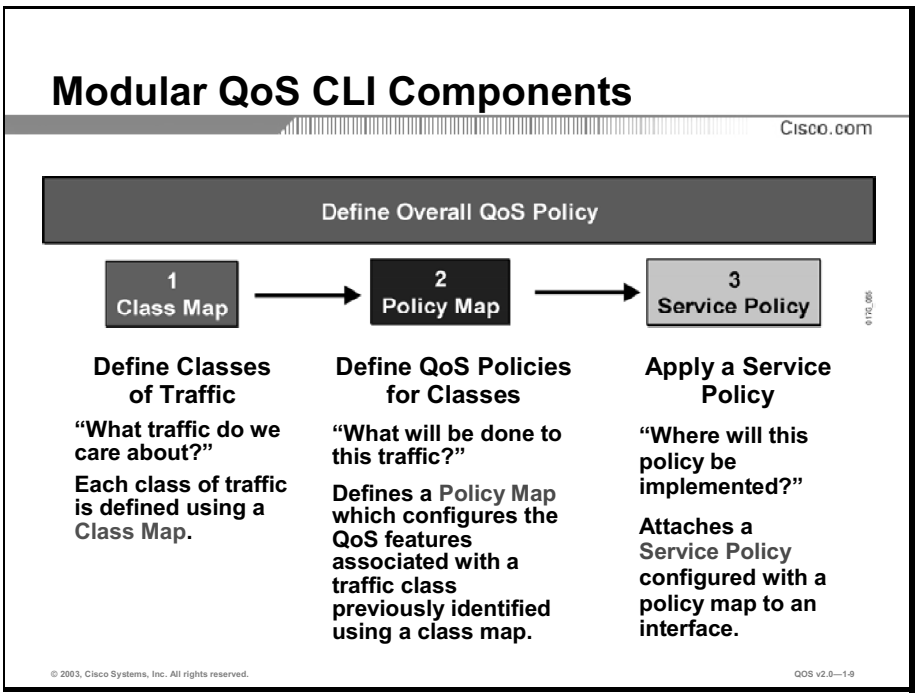

Implementing QoS by using the MQC consists of three steps:

- First, configure classification by using the **class-map** command.
- Second, configure traffic policy by associating the traffic class with one or more QOS features using the **policy-map** command.
- Finally, attach the traffic policy to inbound or outbound traffic on interfaces, subinterfaces, or virtual circuits by using the **service-policy** command.

Example: Configuring MQC

Consider a network with voice telephony:

- First, classify traffic as Voice, High Priority, Low Priority, and Browser in **class-maps**.
- Second, build a single **policy-map** that defines three different traffic policies (different bandwidth and delay requirements for each traffic class): *NoDelay*, *BestService*, and *Whenever,* and assign the already defined classes of traffic to the policies. Voice is assigned to *NoDelay*. High Priority traffic is assigned to *BestService*. Both Low Priority, and Browser traffic are assigned to *Whenever*.
- Finally, assign the **policy-map** to selected router and switch interfaces.

Class Maps

This topic describes the use of Class Maps.

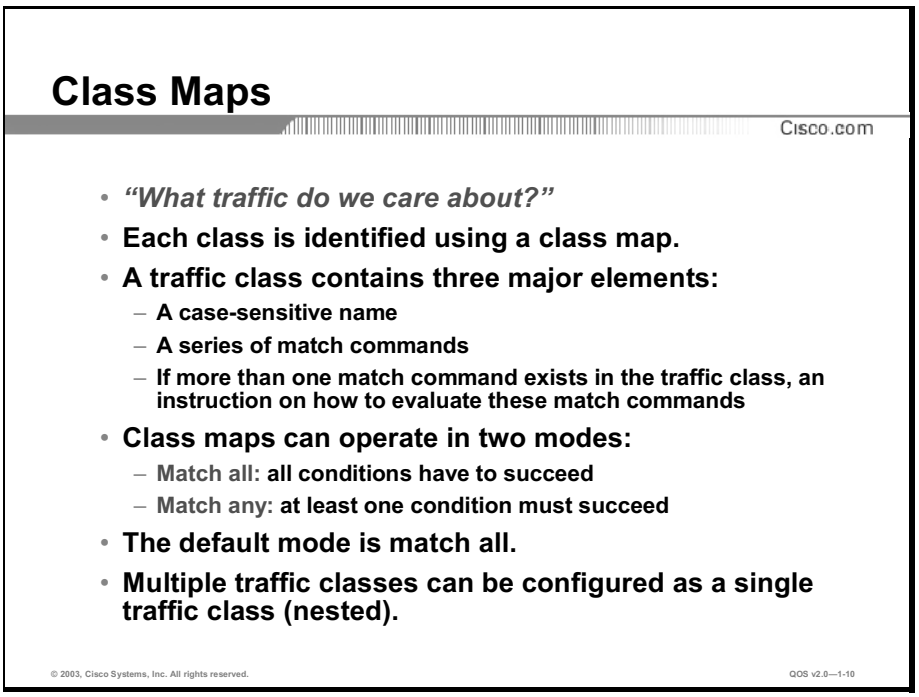

Class maps are used to create classification templates that are later used in policy maps where QoS mechanisms are bound to classes.

Routers can be configured with a large number of class maps (currently limited to 256). Each traffic policy, however, may support a limited number of classes; for example, class-based weighted fair queuing (CBWFQ) and class-based low-latency queuing (LLQ) are limited to 64 classes.

A class map is created using the **class-map** global configuration command. Class maps are identified by case-sensitive names. Each class map contains one or more conditions that determine if the packet belongs to the class.

There are two ways of processing conditions when there is more than one condition in a class map:

- **Match all:** All conditions have to be met to bind a packet to the class.
- **Match any:** At least one condition has to be met to bind the packet to the class.

The default match strategy of class maps is "match all."

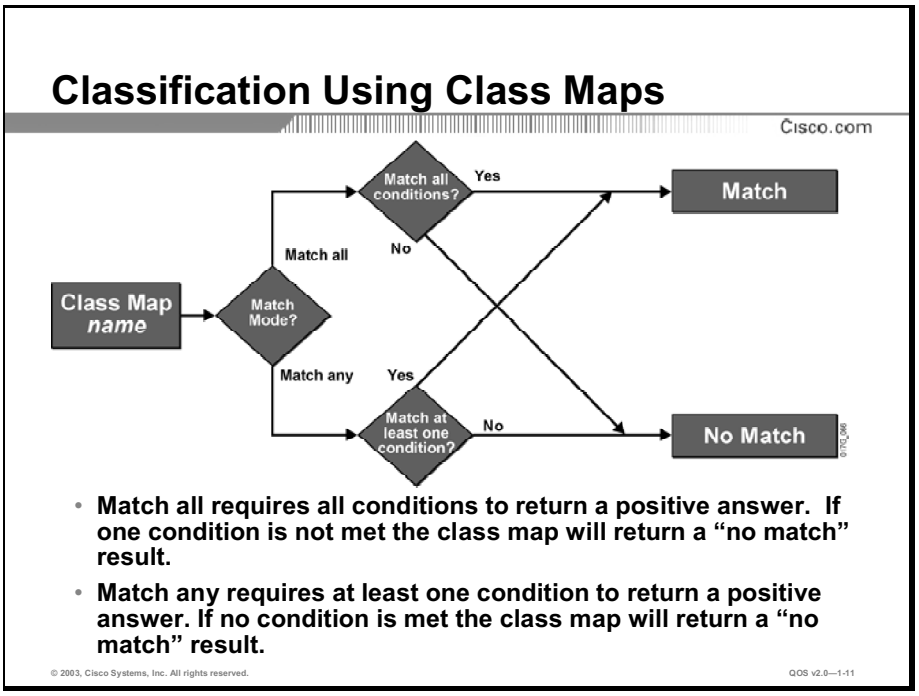

The figure illustrates the full process of determining if a packet belongs to a class (match) or not (no match).

The process goes through the list of conditions and:

- Returns a "match" result if one of the conditions is met and the match-any strategy is used
- Returns a "match" result if all conditions are met and the match-all strategy is used

If either of these conditions is not met it returns "no match."

Configuring and Monitoring Class Maps

This topic explains the commands that are necessary to configure and monitor class maps.

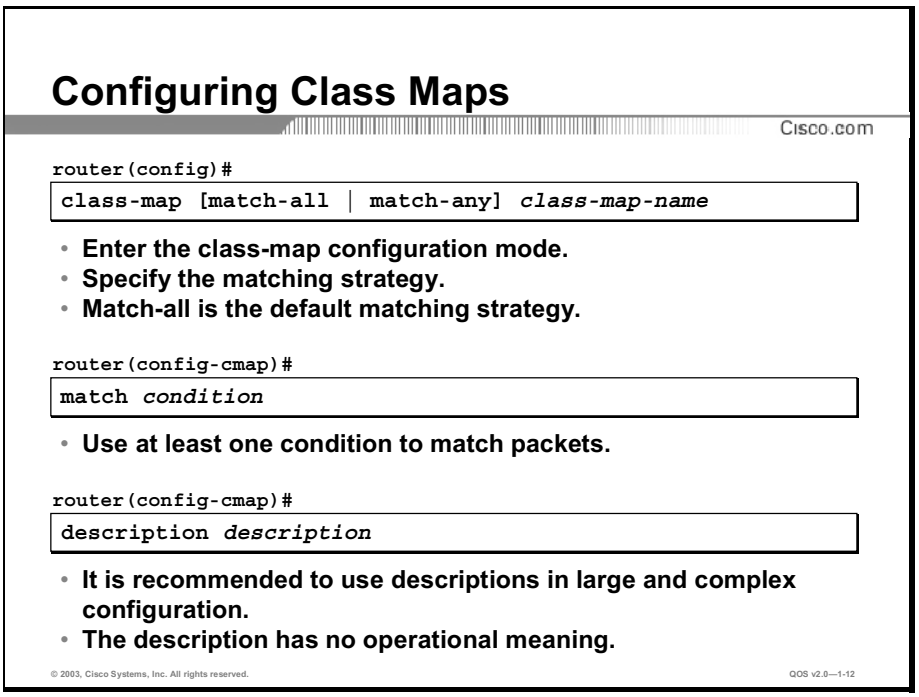

Use the **class-map** global configuration command to create a class map and enter the class map configuration mode. A class map is identified by a case-sensitive name; therefore, all subsequent references to the class map must use exactly the same name.

At least one match command should be used within the class map configuration mode (**match none** is the default).

The **description** command is used for documenting a comment about the class map.

Example: Class Map Configuration

The following example shows a traffic class configured with the **class-map match-all** command:

```
Router(config)# class-map match-all cisco1
Router(config-cmap)# match protocol ip
Router(config-cmap)# match qos-group 4
Router(config-cmap)# match access-group 101
```

If a packet arrives on a router with traffic class called cisco1 configured on the interface, the packet is evaluated to determine if it matches the IP protocol, QoS group 4, *and* access group 101. If all three of these match criteria are met, the packet matches traffic class cisco1.

Configuring Classification Using Special Options

Cisco.com

```
router(config-cmap)#
match not condition
```

- **The "not" keyword inverts the condition.**

```
router(config-cmap)#
match class-map class-map-name
```

- **One class map can use another class map for classification.**
- **Nested class maps allow generic template class maps to be used in other class maps.**

```
router(config-cmap)#
match any
```

- **The "any" keyword can be used to match all packets.**

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. The MQC does not necessarily require that users associate a single traffic class to one traffic policy. Multiple traffic classes can be associated with a single traffic policy using the **match any** command.

The **match not** command inverts the condition specified. It specifies a match criterion value that prevents packets from being classified as members of a specified traffic class. All other values of that particular match criterion belong to the class.

The MQC allows multiple traffic classes (nested traffic classes, which are also called nested class maps) to be configured as a single traffic class. This nesting can be achieved with the use of the **match class-map** command. The only method of combining match-any and match-all characteristics within a single traffic class is with the **match class-map** command.

Example: Using the match Command

The following example shows a traffic class configured with the **class-map match-any** command:

```
Router(config)# class-map match-any cisco2
Router(config-cmap)# match protocol ip
Router(config-cmap)# match qos-group 4
Router(config-cmap)# match access-group 101
```

In traffic class called cisco2, the match criteria are evaluated consecutively until a successful match criterion is located. The packet is first evaluated to determine whether IP protocol can be used as a match criterion. If IP protocol is not a successful match criterion, then QoS group 4 is evaluated as a match criterion. Each matching criterion is evaluated to see if the packet matches that criterion. When a successful match occurs, the packet is classified as a member of traffic class cisco2. If the packet matches none of the specified criteria, the packet is classified as a member of the traffic class.

Example: Nested Traffic Class to Combine match-any and match-all Characteristics in One Traffic Class

The only method of including both match-any and match-all characteristics in a single traffic class is to use the **match class-map** command. To combine match-any and match-all characteristics into a single class, a traffic class created with the **match-any** instruction must use a class configured with the **match-all** instruction as a match criterion (through the **match class-map** command), or vice versa.

The following example shows how to combine the characteristics of two traffic classes, one with match-any and one with match-all characteristics, into one traffic class with the **match class-map** command. The result of traffic class class4 requires a packet to match one of the following three match criteria to be considered a member of traffic class class4: IP protocol and QoS group 4, destination MAC address 1.1.1, or access group 2.

In this example, only the traffic class called class4 is used with the traffic policy called policy1:

```
Router(config)# class-map match-all class3
Router(config-cmap)# match protocol ip
Router(config-cmap)# match qos-group 4
Router(config-cmap)# exit
Router(config)# class-map match-any class4
Router(config-cmap)# match class-map class3
Router(config-cmap)# match destination-address mac 1.1.1
Router(config-cmap)# match access-group 2
Router(config-cmap)# exit
Router(config)# policy-map policy1
Router(config-pmap)# class class4
Router(config-pmap-c)# police 8100 1500 2504 conform-action
transmit exceed-action set-qos-transmit 4
Router(config-pmap-c)# exit
```

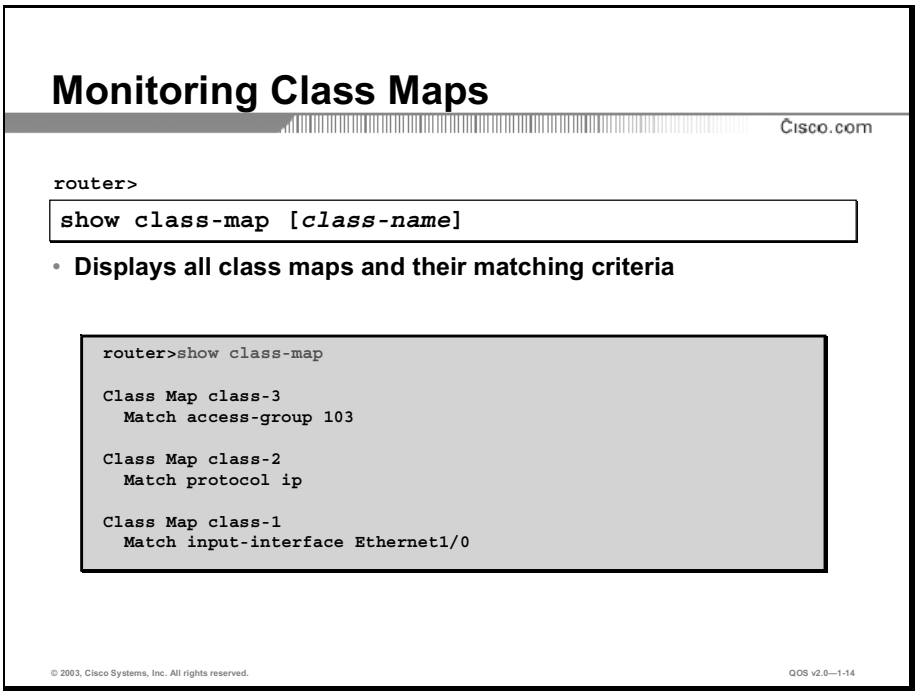

The **show class-map** command lists all class maps with their match statements.

The **show class-map** command with a name of a class map displays the configuration of the selected class map.

The example of **show class-map** in the illustration shows three class maps:

- The first, class-3, will match any packet to access-group 103.
- The second, class-2, matches IP packets.
- The third matches any input from interface Ethernet 1/0.

Policy Maps

This topic describes how to implement QoS policies using policy maps.

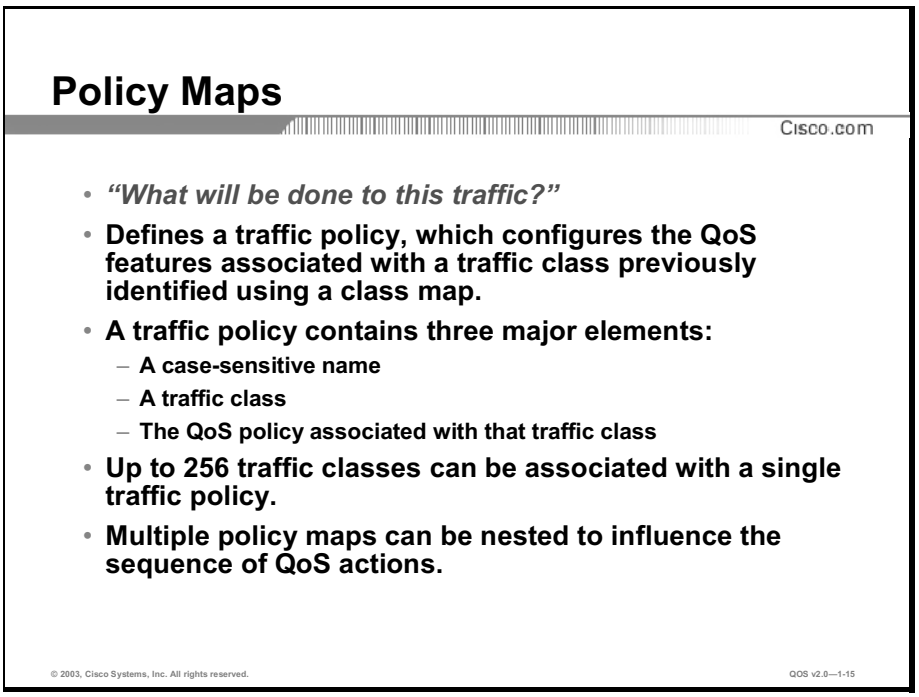

The **policy-map** command is used to create a traffic policy. The purpose of a traffic policy is to configure the QoS features that should be associated with the traffic that has been classified in a user-specified traffic class or classes. A traffic policy contains three elements: a case-sensitive 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). After the policy map command is issued, the user is placed into policy map configuration mode. The name of a traffic class can then be entered, and the user enters policy map class configuration mode. Here is where the user enters QoS features to apply to the traffic that matches this class.

The MQC does not necessarily require that users associate only one traffic class to a single traffic policy. When packets match to more than one match criterion, multiple traffic classes can be associated with a single traffic policy.

Note: A packet can match only one traffic class within a traffic policy. If a packet matches more than one traffic class in the traffic policy, the first traffic class defined in the policy will be used.

Configuring and Monitoring Policy Maps

This topic describes the commands that are necessary to configure and monitor policy maps.

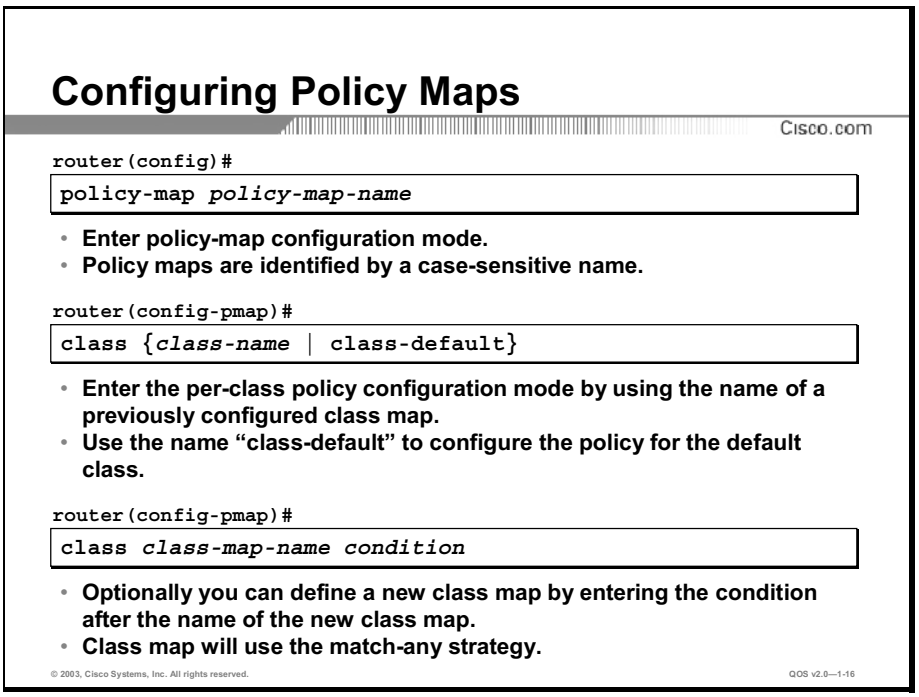

Service policies are configured using the **policy-map** command. Up to 256 classes can be used within one policy map using the **class** command with the name of a preconfigured class map.

A nonexistent class can also be used within the policy map configuration mode if the match condition is specified after the name of the class. The running configuration will reflect such a configuration by using the match-any strategy and inserting a full class map configuration.

The following table shows starting and resulting configuration modes for the **class-map**, **policy-map**, and **class** commands:

Configuration Modes

Starting configuration mode	Command	Configuration mode
Router(config)#	class-map	Router(config-cmap)#
Router(config)#	policy-map	Router(config-pmap)#
Router(config-pmap)#	class	Router(config-pmap-c)#

All traffic that is not classified by any of the class maps that are used within the policy map is part of the default class **class-default**. This class has no QoS guarantees by default. The default class, when used on output, can use one FIFO queue or flow-based weighted fair queuing (WFQ). The default class is part of every policy map even if not configured.

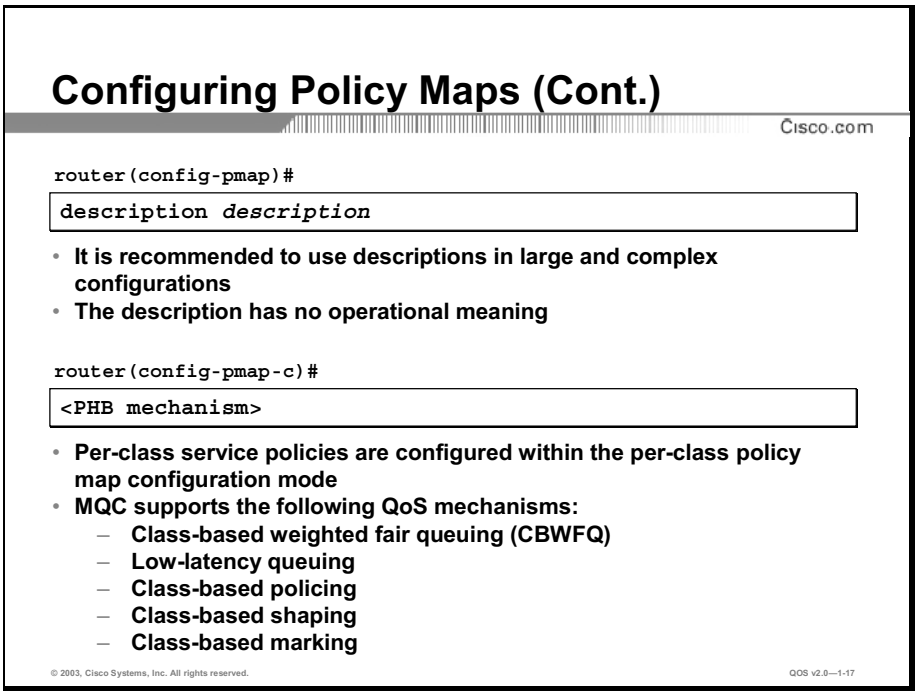

Policy maps, like class maps, should use descriptions in large QoS implementations where a large number of different policy maps are used.

Renaming a policy map would normally require the renaming of all the references to the policy map. Using the **rename** command simplifies the renaming process by automatically renaming all references.

Example: Policy Map Example

The example shows the configuration of a policy map using three classes. The first two classes were separately configured using the **class-map** command. The third class was configured by specifying the match condition after the name of the class:

```
class-map match-all Test1
 match protocol http
 match access-group 100
class-map match-any Test2
 match protocol http
 match access-group 101
!
policy-map Test
 class Test1
bandwidth 100
 class Test2
bandwidth 200
 class Test3 access-group 100
 bandwidth 300
!
access-list 100 permit tcp any host 10.1.1.1
access-list 101 permit tcp any host 10.1.1.2
```

Class Test1 has two match conditions evaluated in the match-all strategy. Classes Test2 and Test3 use the match-any strategy.

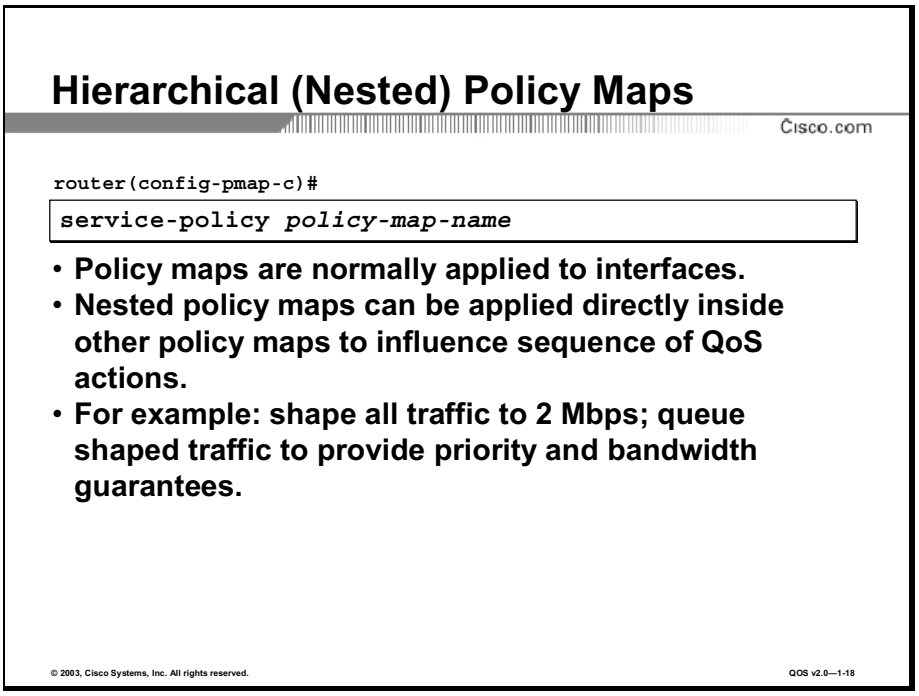

The **service-policy** *policy-map-name* command is used to create hierarchical service policies in policy map class configuration mode.

The **service-policy** [**input** | **output**] *policy-map-name* command is a different command that is used in interface configuration mode. The purpose of the **service-policy** [**input** | **output**] *policy-map-name* is to attach service policies to interfaces.

The child policy is the previously defined service policy that is now associated with the new service policy through the use of the **service-policy** command. The new service policy that uses the preexisting service policy is called the parent policy. In the hierarchical policy maps example below, the service policy named, "*child*" is the child policy and service policy named, "*parent*" is the parent policy.

The **service-policy** *policy-map-name* command has the following restrictions:

- The **set** command is not supported on the child policy.
- The **priority** command can be used in either the parent or the child policy, but not both policies simultaneously.
- The **fair-queue** command cannot be defined in the parent policy.

Example: Hierarchical Policy Maps

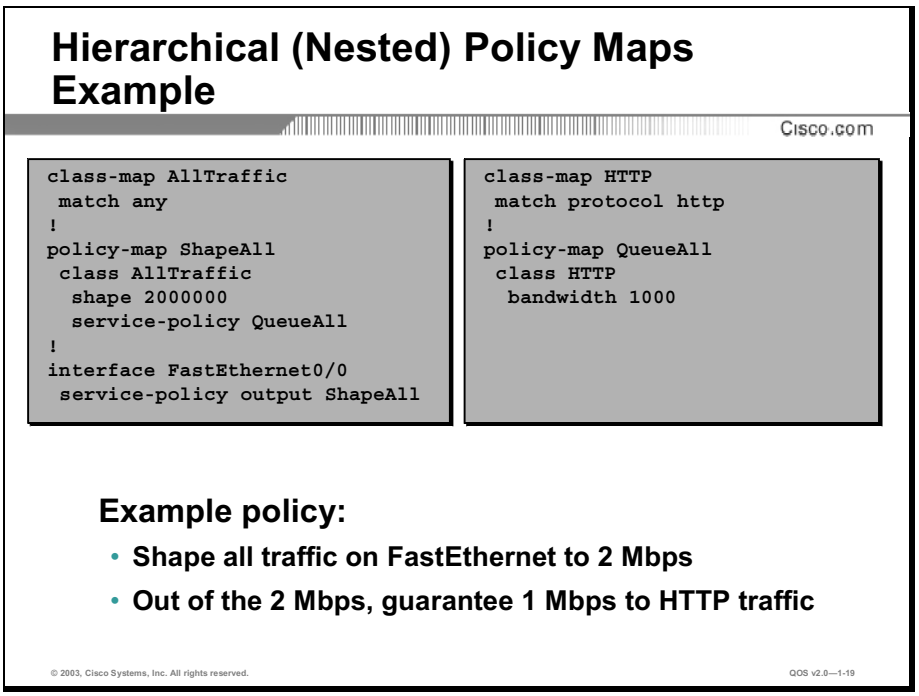

In the example diagram, a child policy-map QueueAll is created, which guarantees bandwidth of 1 Mbps to HTTP traffic.

The QueueAll policy map is then nested within a parent policy map named ShapeAll.

Finally, the parent policy map ShapeAll is applied to the FastEthernet interface.

Traffic out of the FastEthernet interface will first be shaped to 2 Mbps and then HTTP traffic will be guaranteed 1 Mbps of the 2 Mbps of shaped traffic.

Note: Additional information on traffic shaping is covered in the "Traffic Policing and Shaping" module in this course.

Example: Hierarchical Policy Map Configuration

Follow these steps to apply a hierarchical policy:

Step 1 Create a child or lower-level policy that configures a queuing mechanism. In the example below, LLQ is configured using the priority command.

```
policy-map child
  class voice
           priority 512
```

Step 2 Create a parent or top-level policy that applies class-based shaping. Apply the child policy as a command under the parent policy because the admission control for the child class is based on the shaping rate for the parent class.

```
policy-map parent
  class class-default
  shape average 2000000
        service-policy child
```

Step 3 Apply the parent policy to the subinterface.

```
  interface ethernet0/0.1
        service-policy output parent
```


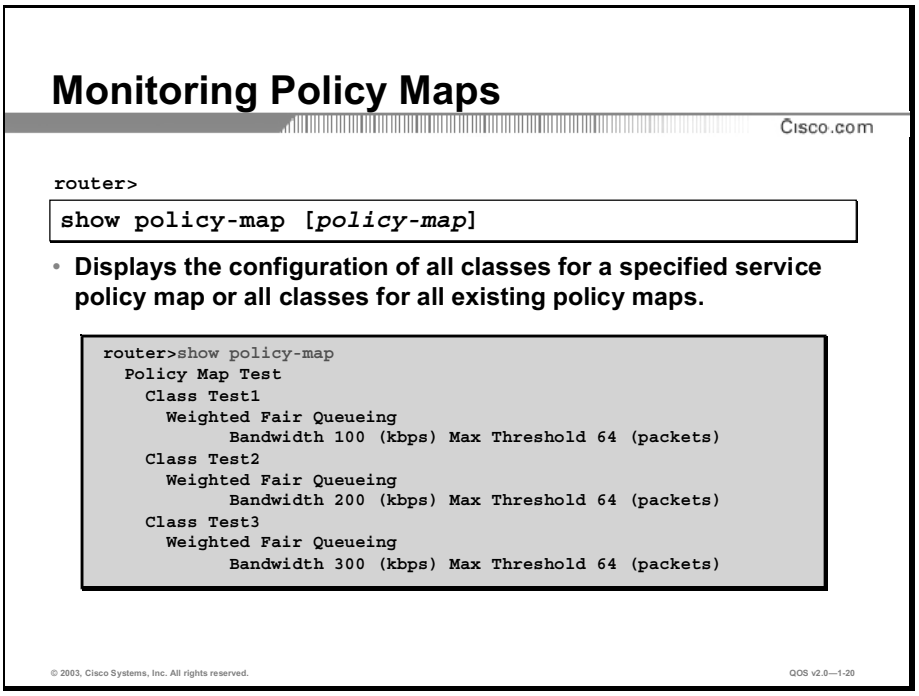

The **show policy-map** command can be used to verify the configuration of a policy map.

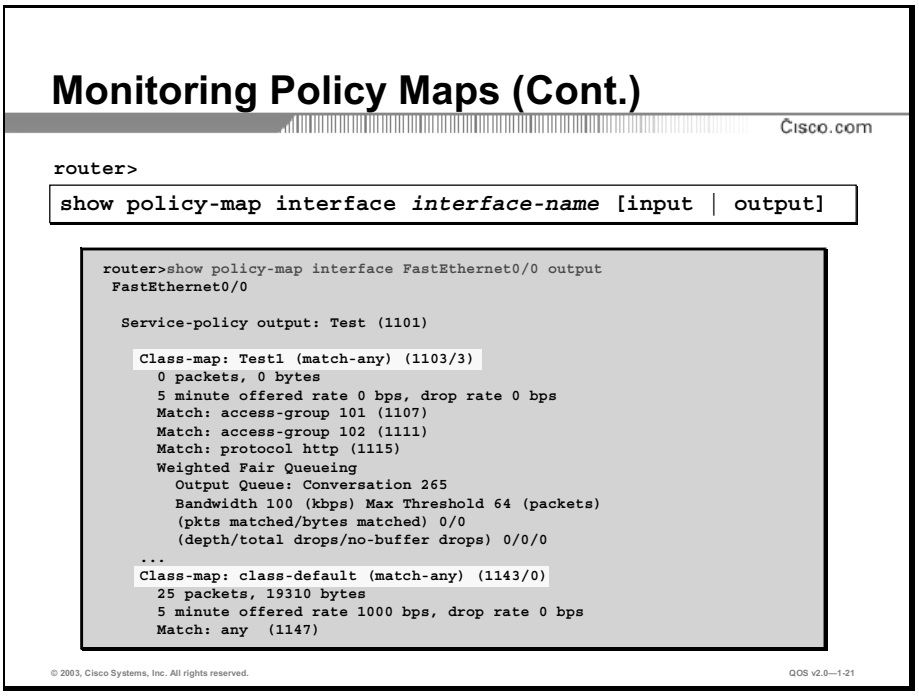

The **show policy-map** command also displays live information if the **interface** keyword is used. The sample output shows the parameters and statistics of the policy map that is attached to outbound traffic on interface FastEthernet0/0.

Service Policy

This topic describes how to attach a QoS policy to an interface using service policies.

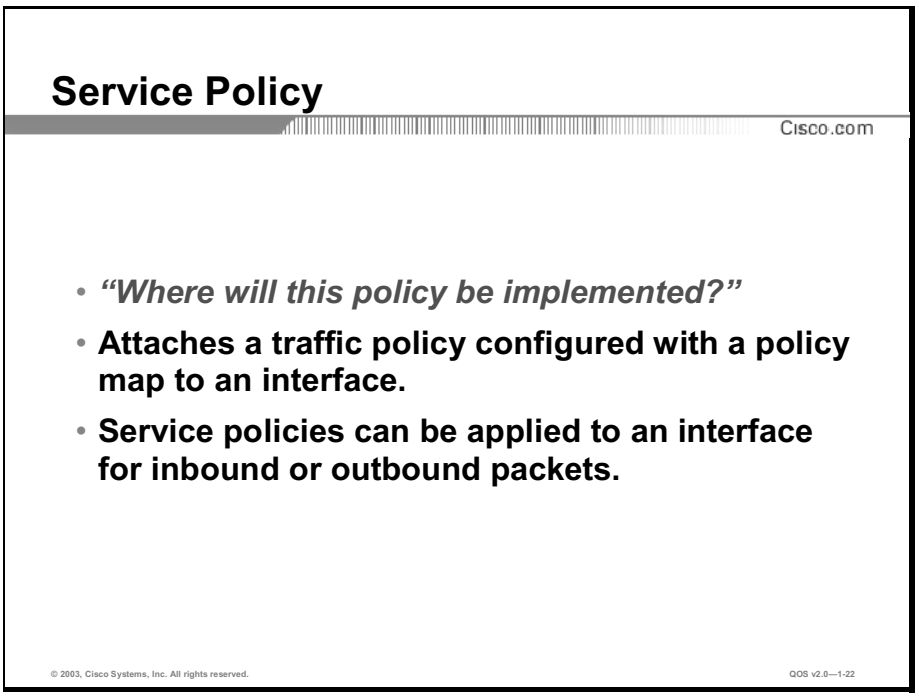

The last configuration step when configuring QoS mechanisms using the MQC is to attach a policy map to the inbound or outbound packets using the **service-policy** command.

Using the **service-policy** command it is possible to assign a single policy map to multiple interfaces or to assign multiple policy maps to a single interface (a maximum of one in each direction, inbound and outbound).

A service policy can be applied for inbound or outbound packets.

Attaching Service Policies to Interfaces

This topic explains how to attach service policies to interfaces.

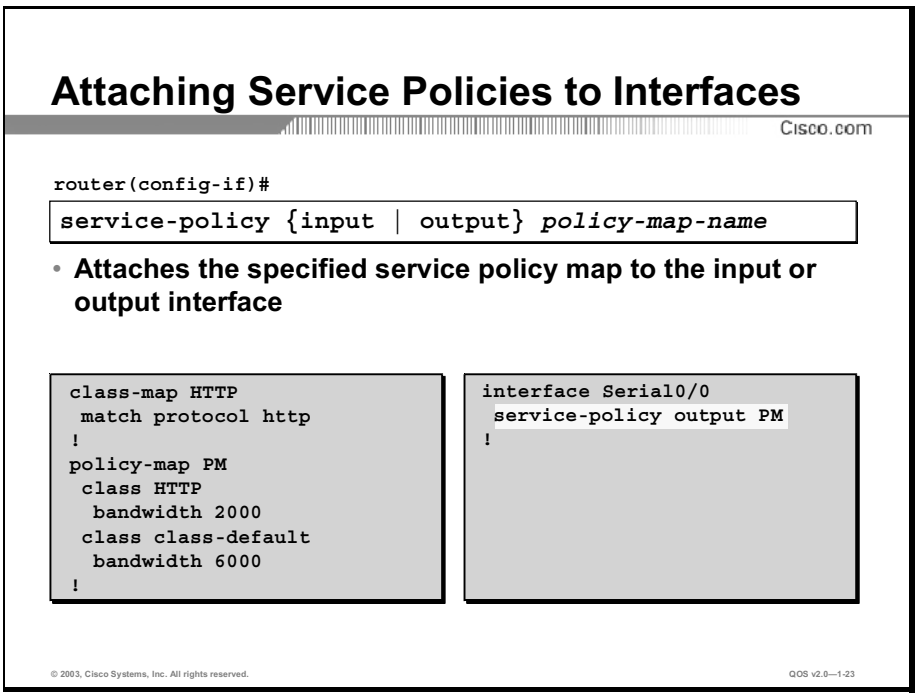

Use the **service-policy** interface configuration command to attach a traffic policy to an interface and to specify the direction in which the policy should be applied (either on packets coming into the interface or packets leaving the interface).

The router immediately verifies the correctness of parameters that are used in the policy map. If there is a mistake in the policy map configuration, the router will display a message explaining what is wrong with the policy map.

The sample configuration shows how a policy map is used to separate HTTP from other traffic. HTTP is guaranteed 2 Mbps. All other traffic belongs to the default class and is guaranteed 6 Mbps.

Example: Complete MQC Configuration

Traffic Classes Defined

In the following example, two traffic classes are created and their match criteria are defined. For the first traffic class called class1, access control list (ACL) 101 is used as the match criterion. For the second traffic class called class2, ACL 102 is used as the match criterion. Packets are checked against the contents of these ACLs to determine if they belong to the class:

```
Router(config)# class-map class1
Router(config-cmap)# match access-group 101
Router(config-cmap)# exit

Router(config)# class-map class2
Router(config-cmap)# match access-group 102
Router(config-cmap)# exit
```


Traffic Policy Created

In the following example, a traffic policy called policy1 is defined to contain policy specifications for the two classes—class1 and class2. The match criteria for these classes were defined in the traffic classes.

For class1, the policy includes a bandwidth allocation request and a maximum packet count limit for the queue reserved for the class. For class2, the policy specifies only a bandwidth allocation request:

```
Router(config)# policy-map policy1
Router(config-pmap)# class class1
Router(config-pmap-c)# bandwidth 3000
Router(config-pmap-c)# queue-limit 30
Router(config-pmap-c)# exit
Router(config-pmap)# class class2
Router(config-pmap-c)# bandwidth 2000
Router(config-pmap-c)# exit
```

Traffic Policy Attached to an Interface

The following example shows how to attach an existing traffic policy (which was created in the preceding section) to an interface. After you define a traffic policy with the **policy-map** command, you can attach it to one or more interfaces to specify the traffic policy for those interfaces by using the **service-policy** command in interface configuration mode. Although you can assign the same traffic policy to multiple interfaces, each interface can have only one traffic policy attached at the input and a single traffic policy attached at the output:

```
Router(config)# interface e1/1
Router(config-if)# service-policy output policy1
Router(config-if)# exit
Router(config)# interface fa1/0/0
Router(config-if)# service-policy output policy1
Router(config-if)# exit
```

Summary

This topic summarizes the key points discussed in this lesson.

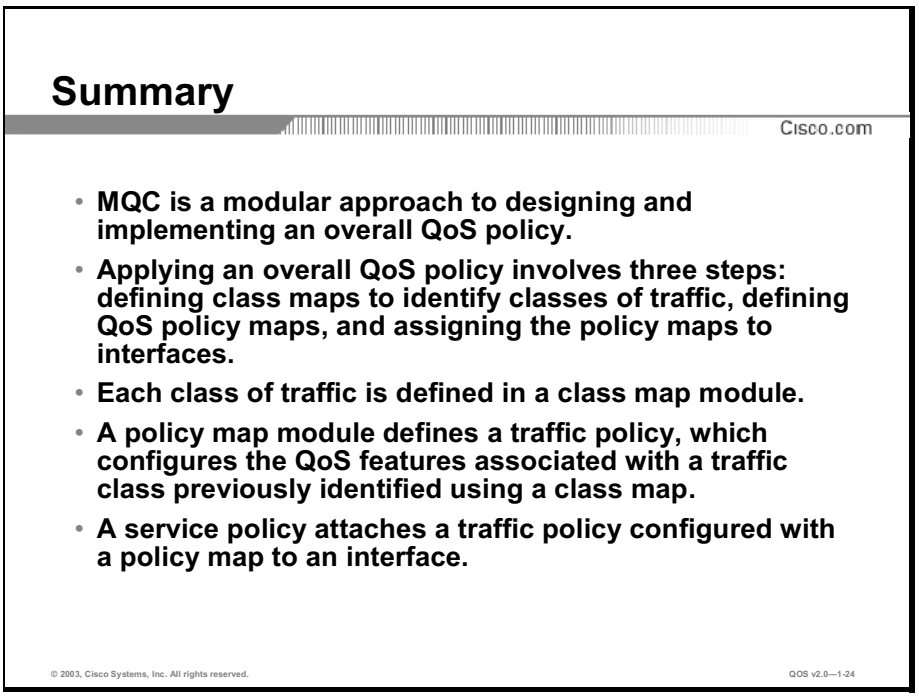

References

For additional information, refer to these resources:

- For more information on the Modular Quality of Service Command-Line Interface, refer to "Modular Quality of Service Command-Line Interface Overview " at the following URL: http://www.cisco.com/en/US/products/sw/iosswrel/ps1835/products_configuration_guide_chapter09186a00800bd908.html
- For more information on the Modular Quality of Service Command-Line Interface, refer to "QC: Part 8: Modular Quality of Service Command-Line Interface" at the following URL: http://www.cisco.com/en/US/products/sw/iosswrel/ps1835/products_configuration_guide_book09186a00800b75e4.html
- For more information on the Modular Quality of Service Command-Line Interface, refer to "Configuring the Modular Quality of Service Command-Line Interface " at the following URL: http://www.cisco.com/en/US/products/sw/iosswrel/ps1835/products_configuration_guide_chapter09186a00800bd909.html

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) Which three of the following are steps in implementing a QoS policy using MQC? (Choose three.)

A) classify traffic

B) configure queuing mechanisms

C) define QoS policies and assign traffic

D) attach policies to interfaces

Q2) Match the MQC implementation step with its associated Cisco IOS command:

A) define classes of traffic

B) define QoS policies for classes

C) apply policy map to interface

_____ 1. policy-map

_____ 2. service-policy

_____ 3. class-map

Q3) Which match type is the default mode for class maps?

A) match none

B) match all

C) match any

D) match first

Q4) When using the match-any conditional when defining a class map, what happens if a packet matches more than one condition?

A) The packet would not be considered a member of the class.

B) The packet would be considered a member of the class.

C) The packet would be moved to the specified alternate match class.

D) The packet would be dropped.

Q5) How many traffic classes can be assigned to a single policy map?

A) 32

B) 64

C) 256

D) 1044

Q6) How do you configure the class default?

A) Using the **class-map default** command.

B) No need to—it is automatically configured.

C) As the last class defined within a class map.

D) As the last class specified in a policy map.

Quiz Answer Key

Q1) A, C, D

Relates to: Modular QoS CLI

Q2) A – 3, B – 1, C - 2

Relates to: Modular QoS CLI Components

Q3) B

Relates to: Class Maps

Q4) B

Relates to: Configuring and Monitoring Class Maps

Q5) C

Relates to: Policy Maps

Q6) B

Relates to: Configuring and Monitoring Policy Maps

Introducing AutoQoS

Overview

Cisco AutoQoS represents innovative technology that simplifies network administration challenges, reducing QoS complexity, deployment time, and cost in enterprise networks. Cisco AutoQoS incorporates value-added intelligence in Cisco IOS software and Cisco Catalyst software to provision and manage large-scale QoS deployments. Cisco AutoQoS provides QoS provisioning for individual routers and switches, simplifying deployment and reducing human error.

The first phase of Cisco AutoQoS offers straightforward capabilities to automate Voice over IP (VoIP) deployments for customers who want to deploy IP telephony, but who lack the expertise and staffing to plan and deploy IP QoS and IP services.

Relevance

AutoQoS is one of the two key methods recommended for implementing QoS on a network. AutoQoS may be the best choice for quickly and easily implementing networks incorporating voice applications.

Objectives

Upon completing this lesson, you will be able to correctly identify capabilities provided by AutoQoS and successfully configure QoS on a network using AutoQoS. This includes being able to meet these objectives:

- Explain how AutoQoS is used to implement QoS policy
- Describe the router environments in which AutoQoS can be used
- Describe the switch environments in which AutoQoS can be used
- Configure AutoQoS on a network using CLI
- Use Cisco IOS commands to examine and monitor a network configuration after AutoQoS has been enabled
- Identify several of the QoS technologies that were automatically implemented on the network using AutoQoS

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Basic understanding of the Cisco IOS command-line interface

Outline

The outline lists the topics included in this lesson.

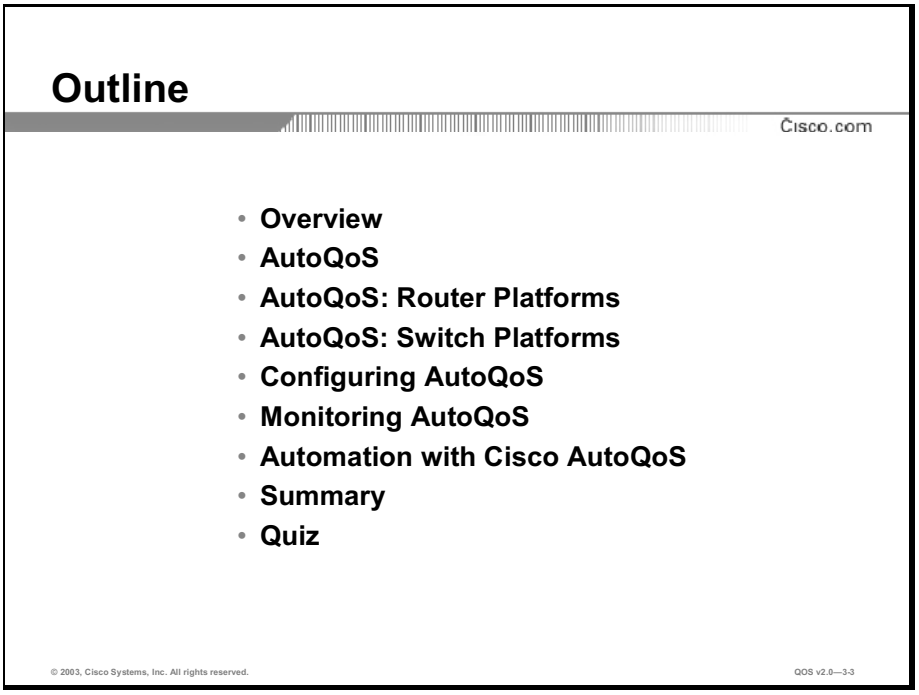

AutoQoS

This topic describes the basic purpose and function of AutoQoS.

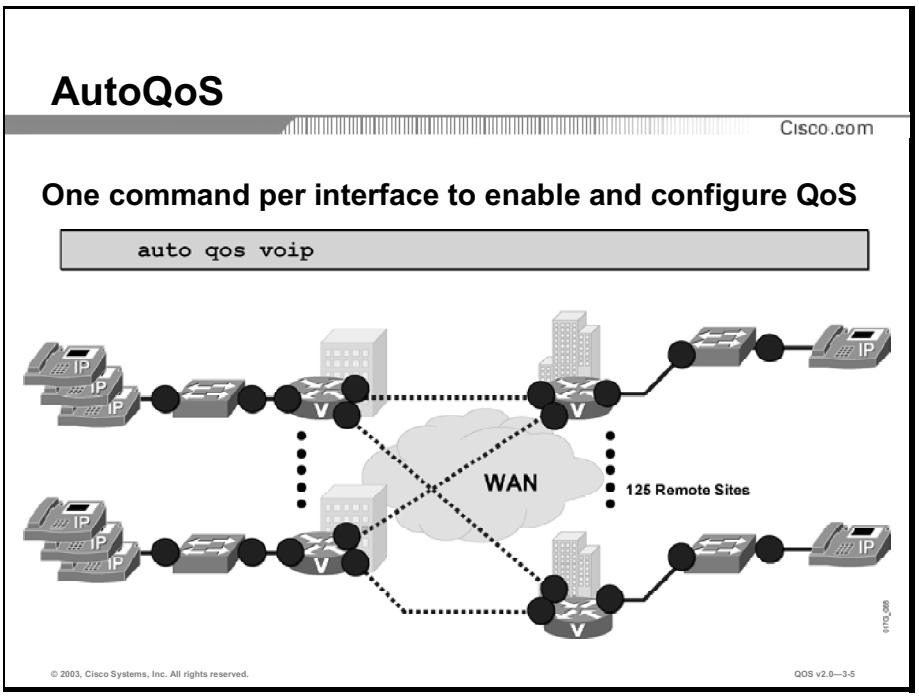

AutoQoS gives customers the ability to deploy QoS features for converged IP telephony and data networks much faster and more efficiently. It simplifies and automates the MQC definition of traffic classes and the creation and configuration of traffic policies. (Cisco AutoQoS generates traffic classes and policy map CLI templates.) Therefore, when AutoQoS is configured at the interface or permanent virtual circuit (PVC), the traffic receives the required QoS treatment automatically. In-depth knowledge of the underlying technologies, service policies, link efficiency mechanisms, and Cisco QoS best practice recommendations for voice requirements is not required to configure AutoQoS.

Cisco AutoQoS can be extremely beneficial for the following scenarios:

- Small- to medium-sized businesses that must deploy IP telephony quickly, but lack the experience and staffing to plan and deploy IP QoS services.
- Large customer enterprises that need to deploy Cisco telephony solutions on a large scale, while reducing the costs, complexity, and timeframe for deployment and ensuring that the appropriate QoS for voice applications is being set in a consistent fashion.
- International enterprises or service providers requiring QoS for VoIP where little expertise exists in different regions of the world and where provisioning QoS remotely and across different time zones is difficult.
- Service providers requiring a template-driven approach to delivering managed services and QoS for voice traffic to large numbers of customer premise devices.

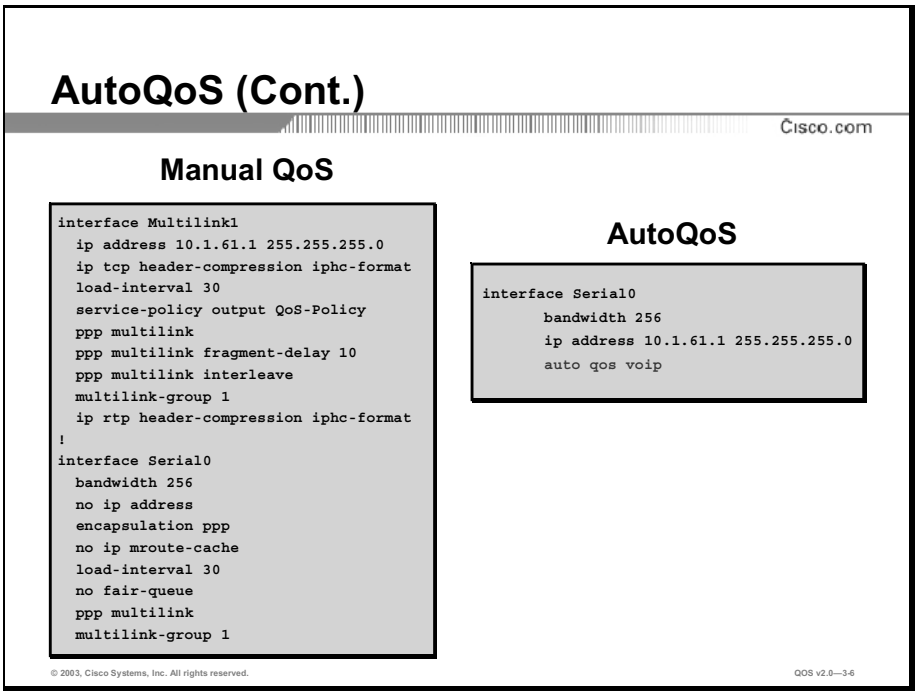

Cisco AutoQoS automatically creates the QoS-specific features required for supporting the underlying transport mechanism and link speed of an interface or PVC type. For example, Frame Relay traffic shaping (FRTS) would be automatically configured and enabled by Cisco AutoQoS for Frame Relay links. Link fragmentation and interleaving (LFI) and compressed Real-Time Transport Protocol (cRTP) would be automatically configured via the Cisco AutoQoS template for slow link speeds (less than 768 kbps). Therefore, it is very important that the bandwidth statement be properly set on the interface *prior* to configuring AutoQoS because the resulting configuration will vary, based on this configurable parameter.

Using Cisco AutoQoS, VoIP traffic is automatically provided with the required QoS template for voice traffic by configuring **auto qos voip** on an interface or PVC. Cisco AutoQoS enables the required QoS based on Cisco best-practices methodologies. (The configuration generated by Cisco AutoQoS can be modified if necessary.)

AutoQoS (Cont.)

Cisco.com

- **Application Classification**
  - Automatically discovers applications and provides appropriate QoS treatment
- **Policy Generation**
  - Automatically generates initial and ongoing QoS policies
- **Configuration**
  - Provides high-level business knobs, and multi-device/domain automation for QoS
- **Monitoring & Reporting**
  - Generates intelligent, automatic alerts and summary reports
- **Consistency**
  - Enables automatic, seamless interoperability among all QoS features and parameters across a network topology – LAN, MAN, and WAN

Cisco AutoQoS simplifies and shortens the QoS deployment cycle. Cisco AutoQoS helps in all five major aspects of successful QoS deployments:

- **Application Classification:** Cisco AutoQoS leverages intelligent classification on routers utilizing Cisco network-based application recognition (NBAR) to provide deep and stateful packet inspection. Cisco AutoQoS uses Cisco Discovery Protocol (CDP) for voice packets, ensuring that the device attached to the LAN is really an IP Phone.
- **Policy Generation:** Cisco AutoQoS evaluates the network environment and generates an initial policy. It automatically determines WAN settings for fragmentation, compression, encapsulation, and Frame Relay-to-ATM Service Interworking (FRF.8), eliminating the need to understand QoS theory and design practices in various scenarios. Customers can meet additional or special requirements by modifying the initial policy as they normally would.

  The first release of Cisco AutoQoS provides the necessary AutoQoS VoIP feature to automate QoS settings for VoIP deployments. This feature automatically generates interface configurations, policy maps, class maps, and ACLs. AutoQoS VoIP will automatically employ Cisco NBAR to classify voice traffic, and mark it with the appropriate differentiated services code point (DSCP) value. AutoQoS VoIP can be instructed to rely on, or trust, the DSCP markings previously applied to the packets.
- **Configuration:** With one command, Cisco AutoQoS configures the port to prioritize voice traffic without affecting other network traffic while still offering the flexibility to adjust QoS settings for unique network requirements.

  Not only will Cisco AutoQoS automatically detect Cisco IP Phones and enable QoS settings, it will disable the QoS settings when a Cisco IP Phone is relocated or moved to prevent malicious activity.

  AutoQoS generated router and switch configurations are customizable using the standard Cisco IOS CLI.

- **Monitoring & Reporting:** Cisco AutoQoS provides visibility into the classes of service deployed via system logging and Simple Network Management Protocol (SNMP) traps, with notification of abnormal events (that is, VoIP packet drops).
- **Consistency:** When deploying QoS configurations using AutoQoS, configurations generated are consistent among router and switch platforms. This level of consistency ensures seamless QoS operation and interoperability within the network.

AutoQoS: Router Platforms

This topic identifies the router platforms on which AutoQoS will operate.

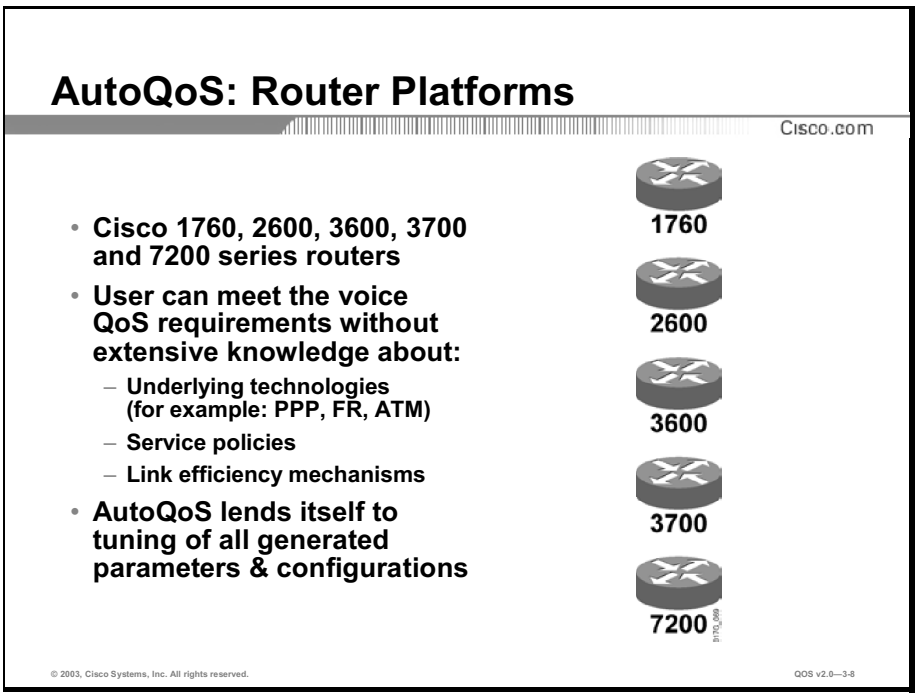

Initial support for AutoQoS includes the Cisco 2600 (including XM models), 3600, 3700, and 7200 series routers. Support for additional platforms will become available.

The Cisco AutoQoS VoIP feature is supported only on the following interfaces and PVCs:

- Serial interfaces with PPP or high-level data link control (HDLC)
- Frame Relay data-link connection identifiers (DLCIs)—PPP subinterfaces only
  - Cisco AutoQoS does not support Frame Relay multipoint interfaces
- ATM PVCs
  - Cisco AutoQoS VoIP is supported on low-speed ATM PVCs on PPP subinterfaces only (link bandwidth less than 768 kbps)
  - Cisco AutoQoS VoIP is fully supported on high-speed ATM PVCs (link bandwidth greater than 768 kbps)

AutoQoS: Switch Platforms

This topic identifies the switch platforms on which AutoQoS will operate.

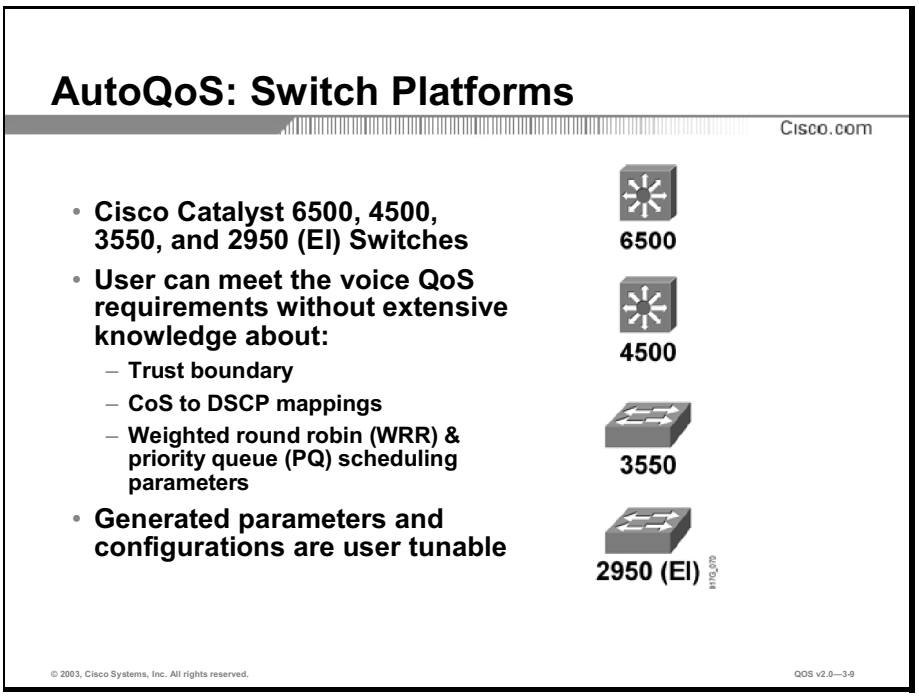

Initial support for AutoQoS includes the Cisco Catalyst 6500, 4500, 3550, and 2950 (EI) series switches. Support for additional platforms including the Cisco Catalyst 4000 will become available.

The Enhanced Image (EI) is required on the Cisco Catalyst 2950 series switches.

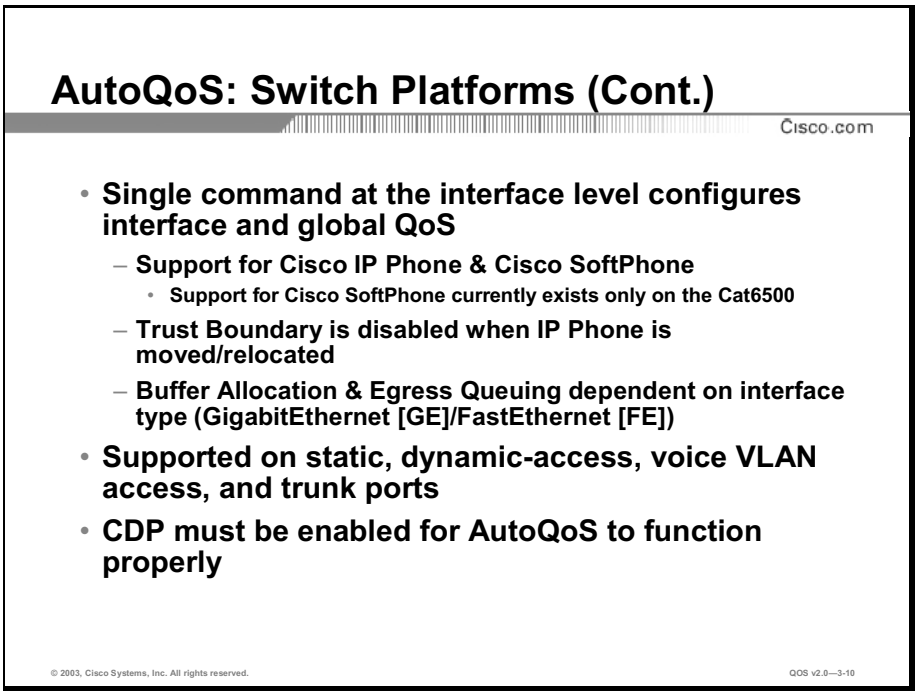

To configure the QoS settings and the trusted boundary feature on the Cisco IP Phone, you must enable CDP version 2 or later on the port. If you enable the trusted boundary feature, a syslog warning message displays if CDP is not enabled or if CDP is running version 1.

You need to enable CDP only for the **ciscoipphone** QoS configuration; CDP does not affect the other components of the automatic QoS features. When you use the **ciscoipphone** keyword with the port-specific automatic QoS feature, a warning displays if the port does not have CDP enabled.

When executing the port-specific automatic QoS command with the **ciscoipphone** keyword without the trust option, the trust-device feature is enabled. The trust-device feature is dependent on CDP. If CDP is not enabled or not running version 2, a warning message displays as follows:

```
Console> (enable) set port qos 4/1 autoqos voip ciscoipphone
Warning: CDP is disabled or CDP version 1 is in use. Ensure
that CDP version 2 is enabled globally, and also ensure that
CDP is enabled on the port(s) you wish to configure autoqos
on.
Port 4/1 ingress QoS configured for ciscoipphone.
It is recommended to execute the "set qos autoqos" global
command if not executed previously.
Console> (enable)
```

Configuring AutoQoS

This topic describes one of the key prerequisites for using AutoQoS.

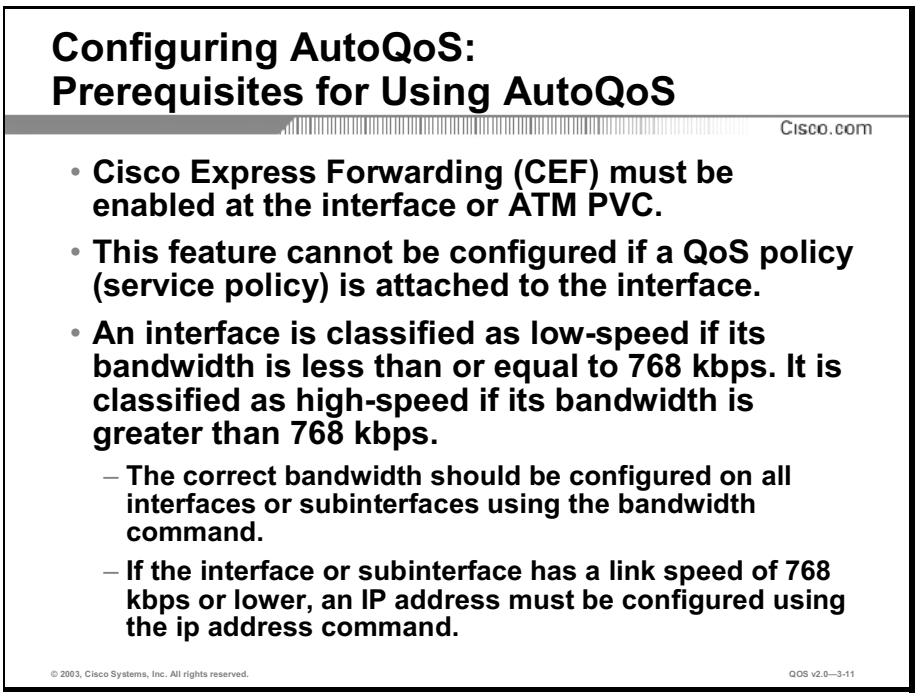

Before configuring AutoQoS, the following prerequisites must be met:

- Cisco Express Forwarding (CEF) must be enabled at the interface or ATM PVC. Cisco AutoQoS uses NBAR to identify various applications and traffic types and CEF is a prerequisite for NBAR.
- Ensure that no QoS policies (service policies) are attached to the interface. This feature cannot be configured if a QoS policy (service policy) is attached to the interface.
- AutoQoS classifies links as either low-speed or high-speed depending upon the link bandwidth. Remember that on a serial interface, if the default bandwidth is not specified it is 1.544 Mbps. Therefore, it is important that the correct bandwidth be specified on the interface or subinterface where AutoQoS is to be enabled.
  - For all interfaces or subinterfaces, be sure to properly configure the bandwidth by using the bandwidth command. The amount of bandwidth that is allocated should be based on the link speed of the interface.
  - If the interface or subinterface has a link speed of 768 kbps or lower, an IP address must be configured on the interface or subinterface using the **ip address** command. By default, AutoQoS will enable Multilink PPP (MLP) and copy the configured IP address to the multilink bundle interface.

In addition to the AutoQoS prerequisites, the following are recommendations and requirements when configuring AutoQoS. Be aware that these may change with Cisco IOS releases and should be verified before implementing AutoQoS in your environment.

- The Cisco AutoQoS VoIP feature is supported only on the following interfaces and PVCs:
  - Serial interfaces with PPP or HDLC
  - Frame Relay DLCIs (PPP subinterfaces only)
    - Cisco AutoQoS does not support Frame Relay multipoint interfaces
  - ATM PVCs
- Configuration template (CLI) generated by configuring Cisco AutoQoS on an interface or PVC can be tuned manually (via CLI configuration) if desired.
- Cisco AutoQoS cannot be configured if a QoS service-policy is already configured and attached to the interface or PVC.
- MLP is configured automatically for a serial interface with low-speed link. The serial interface must have an IP address and this IP address is removed and put on the MLP bundle. Cisco AutoQoS VoIP must also be configured on the other side of the link
- The **no auto qos voip** command removes Cisco AutoQoS. However, if the interface or PVC Cisco AutoQoS generated QoS configuration is deleted without configuring the **no auto qos voip** command, Cisco AutoQoS VoIP will not be completely removed from the configuration properly.
- Cisco AutoQoS SNMP traps are only delivered when an SNMP server is used in conjunction with Cisco AutoQoS.
- The SNMP community string “AutoQoS” should have “write” permissions.
- If the device is reloaded with the saved configuration after configuring Cisco AutoQoS and saving the configuration to NVRAM, some warning messages may be generated by Remote Monitoring (RMON) threshold commands. These warning messages can be ignored. (To avoid further warning messages, save the configuration to NVRAM again without making any changes to the QoS configuration.)
- By default, Cisco 7200 series routers and below that support MQC QoS, reserve up to 75 percent of the interface bandwidth for user-defined classes. The remaining bandwidth is used for the default class. However, the entire remaining bandwidth is *not* guaranteed to the default class. This bandwidth is shared proportionately between the different flows in the default class and excess traffic from other bandwidth classes. At least one percent of the available bandwidth is reserved and guaranteed for class default traffic by default on Cisco 7500 series routers. (Up to 99 percent can be allocated to the other classes.)

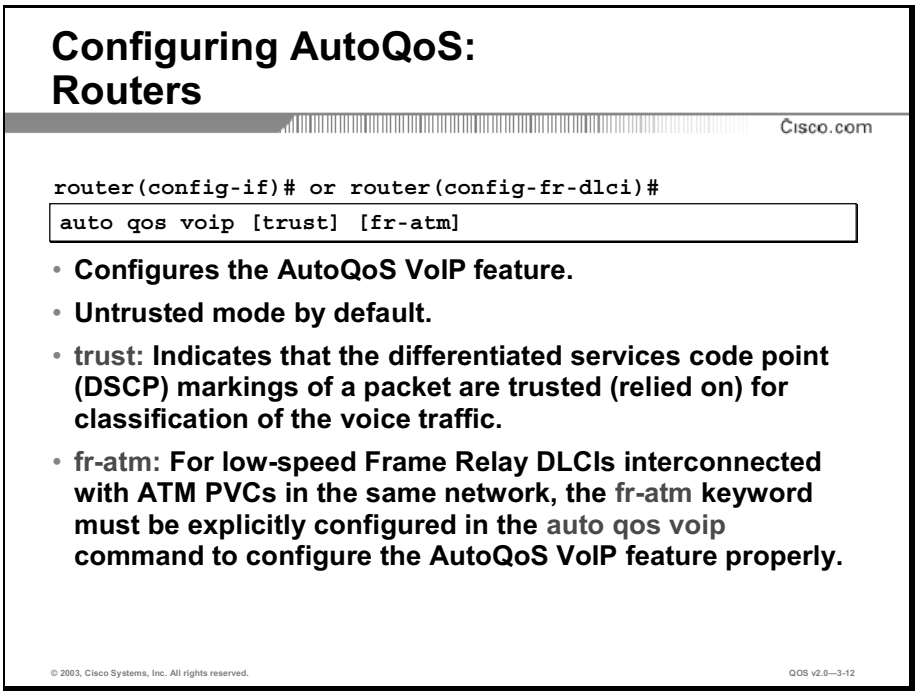

To configure the AutoQoS VoIP feature on an interface, use the **auto qos voip** command in interface configuration mode or Frame Relay DLCI configuration mode. To remove the AutoQoS VoIP feature from an interface, use the **no** form of the **auto qos voip** command.**auto qos voip [trust] [fr-atm]**

no auto qos voip [trust] [fr-atm]

Syntax Description

Parameter	Description
trust	(Optional) Indicates that the DSCP markings of a packet are trusted (relied on) for classification of the voice traffic. If the optional trust keyword is not specified, the voice traffic is classified using NBAR, and the packets are marked with the appropriate DSCP value.
fr-atm	(Optional) Enables the AutoQoS VoIP feature for the Frame Relay-to-ATM links. This option is available on the Frame Relay DLCIs for Frame Relay-to-ATM interworking only.

The bandwidth of the serial interface is used to determine the link speed. The link speed is one element that is used to determine the configuration generated by the AutoQoS VoIP feature. The AutoQoS VoIP feature uses the bandwidth at the time the feature is configured and does not respond to changes made to bandwidth after the feature is configured.

For example, if the **auto qos voip** command is used to configure the AutoQoS VoIP feature on an interface with 1000 kbps, the AutoQoS VoIP feature generates configurations for high-speed interfaces. However, if the bandwidth is later changed to 500 kbps, the AutoQoS VoIP feature will not use the lower bandwidth. The AutoQoS VoIP feature retains the higher bandwidth and continues to use the generated configurations for high-speed interfaces.

To force the AutoQoS VoIP feature to use the lower bandwidth (and thus generate configurations for the low-speed interfaces), use the **no auto qos voip** command to remove the AutoQoS VoIP feature and then reconfigure the feature.

Example: Configuring the AutoQoS VoIP Feature on a High-Speed Serial Interface

In this example, the AutoQoS VoIP feature is configured on the high-speed serial interface s1/2:

```
Router> enable
Router# configure terminal
Router(config)# interface s1/2
Router(config-if)# bandwidth 1540
Router(config-if)# auto qos voip
Router(config-if)# exit
```

Example: Configuring the AutoQoS VoIP Feature on a Low-Speed Serial Interface Example

In this example, the AutoQoS VoIP feature is configured on the low-speed serial interface s1/3:

```
Router# configure terminal
Router(config)# interface s1/3
Router(config-if)# bandwidth 512
Router(config-if)# ip address 10.10.100.1 255.255.255.0
Router(config-if# auto qos voip
Router(config-if)# exit
```

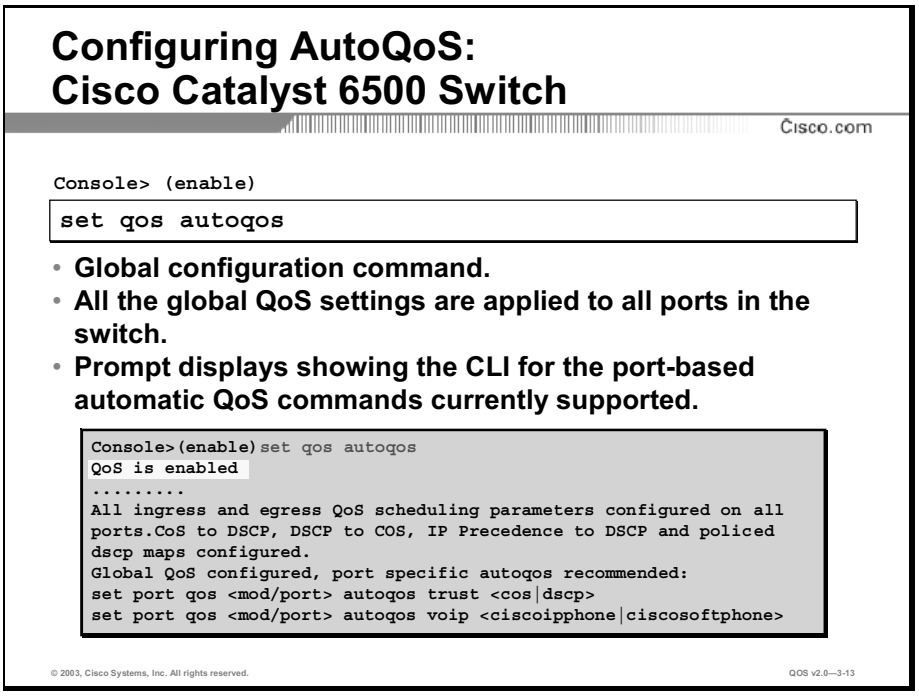

When you execute the global automatic QoS macro, all the global QoS settings are applied to all ports in the switch. After completion, a prompt will display showing the CLI for the port-based automatic QoS commands that are currently supported.

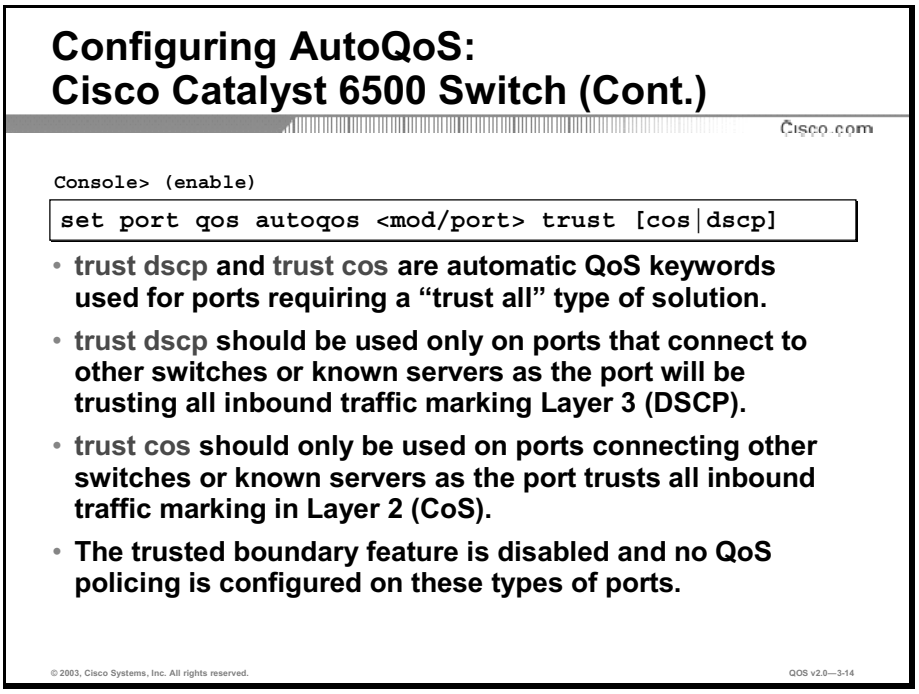

The port-specific automatic QoS macro handles all inbound QoS configuration that is specific to a particular port.

The QoS ingress port specific settings include port trust, default class of service (CoS), classification, and policing, but does not include scheduling. Input scheduling is programmed through the global automatic QoS macro. Together with the global automatic QoS macro command, all QoS settings are configured properly for a specific QoS traffic type.

Any existing QoS ACLs that are already associated with a port are removed when AutoQoS modifies ACL mappings on that port. The ACL names and instances are not changed.

Configuring AutoQoS: Cisco Catalyst 6500 Switch (Cont.)

Cisco.com

```
Console> (enable)
set port qos autoqos <mod/port> voip [ciscosoftphone
| ciscoipphone]
```

ciscosoftphone

- **The trusted boundary feature must be disabled for Cisco SoftPhone ports.**
- **QoS settings must be configured to trust the Layer 3 markings of the traffic that enters the port.**
- **Only available on Catalyst 6500.**

ciscoipphone

- **The port is set up to trust-cos as well as to enable the trusted boundary feature.**
- **Combined with the global automatic QoS command, all settings are configured on the switch to properly handle the signaling and voice bearer and PC data entering and leaving the port.**
- **CDP must be enabled for the ciscoipphone QoS configuration.**

The port-specific automatic QoS macro accepts a *mod/port* combination and must include a Cisco Architecture for Voice, Video and Integrated Date (AVVID) type of keyword. The **ciscoipphone**, **ciscosoftphone**, and **trust** keywords are supported.

With the **ciscoipphone** keyword, the port is set up to trust-cos as well as to enable the trusted boundary feature. Combined with the global automatic QoS command, all settings are configured on the switch to properly handle the signaling and voice bearer and PC data entering and leaving the port.

In addition to the switch-side QoS settings that are covered by the global automatic QoS command, the IP Phone has a few QoS features that need to be configured for proper labeling to occur. QoS configuration information is sent to the IP Phone through CDP from the switch. The QoS values that need to be configured are the trust settings of the “PC port” on the IP Phone (trust or untrusted) and the CoS value that is used by the IP Phone to remark packets in case the port is untrusted (ext-cos).

Only the Catalyst 6500 supports AutoQoS for Cisco SoftPhone. On the ports that connect to a Cisco SoftPhone, QoS settings must be configured to trust the Layer 3 markings of the traffic that enters the port. Trusting all Layer 3 markings is a security risk because PC users could send non-priority traffic with DSCP 46 and gain unauthorized performance benefits. Although not configured by AutoQos, policing on all inbound traffic can be used to prevent malicious users from obtaining unauthorized bandwidth from the network. Policing is accomplished by rate limiting the DSCP 46 (EF) inbound traffic to the codec rate used by the Cisco SoftPhone application (worst case G.722). Any traffic that exceeds this rate is marked down to the default traffic rate (DSCP 0 - BE). Signaling traffic (DSCP 24) is also policed and marked down to zero if excess signaling traffic is detected. All other inbound traffic types are reclassified to default traffic (DSCP 0 - BE).

Note: You must disable the trusted boundary feature for Cisco SoftPhone ports.

Example: Using the Port-Specific AutoQoS Macro

This example shows how to use the **ciscoipphone** keyword:

```
Console> (enable) set port qos 3/1 autoqos help
Usage: set port qos <mod/port> autoqos trust <cos|dscp>
set port qos <mod/port> autoqos voip
<ciscoipphone|ciscosoftphone>
Console> (enable) set port qos 3/1 autoqos voip ciscoipphone
Port 3/1 ingress QoS configured for Cisco IP Phone.
It is recommended to execute the "set qos autoqos" global
command if not executed previously.
Console> (enable)
```

This example shows how to use the **ciscosoftphone** keyword:

```
Console> (enable) set port qos 3/1 autoqos voip ciscosoftphone
Port 3/1 ingress QoS configured for Cisco Softphone.
It is recommended to execute the "set qos autoqos" global
command if not executed previously.
Console> (enable)
```

This example shows how to use the **trust cos** keyword:

```
Console> (enable) set port qos 3/1 autoqos trust cos
Port 3/1 QoS configured to trust all incoming CoS marking.
It is recommended to execute the "set qos autoqos" global
command if not executed previously.
Console> (enable)
```

This example shows how to use the **trust dscp** keyword:

```
Console> (enable) set port qos 3/1 autoqos trust dscp
Port 3/1 QoS configured to trust all incoming DSCP marking.
It is recommended to execute the "set qos autoqos" global
command if not executed previously.
Console> (enable)
```

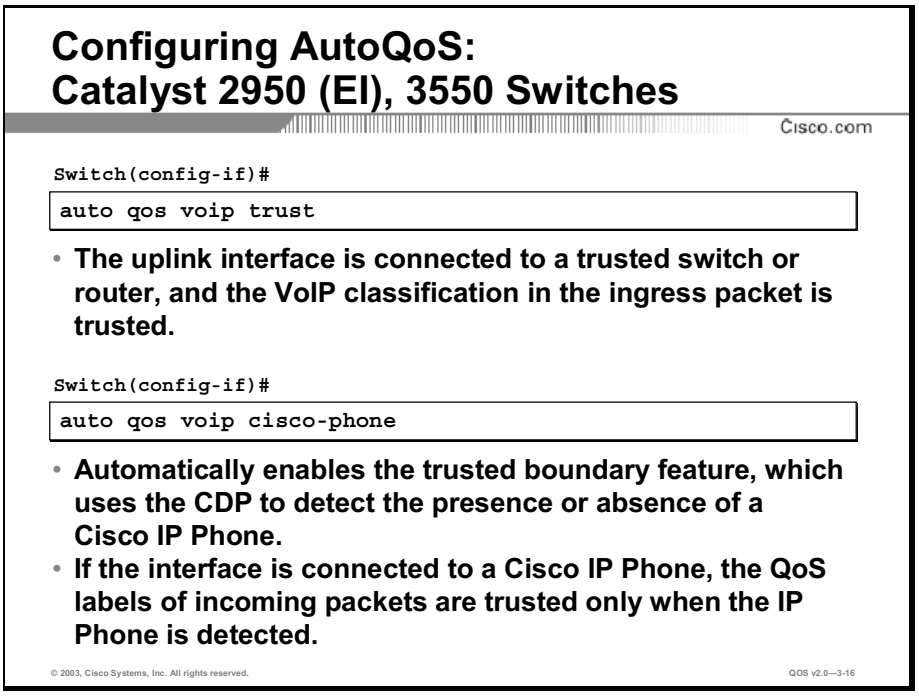

When you enable the AutoQoS feature on the first interface, QoS is globally enabled (**mls qos** global configuration command).

When you enter the **auto qos voip trust** interface configuration command, the ingress classification on the interface is set to trust the CoS QoS label received in the packet, and the egress queues on the interface are reconfigured. QoS labels in ingress packets are trusted.

When you enter the **auto qos voip cisco-phone** interface configuration command, the trusted boundary feature is enabled. It uses the CDP to detect the presence or absence of a Cisco IP Phone. When a Cisco IP Phone is detected, the ingress classification on the interface is set to trust the QoS label received in the packet. When a Cisco IP Phone is absent, the ingress classification is set to not trust the QoS label in the packet. The egress queues on the interface are also reconfigured. This command extends the trust boundary if IP Phone is detected.

Monitoring AutoQoS

This topic describes the commands that are used to monitor AutoQoS configurations.

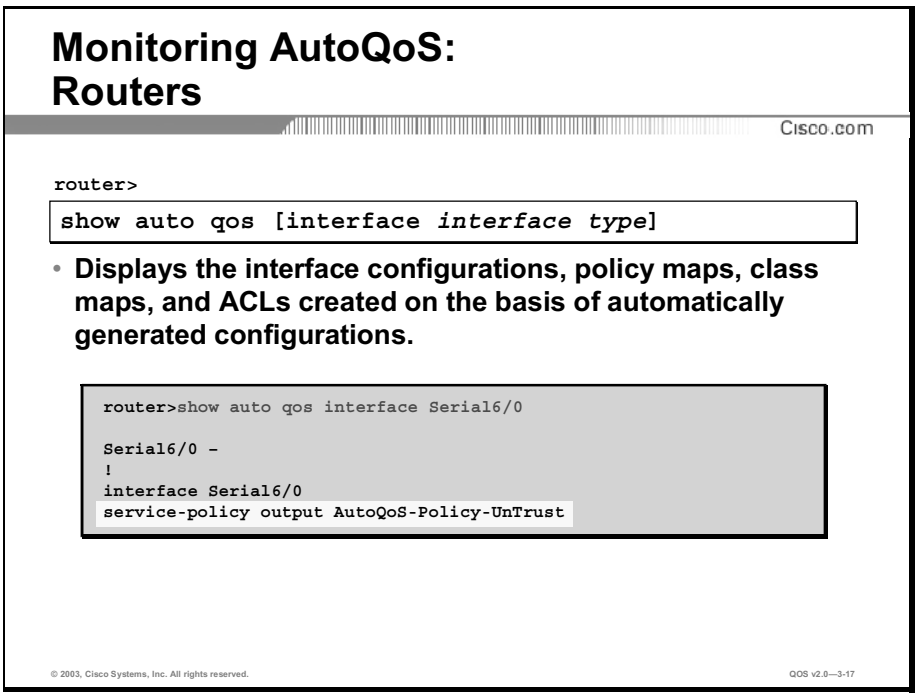

When the **auto qos voip** command is used to configure the AutoQoS VoIP feature, configurations are generated for each interface or PVC. These configurations are then used to create the interface configurations, policy maps, class maps, and ACLs. The **show auto qos** command can be used to verify the contents of the interface configurations, policy maps, class maps, and ACLs.

The **show auto qos interface** command can be used with Frame Relay DLCIs and ATM PVCs.

When the **interface** keyword is used along with the corresponding interface type argument, the **show auto qos interface** [*interface type*] command displays the configurations created by the AutoQoS VoIP feature on the specified interface.

When the **interface** keyword is used but an interface type is not specified, the **show auto qos interface** command displays the configurations created by the AutoQoS VoIP feature on all the interfaces or PVCs on which the AutoQoS VoIP feature is enabled.

Example: Show Auto QoS and Show Auto QoS Interface

The **show auto qos** command displays all of the configurations created by the AutoQoS VoIP feature:

```
Router# show auto qos
Serial6/1.1: DLCI 100 -
!
interface Serial6/1
frame-relay traffic-shaping
!
interface Serial6/1.1 point-to-point
frame-relay interface-dlci 100
class AutoQoS-VoIP-FR-Serial6/1-100
frame-relay ip rtp header-compression
!
map-class frame-relay AutoQoS-VoIP-FR-Serial6/1-100
frame-relay cir 512000
frame-relay bc 5120
frame-relay be 0
frame-relay mincir 512000
service-policy output AutoQoS-Policy-UnTrust
frame-relay fragment 640
```


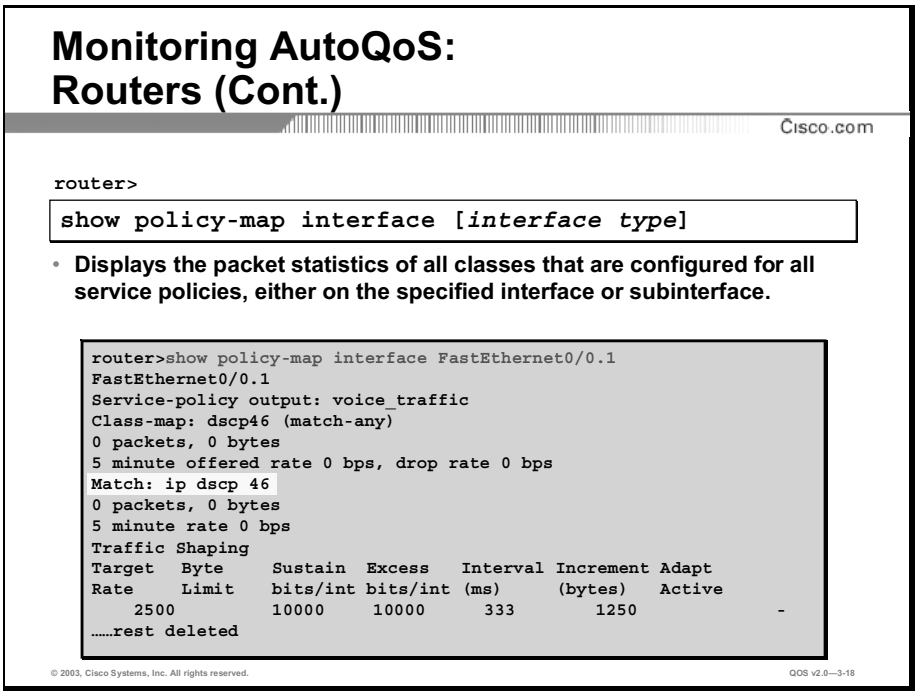

To display the configuration of all classes configured for all service policies on the specified interface, or to display the classes for the service policy for a specific PVC on the interface, use the **show policy-map interface** EXEC or privileged EXEC command.

show **policy-map interface** interface-name **[vc [*vpi*/] *vci*][dlci *dlci*] [input | output]**

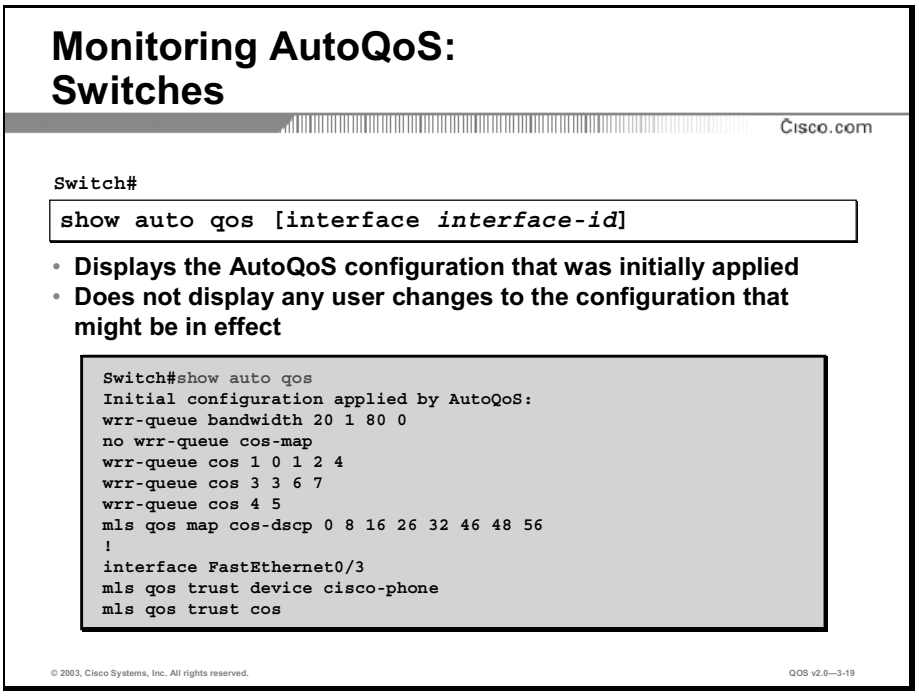

To display the inital AutoQoS configuration, use the **show auto qos** [**interface** [*interface-id*]] privileged EXEC command. To display any user changes to that configuration, use the **show running-config** privileged EXEC command. You can compare the **show auto qos** and the **show running-config** command output to identify the user-defined QoS settings.

Monitoring AutoQoS: Switches (Cont.)

Cisco.com

```
Switch#
show mls qos interface [interface-id | vlan vlan-id]
[buffers | policers | queueing | statistics]
[ | {begin | exclude | include} expression]
```

- **Displays QoS information at the interface level**

```
Switch#show mls qos interface gigabitethernet0/1 statistics
Ingress
    dscp: incoming  no_change     classified      policed        dropped (in bytes)
      1 : 0         0             0               0              0
Others: 203216935   24234242      178982693       0              0

Egress
    dscp: incoming  no_change     classified      policed        dropped (in bytes)
      1 : 0         n/a           n/a             0              0

WRED drop counts:
      qid           thresh1   thresh2   FreeQ
     1 :  0         0         1024
     2 :  0         0         1024
……rest deleted
```


The **show mls qos interface** command is used to display QoS information at the interface level, including the configuration of the egress queues and the CoS-to-egress-queue map, the interfaces that have configured policers, and ingress and egress statistics (including the number of bytes dropped).

If no keyword is specified with the **show mls qos interface** command, the port QoS mode (DSCP trusted, CoS trusted, untrusted, and so forth), default CoS value, DSCP-to-DSCP-mutation map (if any) that is attached to the port, and policy map (if any) that is attached to the interface, are displayed. If a specific interface is not specified, the information for all interfaces is displayed.

Expressions are case sensitive. For example, if you enter | **exclude output**, the lines that contain *output* are not displayed, but the lines that contain *output* are displayed.

Monitoring AutoQoS: Switches (Cont.)

Cisco.com

```
Switch#
show mls qos maps [cos-dscp | dscp-cos | dscp-
mutation dscp-mutation-name | dscp-switch-priority |
ip-prec-dscp | policed-dscp] [ | {begin | exclude |
include} expression
```

- **Maps are used to generate an internal DSCP value, which represents the priority of the traffic.**

```
Switch#show mls qos maps dscp-cos

Dscp-cos map:
   dscp: 0  8 10 16 18 24 26 32 34 40 46 48 56
         ------------------------------------------------
    cos: 0  1  1  2  2  3  3  4  4  5  5  6  7
```


To generate an internal DSCP value representing the priority of the traffic, display the QoS mapping information.

Automation with Cisco AutoQoS

This topic identifies several of the QoS technologies that are automatically implemented on the network when using AutoQoS.

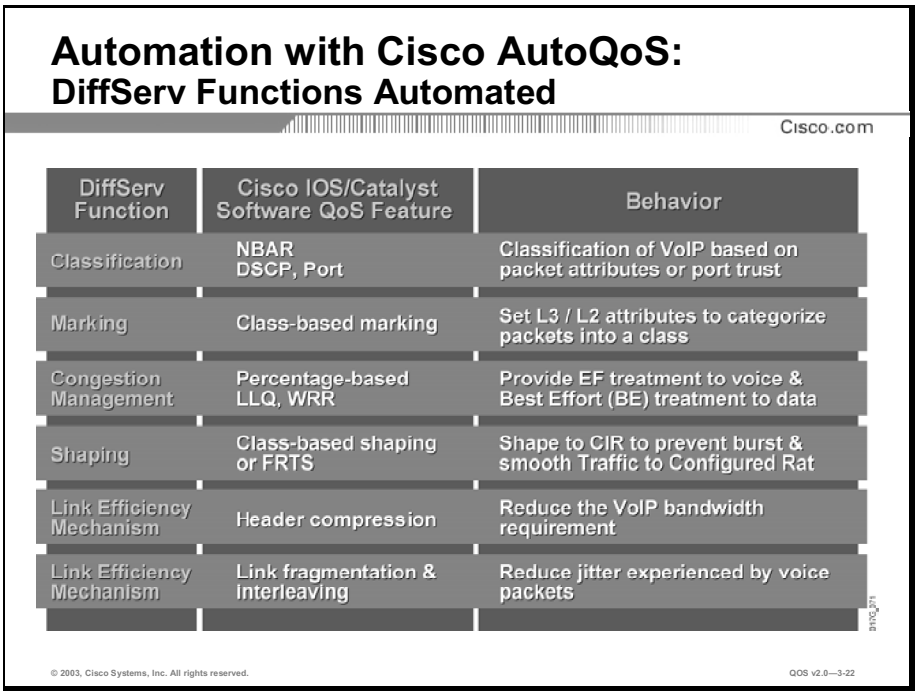

Cisco AutoQoS performs the following functions:

WAN:

- Automatically classify RTP payload and VoIP control packets (H.323, H.225 Unicast, Skinny, session initiation protocol (SIP), Media Gateway Control protocol (MGCP).
- Build service policies for VoIP traffic that are based on Cisco MQC.
- Provision LLQ—priority queuing (PQ) for VoIP bearer and bandwidth guarantees for control traffic.
- Enable WAN traffic shaping that adheres to Cisco best practices, where required.
- Enable link efficiency mechanisms, such as LFI and cRTP where required.
- Provide SNMP and syslog alerts for VoIP packet drops.

LAN:

- Enforce the trust boundary on Cisco Catalyst switch access ports and uplinks/downlinks.
- Enable Cisco Catalyst strict PQ (also known as expedited queuing) with WRR scheduling for voice and data traffic, where appropriate.
- Configure queue admission criteria (map CoS values in incoming packets to the appropriate queues).
- Modify queue sizes and weights where required.

Summary

This topic summarizes the key points discussed in this lesson.

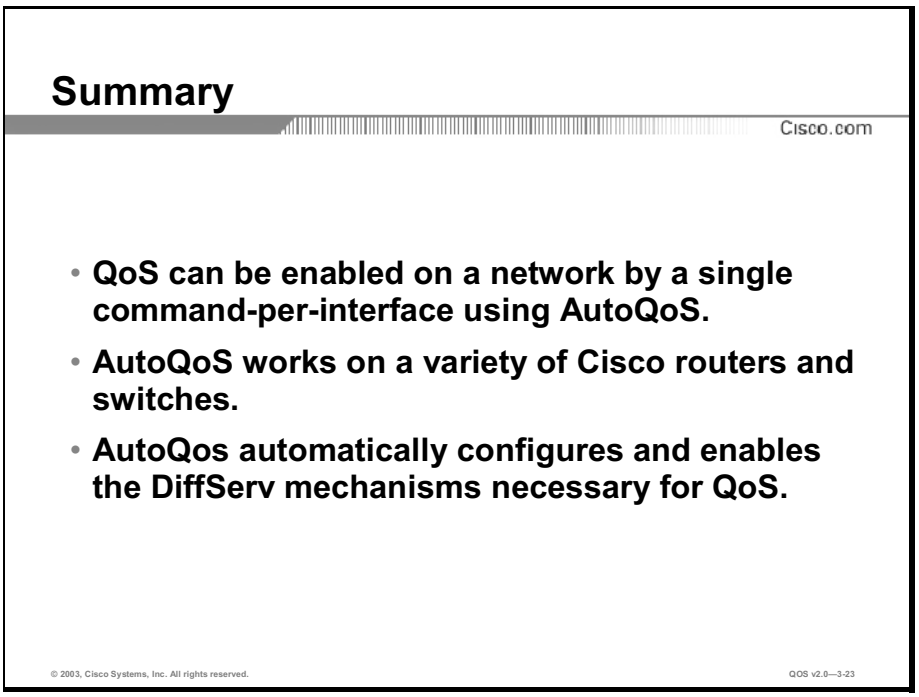

References

For additional information, refer to these resources:

- For more information on Cisco AutoQos, refer to “Cisco AutoQoS Whitepaper” at the following URL: http://www.cisco.com/en/US/tech/tk543/tk759/technologies_white_paper09186a00801348bc.shtml
- For more information on Cisco AutoQos, refer to “Configuring Automatic QoS” at the following URL: http://www.cisco.com/en/US/products/hw/switches/ps708/products_configuration_guide_chapter09186a0080121d11.html#1032637
- For more information on Cisco AutoQos, refer to “Configuring QoS” at the following URL: http://www.cisco.com/en/US/products/hw/switches/ps646/products_configuration_guide_caapter09186a0080115928.html
- For more information on Cisco AutoQos, refer to “AutoQos – VoIP” at the following URL: http://www.cisco.com/en/US/products/sw/iosswrel/ps1839/products_feature_guide09186a0080153ece.html#73342
- For more information on Cisco AutoQos, refer to “Cisco IOS Quality of Service Configuration Guide, Release 12.3” at the following URL: http://www.cisco.com/en/US/products/sw/iosswrel/ps5187/prod_configuration_guide09186a008017d8e5.html

Next Steps

For the associated lab exercise, refer to the following section of the course Lab Guide:

- Lab Exercise 3-1: Configuring QoS with AutoQos

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) For which four of the following scenarios could Cisco AutoQos be most helpful? (Choose four.)

A) small- to medium-sized businesses that need to deploy IP telephony quickly

B) service providers requiring a template-driven approach to delivering managed services

C) international enterprises or service providers requiring QoS for VoIP where little expertise exists in different regions of the world

D) large customer enterprises needing to ensure that the appropriate QoS for voice applications is being set in a consistent fashion

E) service providers with highly specialized customer requirements and large, well-trained staffs

Q2) Which three of the following aspects of QoS deployment does AutoQoS accomplish? (Choose three.)

A) configuration

B) packet marking

C) monitoring and reporting

D) application classification

Q3) Which two of the following series of routers support AutoQoS? (Choose two.)

A) 2600

B) 12000

C) 7600

D) 7200

Q4) Which two of the following series of switches support AutoQoS? (Choose two.)

A) Catalyst 3550

B) Catalyst 1900-E

C) Catalyst 9509

D) Catalyst 6500

Q5) Which Cisco feature must be enabled on a switch to use AutoQoS?

A) CDM

B) ADE

C) QPM

D) CDP

Q6) Why must CEF be enabled to use AutoQoS?

A) to enable fast switching for interface speeds

B) CEF forwarding is used by AutoQoS for configuration

C) AutoQoS examines FIB tables to determine best policies

D) CEF is required for NBAR which is used by AutoQoS

Quiz Answer Key

Q1) A, B, C, D

Relates to: AutoQoS

Q2) A, B, D

Relates to: AutoQoS

Q3) A, D

Relates to: AutoQoS: Router Platforms

Q4) A, D

Relates to: AutoQoS: Switch Platforms

Q5) D

Relates to: AutoQoS: Switch Platforms

Q6) D

Relates to: Configuring AutoQoS

Module Assessment

Overview

Use this assessment to test what you learned in this module. The correct answers and solutions are found in the Module Assessment Answer Key.

Quiz: Introduction to Modular QoS CLI and AutoQoS

Complete the Quiz to assess what you have learned in the module.

Objectives

This activity tests your knowledge on how to meet these objectives:

- Explain how to implement a QoS policy using MQC
- Correctly identify capabilities provided by AutoQoS and successfully configure QoS on a network using AutoQoS

Instructions

Complete these steps:

Step 1 Answer all questions in this quiz by selecting the best answer(s) to each question.

Step 2 Verify your results against the answer key located at the end of this section.

Step 3 Review the topics in this module that relate to the questions that you answered incorrectly.

Q1) Which command would you use to attach a QoS policy to an interface?

A) **policy-set-interface**
B) **policy-map**
C) **policy-interface**
D) **service-policy**

Q2) In what manner can a service policy be attached to an interface?

A) for inbound packets only
B) for outbound packets only
C) for inbound or outbound, not both
D) for inbound only, outbound only, or for both inbound and outbound

Q3) What is “trusted” when the auto **qos voip** command is configured with the “trust” parameter?

A) source address
B) MAC address of sender
C) DES keyword
D) DSCP

Q4) Which three of the following terms are displayed by the **show auto qos interface** command? (Choose three.)

A) ACLs
B) class maps
C) policy maps
D) service maps

Q5) Which command would you use on a Catalyst switch to display the configuration of the egress queues?

A) **show mls qos maps**
B) **show auto qos**
C) **show auto qos inteface**
D) **show mls qos interface**

Q6) Which three of the following does AutoQoS VoIP automatically do when used to automatically configure a WAN interface? (Choose three.)

A) enable payload compression
B) provision Low Latency Queuing (LLQ)
C) automatically classify RTP payload and VoIP control packets
D) enable Link Fragmentation and Interleaving (LFI) where required

Scoring

You have successfully completed the quiz for this lesson when you earn a score of 80 percent or better.

Module Assessment Answer Key

Q1) D

Relates to: Introducing Modular QoS CLI

Q2) D

Relates to: Introducing Modular QoS CLI

Q3) D

Relates to: Introducing AutoQoS

Q4) A, B, C

Relates to: Introducing AutoQoS

Q5) D

Relates to: Introducing AutoQoS

Q6) B, C, D

Relates to: Introducing AutoQoS

Module Summary

This topic summarizes the key points discussed in this module.

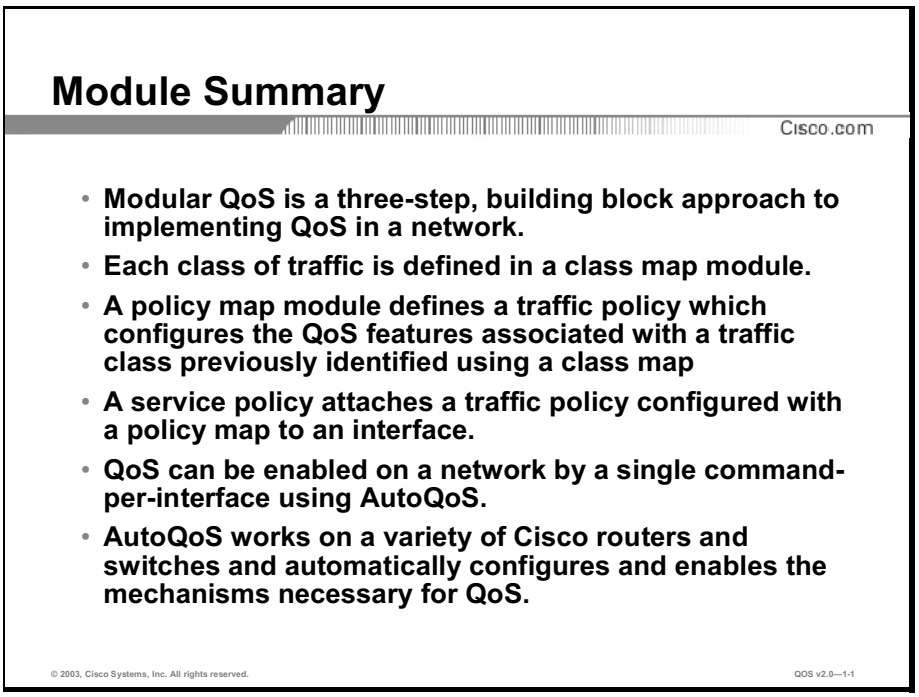

Both the MQC and Cisco AutoQoS were designed to aid in more rapid and consistent design, implementation, and maintenance of QoS policies for converged networks. The MQC offers a three-step, building-block approach to implementing extremely modular QoS policies for network administrators with the requirement to carefully manage large and complex networks. Cisco AutoQoS provides an easy-to-use, mostly automated means to provide consistent QoS policies throughout a network with a minimal design and implementation effort.

Module 4

Classification and Marking

Overview

In any network where networked applications require differentiated levels of service, traffic must be sorted into different classes upon which quality of service (QoS) is applied. Classification and marking are two critical functions of any successful QoS implementation. Classification allows network devices to identify traffic as belonging to a specific class with specific QoS requirements as determined by an administrative QoS policy. After network traffic is sorted, individual packets are colored or marked so that other network devices can apply QoS features uniformly to those packets in compliance with the defined QoS policy.

This module introduces classification and marking and the different methods of performing these critical QoS functions on Cisco routers and switches.

Module Objectives

Upon completing this module, you will be able to successfully classify and mark network traffic to implement a policy defining QoS requirements.

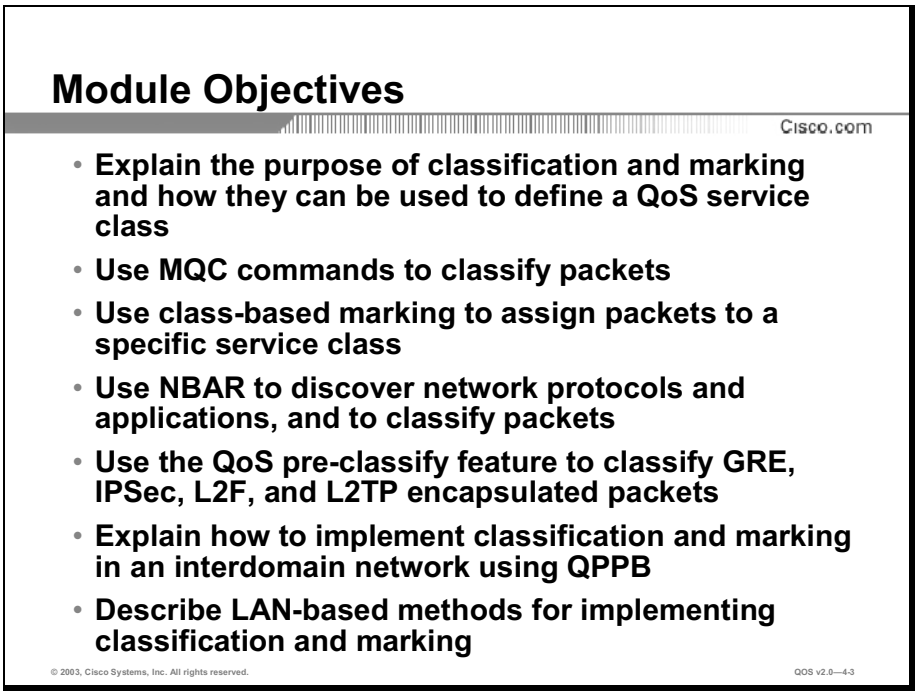

Module Outline

The outline lists the components of this module.

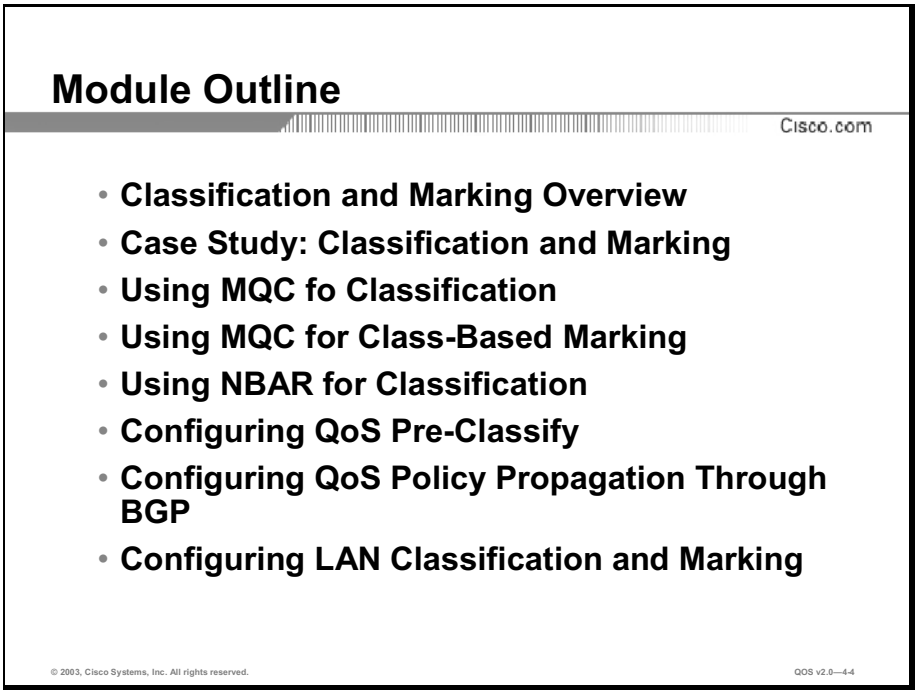

Classification and Marking Overview

Overview

By its very definition, QoS is the ability to provide differential levels of treatment to specific classes of traffic. Before any QoS applications or mechanisms can be applied, traffic must be identified and sorted into different classes. It is upon these different traffic classes to which QoS is applied. Network devices use classification to identify traffic as belonging to a specific class. After network traffic is sorted, marking can be used to color (tag) individual packets so that other network devices can apply QoS features uniformly to those packets as they travel through the network.

This lesson introduces the concepts of classification and marking, explains the different markers that are available at the data link and network layer, and identifies where classification and marking should be used in a network. The concept of a QoS service class and how a service class can be used to represent an application or set of applications is also discussed.

Relevance

Classification and marking are the foundations for any network deployment of QoS. As such, it is of key importance to understand the classification and marking mechanisms and how they are used in implementing QoS.

Objectives

Upon completing this lesson, you will be able to explain the purpose of classification and marking and how they can be used to define a QoS service class. This includes being able to meet these objectives:

- Explain the purpose of packet classification
- Explain the purpose of packet marking
- Describe IP packet classification and marking at the data link layer
- Describe IP packet classification and marking at the network layer
- Describe data link to network layer interoperability between QoS markers
- Define the term “QoS service class” and describe how service classes can be used to create a service policy throughout a network
- Explain how link layer and network layer markings are used to define service classes and the different applications represented by each of these service classes
- Explain the concept of trust boundaries and how they are used with classification and marking

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Basic knowledge of internetworking with TCP/IP concepts

Outline

The outline lists the topics included in this lesson.

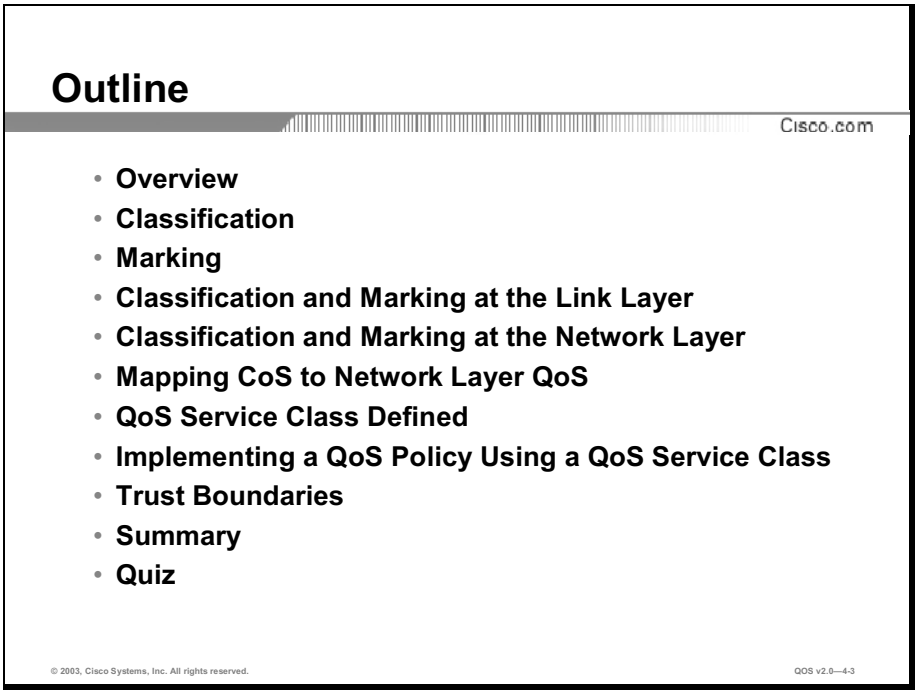

Classification

This topic describes the purpose of packet classification.

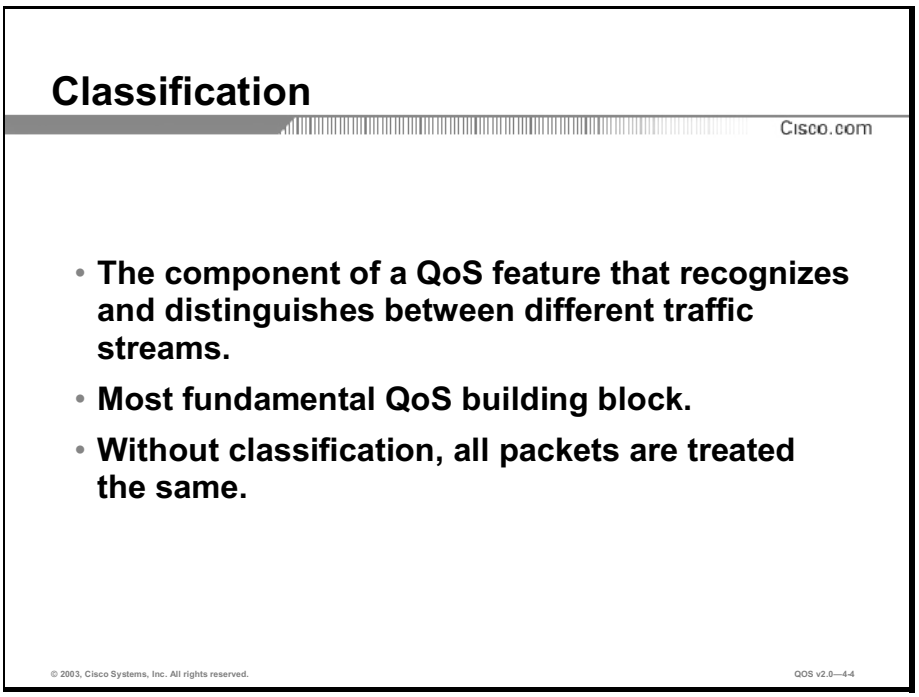

Classification is the process of identifying traffic and categorizing it into different classes. Packet classification uses a traffic descriptor to categorize a packet within a specific group in order to define that packet. Typically used traffic descriptors include: incoming interface, IP precedence, differentiated services code point (DSCP), source or destination address, and application. After the packet has been defined (that is, classified), the packet is then accessible for QoS handling on the network.

Using packet classification, network administrators can partition network traffic into multiple priority levels or classes of service. When traffic descriptors are used to classify traffic, the source agrees to adhere to the contracted terms and the network promises a QoS. Different QoS mechanisms, such as traffic policing, traffic shaping, and queuing techniques, use the traffic descriptor of the packet (that is, the classification of the packet) to ensure adherence to that agreement.

Classification should take place at the network edge, typically in the wiring closet, within IP Phones or at network endpoints. It is preferred that classification occur as close to the source of the traffic as possible.

Marking

This topic describes the purpose of packet marking.

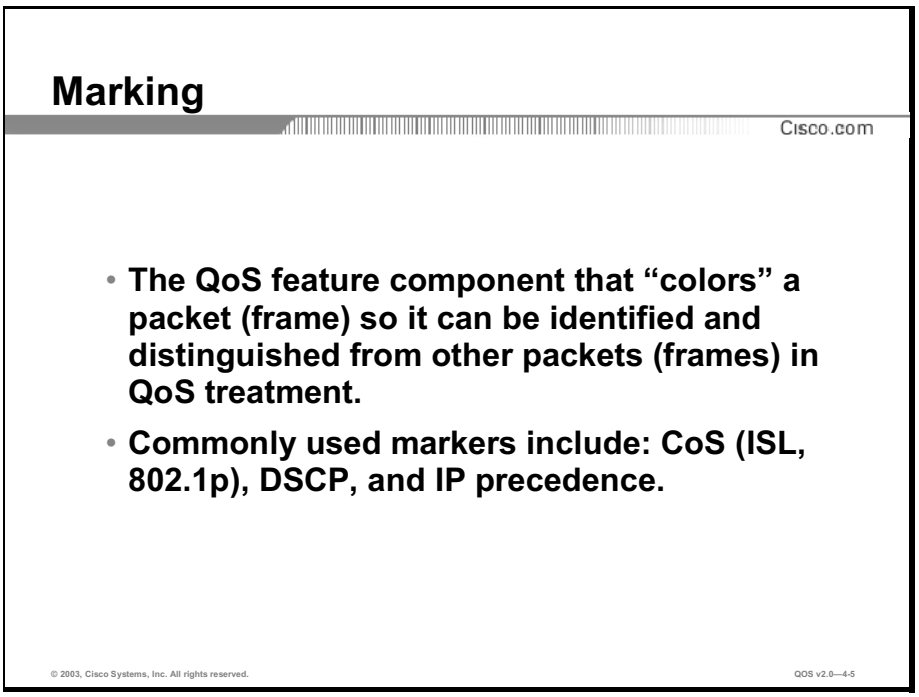

Marking is related to classification. Marking allows network devices to classify a packet or frame based on a specific traffic descriptor. Typically used traffic descriptors include: class of service (CoS), DSCP, IP precedence, QoS group, and Multiprotocol Label Switching (MPLS) experimental bits. Marking can be used to set information in the Layer 2 or Layer 3 packet headers.

Marking a packet or frame with its classification allows network devices to easily distinguish the marked packet or frame. Marking is a useful feature because it allows network devices to easily identify packets or frames as belonging to a specific class. After they are identified as belonging to a specific class, QoS mechanisms can be uniformly applied to ensure compliance with administrative QoS policies.

Classification and Marking at the Link Layer

This topic describes different classification and marking options that are available at the data link layer.

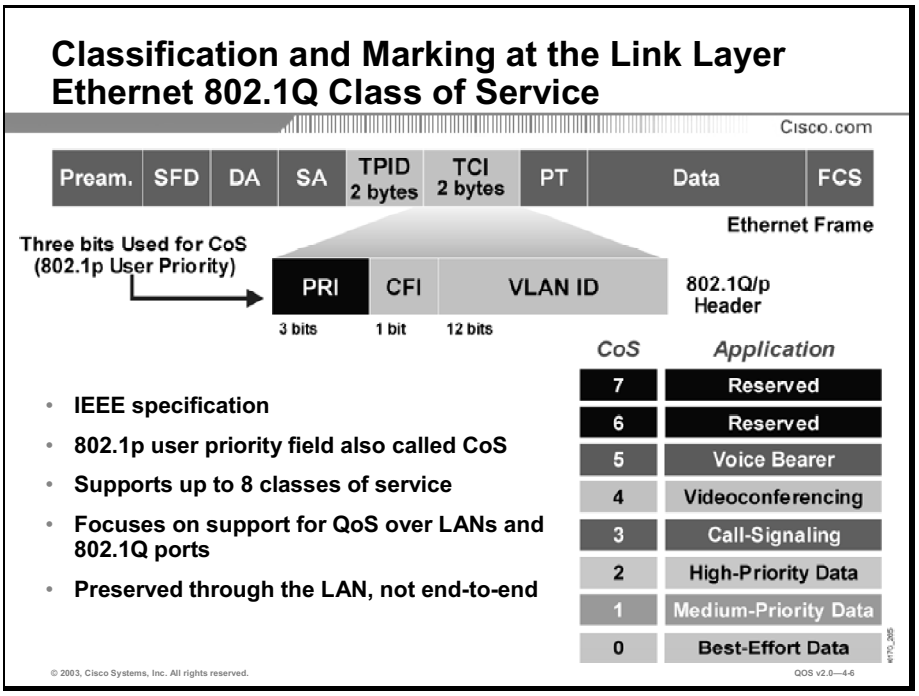

The 802.1Q standard is an IEEE specification for implementing virtual LANs (VLANs) in Layer 2 switched networks. The 802.1Q specification defines two 2-byte fields (Tag Protocol Identifier [TPID]) and Tag Control Information [TCI]) that are inserted within an Ethernet frame following the source address field. The TPID field is currently fixed and assigned the value 0x8100. The TCI field is composed of three fields as follows:

- **User Priority Bits (3 bits):** The specifications of this 3-bit field are defined by the IEEE 802.1p standard. These bits can be used to mark packets as belonging to a specific CoS. The CoS marking uses the three 802.1p user priority bits and allows a Layer 2 Ethernet frame to be marked with 8 different levels of priority (values 0-7). Three bits allow for 8 levels of classification, allowing a direct correspondence with IPv4 (IP precedence) type of service (ToS) values. The IEEE 802.1p specification defines these standard definitions for each CoS:
  - CoS 7 (111): network
  - CoS 6 (110): internet
  - CoS 5 (101): critical
  - CoS 4 (100): flash-override
  - CoS 3 (011): flash
  - CoS 2 (010): immediate
  - CoS 1 (001): priority
  - CoS 0 (000): routine

One disadvantage of using CoS markings is that frames will lose their CoS markings when transiting a non-802.1Q/p link. Therefore, a more ubiquitous permanent marking should be used for network transit. This is typically accomplished through translating a CoS marking into another marker or simply using a different marking mechanism.

- **Canonical Format Identifier (CFI) (1 bit):** This bit indicates whether the bit order is canonical or noncanonical. The CFI bit is used for compatibility between Ethernet and Token Ring networks.
- **VLAN Identifier (VLAN ID) (12 bits):** The VLAN ID field is a 12-bit field that defines the VLAN used by 802.1Q. The fact that the field is 12 bits restricts the number of VLANs supported by 802.1Q to 4096. For most enterprise customers, 4096 VLANs is adequate. For service provider applications, 4096 VLANs may not be enough.

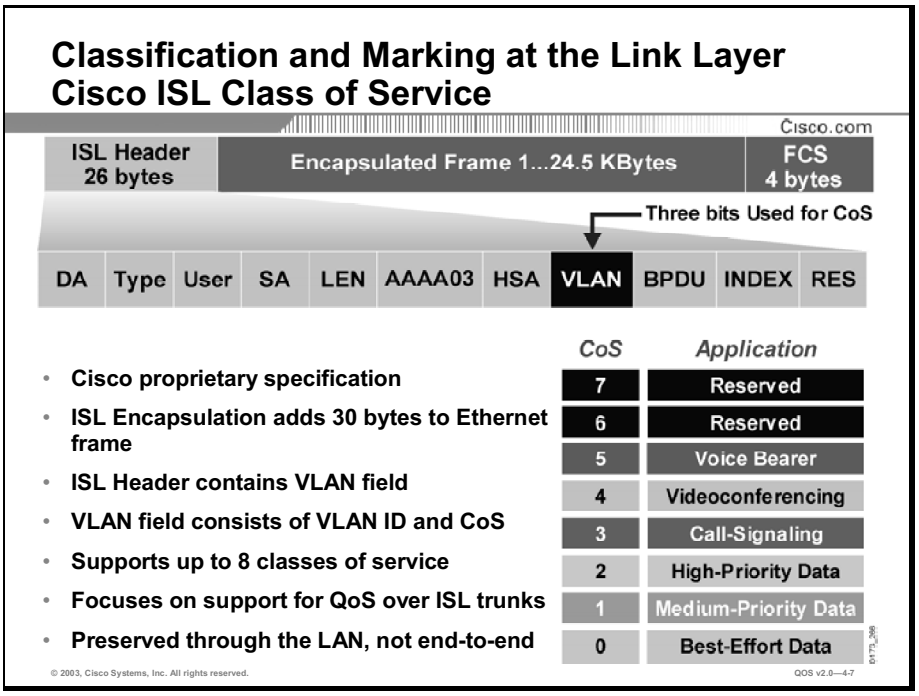

Inter-Switch Link (ISL) is a proprietary Cisco protocol for interconnecting multiple switches and maintaining VLAN information as traffic goes between switches. ISL was created prior to the standardization of 802.1Q. However, ISL is compliant with the IEEE 802.1p standard.

- The ISL frame header contains a 1-byte User field that carries an IEEE 802.1p CoS values in the three least significant bits. When an ISL frame is marked for priority, the three 802.1p CoS bits are set to a value 0-7. In compliance with the IEEE 802.1p specification, ISL follows the standard definitions for each CoS:
  - CoS 7 (111): network
  - CoS 6 (110): internet
  - CoS 5 (101): critical
  - CoS 4 (100): flash-override
  - CoS 3 (011): flash
  - CoS 2 (010): immediate
  - CoS 1 (001): priority
  - CoS 0 (000): routine

Similar to 802.1Q, ISL CoS markings are not maintained end-to-end if a non-ISL or 802.1Q trunk is transited. As a result, network administrators typically translate CoS markings into another marker or simply use a different marking mechanism altogether.

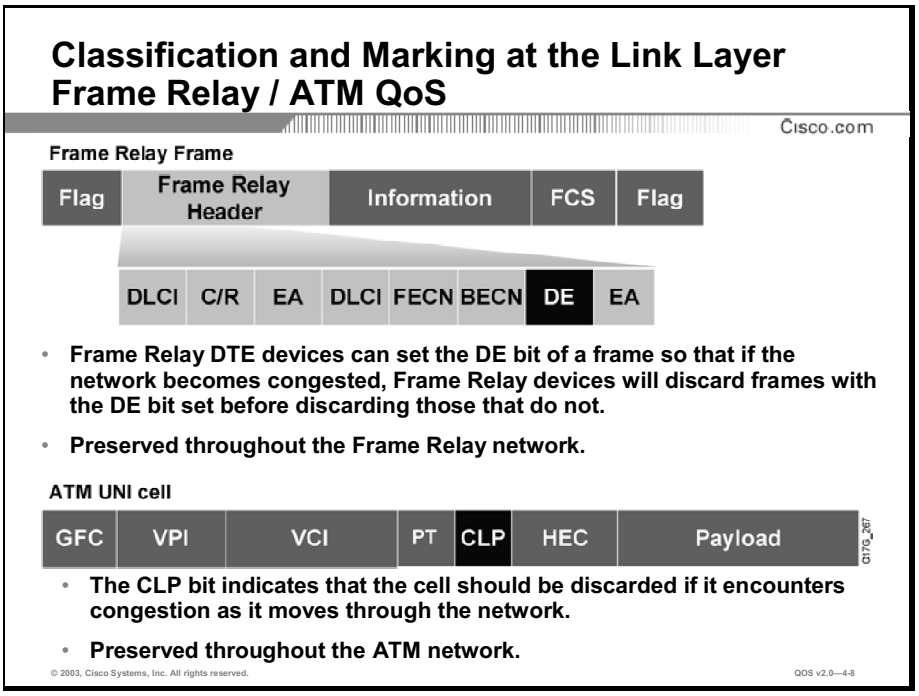

Before the Internet Engineering Task Force (IETF) defined QoS methods for the network layer, the ITU-T (International Union for Telecommunications), ATM Forum, and the Frame Relay Forum (FRF) had already derived standards for link layer QoS in ATM and Frame Relay networks.

ATM standards define a very rich QoS infrastructure by supporting traffic contracts, many adjustable QoS knobs (such as peak cell rate [PCR], minimum cell rate [MCR], and so on), signaling, and admission control. Frame Relay provides a simpler set of QoS mechanisms to ensure a committed information rate (CIR), congestion notification, and Frame Relay fragmentation (FRF.12).

One component of Frame Relay QoS is packet discard when congestion is experienced in the network. Frame Relay will allow network traffic to be sent at a rate exceeding its CIR. Frames sent that exceed the committed rate can be marked as discard eligible (DE). If congestion occurs in the network, frames marked DE will be discarded prior to discarding frames that do not.

ATM cells consist of 48 bytes of payload and 5 bytes of header. The ATM header includes the 1-bit cell loss priority (CLP) field, which indicates the drop priority of the cell if it encounters extreme congestion as it moves through the ATM network. The CLP bit represents two values: 0 to indicate higher priority and 1 to indicate lower priority. Setting the CLP bit to 1 lowers the priority of the cell, increasing the likelihood that the cell will be dropped when the ATM network experiences congestion.

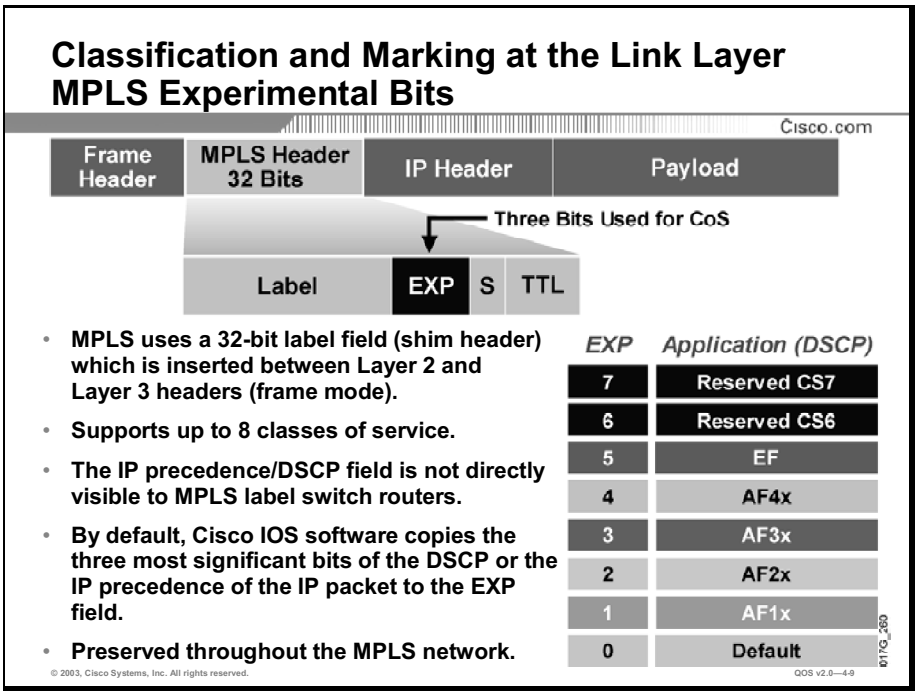

When a customer transmits IP packets from one site to another, the IP precedence field (the first 3 bits of the DSCP field in the header of an IP packet) specifies the CoS. Based on the IP precedence marking, the packet is given the desired treatment, such as guaranteed bandwidth or latency. If the service provider network is an MPLS network, then the IP precedence bits are copied into the MPLS experimental field at the edge of the network. However, the service provider might want to set an MPLS packet QoS to a different value that is determined by the service offering.

The MPLS experimental field allows the service provider to provide QoS without overwriting the value in the customer IP precedence field. The IP header remains available for customer use; the IP packet marking is not changed as the packet travels through the MPLS network.

Classification and Marking at the Network Layer

This topic describes the different classification and marking options that are available at the network layer.

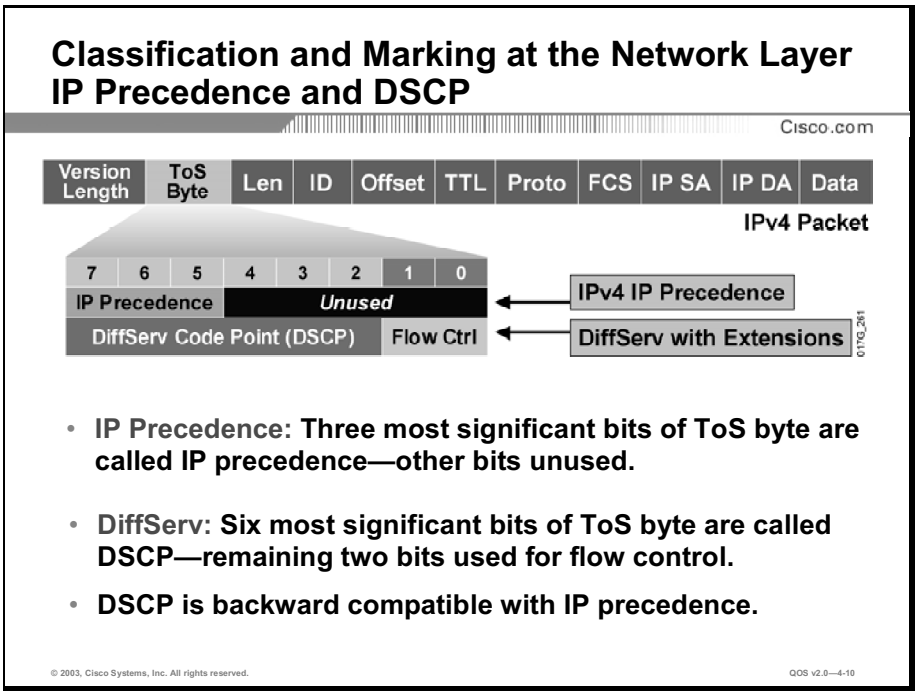

At the network layer, IP packets are typically classified based on source or destination IP address, packet length, or the contents of the ToS byte. Link layer media often changes as a packet travels from its source to its destination. Because a CoS field does not exist in a standard Ethernet frame, CoS markings at the link layer are not preserved as packets traverse the network. Using marking at the network layer provides a more permanent marker that is preserved from source to destination. The network layer markers most typically used are IP precedence and DSCP.

The header of an IPv4 packet contains the ToS byte. IP precedence uses three precedence bits in the ToS field of the IPv4 header to specify CoS for each packet. IP precedence values range from 0 to 7 and allow network administrators to partition traffic in up to six useable classes of service. (Settings 6 and 7 are reserved for internal network use.)

Differentiated Services (DiffServ) is a new model that supersedes—and is backward compatible with—IP precedence. DiffServ redefines the ToS byte and uses six prioritization bits that permits classification of up to 64 values (0 to 63) of which 32 are commonly used. A DiffServ value is called a DSCP.

With DiffServ, packet classification is used to partition network traffic into multiple priority levels or classes of service. Packet classification uses the DSCP traffic descriptor to categorize a packet within a specific group to define that packet. After the packet has been defined (classified), the packet is then accessible for QoS handling on the network.

Mapping CoS to Network Layer QoS

This topic describes the different QoS markers that can be used for interoperability between data link layer and network layer QoS.

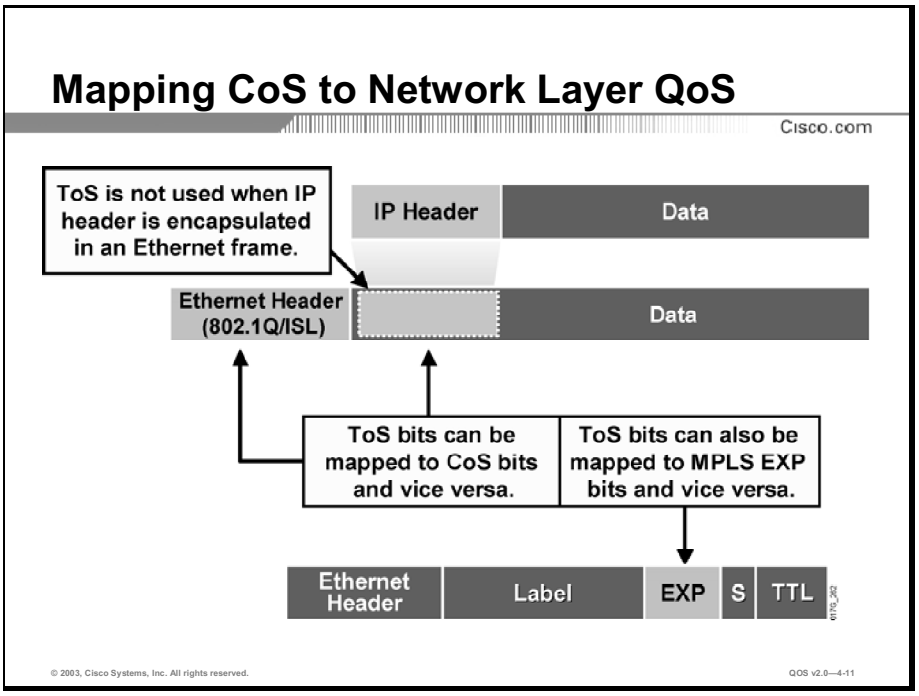

IP headers are preserved end-to-end when IP packets are transported across a network; data link layer headers are not. This means that the IP layer is the most logical place to mark packets for end-to-end QoS. However, there are edge devices that can only mark frames at the data link layer and there are many other network devices that only operate at the data link layer. To provide true end-to-end QoS, the ability to map QoS marking between the data link layer and the network layer is essential.

Enterprise networks typically consist of a number of remote sites connected to the headquarters campus via a WAN. Remote sites typically consist of a switched LAN, and the headquarters campus network is both routed and switched. Providing end-to-end QoS through such an environment requires that CoS markings that are set at the LAN edge be mapped into QoS markings (such as IP precedence or DSCP) for transit through Campus or WAN routers. Campus and WAN routers can also map the QoS markings to new data link headers for transit across the LAN. In this way, QoS can be preserved and uniformly applied across the enterprise.

Service providers offering IP services have a requirement to provide robust QoS solutions to their customers. The ability to map network layer QoS to link layer CoS allows these providers to offer a complete end-to-end QoS solution that does not depend on any specific link layer technology.

Compatibility between an MPLS transport and network layer QoS is also achieved by mapping between MPLS experimental (EXP) bits and the IP precedence or DSCP bits. A service provider can map the customer network layer QoS marking as-is, or change them to fit an agreed upon service level agreement (SLA). The information in the MPLS EXP bits can be

carried end-to-end in the MPLS network, independent of the transport media. In addition, the network layer marking can remain unchanged so that when the packet leaves the service provider MPLS network, the original QoS markings remain intact. Thus, a service provider with an MPLS network can help provide a true end-to-end QoS solution.

QoS Service Class Defined

This topic defines the term, “QoS service class” and describes how service classes can be used to create a service policy throughout a network.

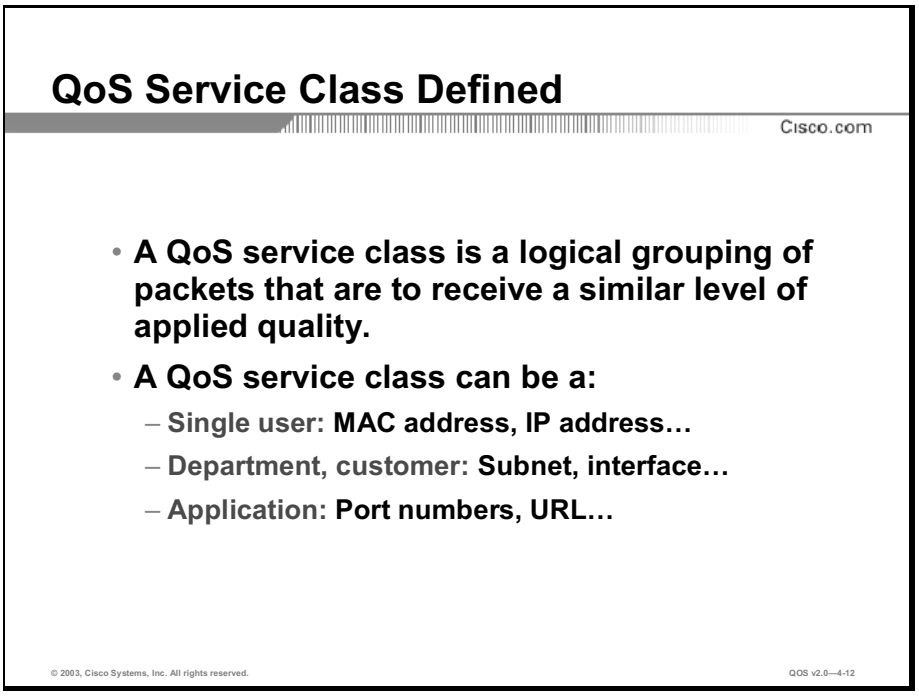

When an administrative policy requiring QoS is created, it must be determined how network traffic is to be treated. As part of that policy definition, network traffic must be associated with a specific service class. QoS classification mechanisms are used to separate traffic and identify packets as belonging to a specific service class. QoS marking mechanisms are used to tag each packet as belonging to the assigned service class. After the packets are identified as belonging to a specific service class, QoS mechanisms such as policing, shaping, queuing techniques can be applied to each service class to meet the specifications of the administrative policy. Packets belonging to the same service class are given the same treatment with regards to QoS.

A QoS service class, being a logical grouping, can be defined in many ways, some of which include:

- Organization or department (Marketing, Engineering, Sales, and so on)
- A specific customer or set of customers
- Specific applications or set of applications (Telnet, FTP, Voice, SAP, Oracle, Video, and so on)
- Specific users or sets of users (based on MAC address, IP address, LAN port, and so on.)
- Specific network destinations (tunnel interfaces, Virtual Private Networks [VPNs], and so on)

Example: Defining QoS Service Classes

A network administrator wishes to apply QoS to the corporate network to better control bandwidth allocation of different network applications. Before QoS can be applied, an administrative QoS policy is first devised as follows:

- Voice traffic is to be given a strict priority over all other traffic types.
- Business applications (FTP, TN3270, and Oracle) should be given priority over web traffic and have a guaranteed bandwidth of 20 percent.
- Web traffic should consume no more than 30 percent of any WAN link.

As a result of this policy, three QoS service classes have been defined:

- **Voice class:** To be treated with a strict priority service.
- **Business applications class:** Requires a guaranteed bandwidth of 20 percent and is to be given priority over web traffic.
- **Web class:** Only allowed to consume up to 30 percent of any WAN link.

Implementing a QoS Policy Using a QoS Service Class

This topic describes how link layer and network layer markers are used to define QoS service classes and the different applications that can be represented by each of these service classes.

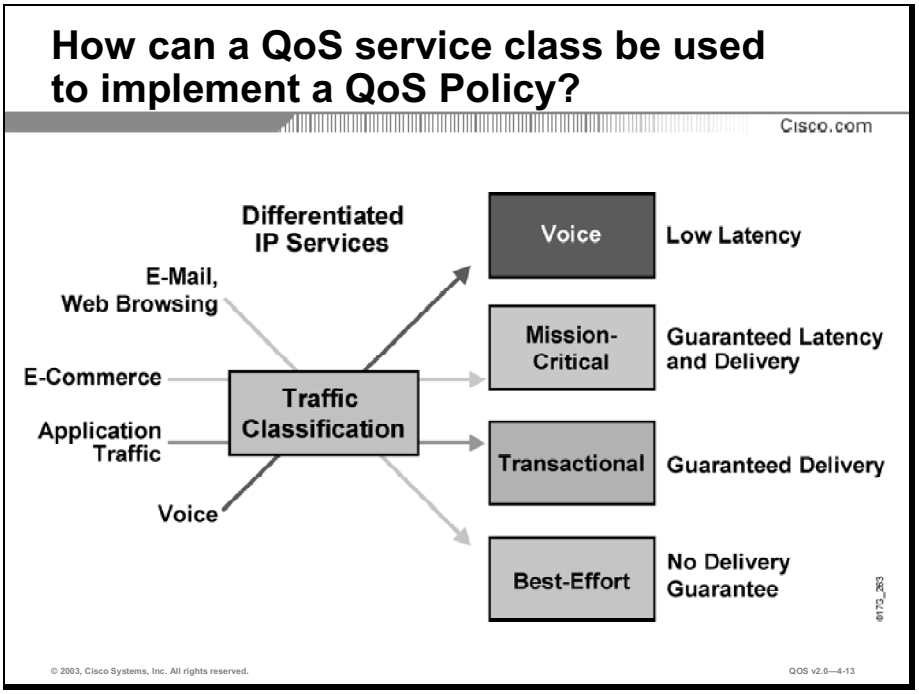

Specifying an administrative policy for QoS requires that a specific set of service classes be defined. QoS mechanisms are uniformly applied to these individual service classes to meet the requirements of the administrative policy. Because the application of QoS mechanisms is applied to different service classes and used to differentiate between applications, users, and traffic, the service class is a key component of a successful QoS implementation.

There are many different methods in which service classes can be used to implement an administrative policy. The first step is to identify the traffic that exists in the network and the QoS requirements for each traffic type. Then, traffic can be grouped into a set of service classes for differentiated QoS treatment in the network.

One popular model for the application of QoS service classes is the customer model which is typically used by service providers when referring to customer traffic. The customer model defines the following service classes (although many variations exist):

- **Voice service class:** Delivers low latency for voice services.
- **Mission-critical service class:** Guarantees latency and delivery for the transport of mission-critical business applications like SNA.
- **Transactional service class:** Guarantees delivery and is used for more general applications that are not as sensitive to delay, like e-commerce.
- **Best-effort service class:** Used to support small business and e-mail and other best-effort applications.

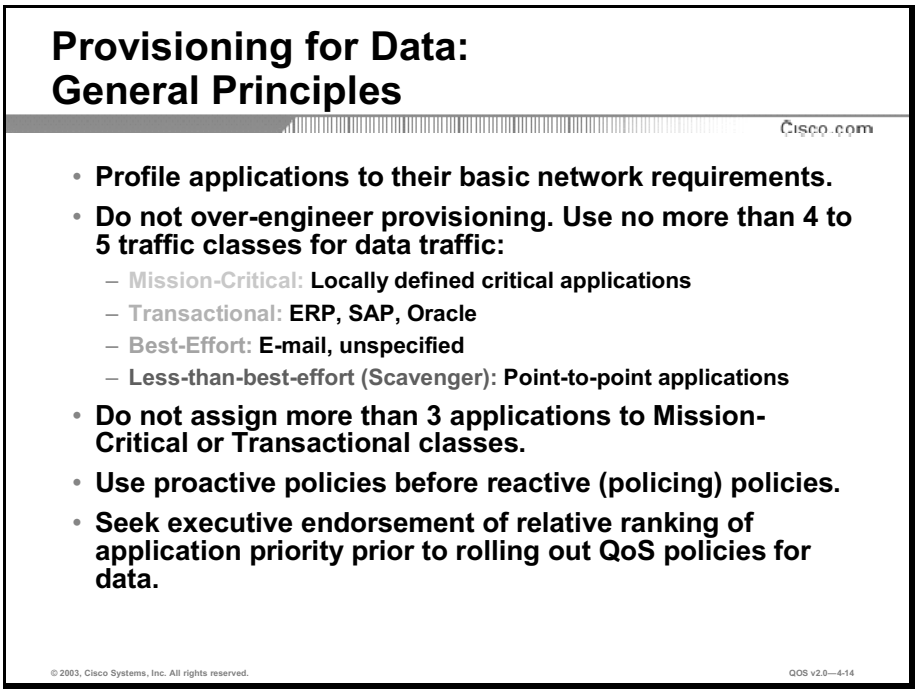

One key element of defining QoS service classes is to understand the basic quality needs of network applications. It is essential that applications be given QoS treatment inline with their needs. For example, improperly specifying voice traffic into a service class with guaranteed bandwidth—without a guaranteed latency (delay)—would not meet the needs of the voice traffic.

While it is important to fully understand network application requirements, it is equally important not to over-provision or over design the administrative policy. An administrative policy should be proactive in nature and require as few service classes as possible. One good rule is to limit the number of service classes to no more than four or five. A typical network has the following application types:

- Mission critical applications—Oracle, SAP, SNA
- Interactive applications—Telnet, TN3270
- Bulk applications—FTP, TFTP, database synchronization/backup
- Best-effort applications—e-mails, Web
- Scavenger applications—Naspter, Kazaa

The QoS requirements of these applications can be met with a few well-designed service classes. The more service classes implemented in support of an administrative QoS policy, the more complex the QoS implementation will be. This complexity also extends to support and troubleshooting as well.

It is also important that the highest-priority classes be reserved for a select few number of applications. Marking 90 percent of network traffic as high priority will render most administrative QoS policies useless.

Example Application Service Classes

Cisco.com

Application	L3 Classification			L2	L2
	IPP	PHB	DSCP	CoS	MPLS EXP
Reserved	7	-	56-63	7	7
Reserved	6	-	48-55	6	6
Voice Bearer	5	EF	46	5	5
Videoconferencing	4	AF41	34	4	4
Mission-Critical Data	3	AF31	26	3	3
Transactional Data	2	AF2x	18,20,22	2	2
Bulk Data	1	AF1x	10,12,14	1	1
Best-Effort Data	0	BE	0	0	0
Less-than-Best-Effort Data	0	-	2,4,6	0	0

0170_264

QOS v2.0—4-15

Although there are several sources of information that can be used as guidelines for determining a QoS policy, none of them can determine exactly what is proper for a specific network. Each network presents its own unique challenges and administrative policies. To properly implement QoS, measurable goals must be declared. Then a plan for achieving these goals must be formulated and implemented.

QoS must be implemented consistently across the entire network. It is not so important whether call signaling is marked as DSCP 34 or 26, but rather that DSCP 34 is treated in a manner that is necessary to accomplish the QoS policy. It is also important that data marked DSCP 34 is treated consistently across the network. If data travels over even a small portion of a network where different policies are applied (or no policies are applied), the entire QoS policy is nullified. Whether the data is crossing slow WAN links or gigabit Ethernet, whether it is being switched by a Layer 2 switch or routed in a Layer 3 router, the policies must be implemented in a way that causes a consistent effect that satisfies the policy requirements.

Trust Boundaries

This topic describes the concept of trust boundaries and how they are used with classification and marking.

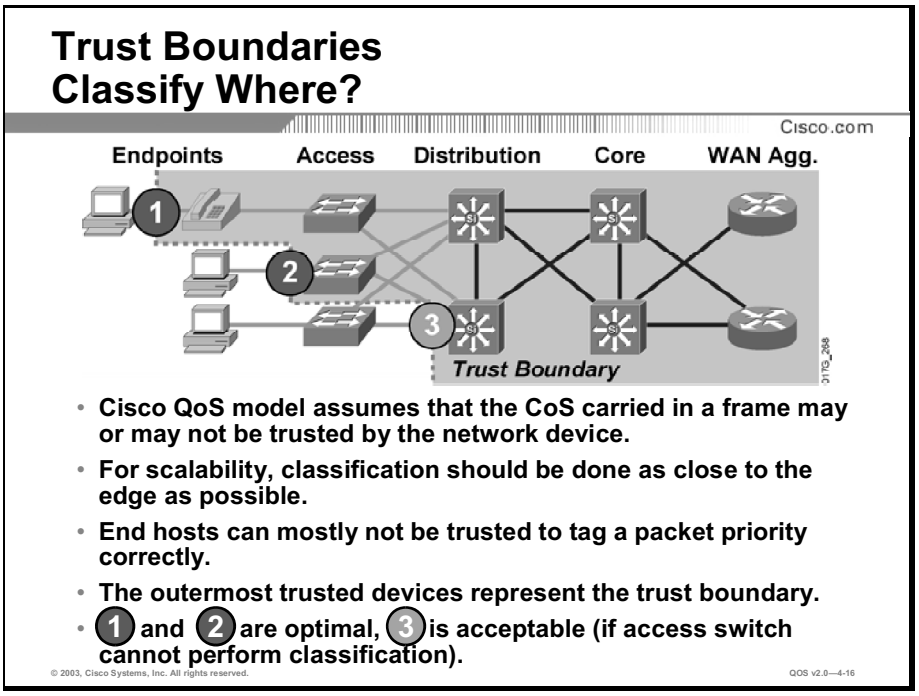

The concept of trust is important and integral to deploying QoS. After the end devices have set CoS or ToS values, the switch has the option of trusting them. If the switch trusts the values, it does not need to reclassify; if it does not trust the values, then it must perform reclassification for the appropriate QoS.

The notion of trusting or not trusting forms the basis for the trust boundary. Ideally, classification should be done as close to the source as possible. If the end device is capable of performing this function, the trust boundary for the network is at the end device. If the device is not capable of performing this function, or the wiring closet switch does not trust the classification done by the end device, the trust boundary might shift. How this shift happens depends on the capabilities of the switch in the wiring closet. If the switch can reclassify the packets, the trust boundary is in the wiring closet. If the switch cannot perform this function, the task falls to other devices in the network, going toward the backbone. In this case, one good rule is to perform reclassification at the distribution layer. This means that the trust boundary has shifted to the distribution layer. It is likely that there is a high-end switch in the distribution layer with features to support this function. If possible, try to avoid performing this function in the core of the network.

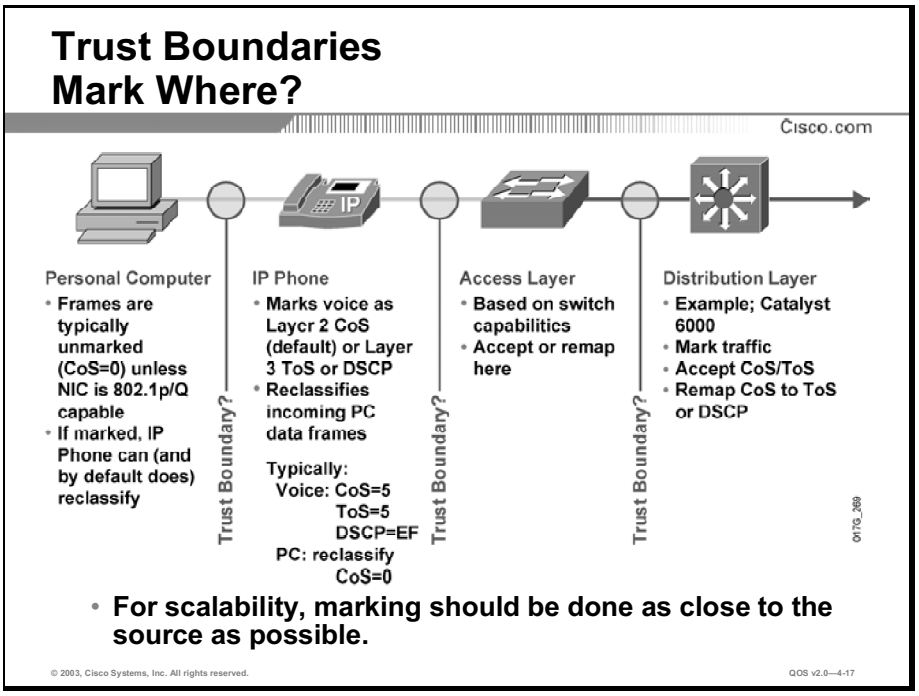

Classification should take place at the network edge, typically in the wiring closet or within endpoints (servers, hosts, video endpoints, or IP telephony devices) themselves.

For example, consider the campus network containing IP telephony and host endpoints. Frames can be marked as important by using link layer CoS settings or the IP precedence/DSCP bits in the ToS/DS field in the IPv4 header. Cisco IP Phones can mark voice packets as high priority using CoS as well as ToS. By default, the IP Phone sends 802.1p tagged packets with the CoS and ToS set to a value of 5 for its voice packets. Because most PCs do not have an 802.1Q capable network interface card (NIC), they send packets untagged. This means that the frames do not have an 802.1p field. Also, unless the applications running on the PC send packets with a specific CoS value, this field is zero.

Note:	A special case exists where the TCP/IP stack in the PC has been modified to send all packets with a ToS value other than zero. Typically this does not happen, and the ToS value is zero.

Even if the PC is sending tagged frames with a specific CoS value, Cisco IP Phones can zero out this value before sending the frames to the switch. This is the default behavior. Voice frames coming from the IP Phone have a CoS of 5 and data frames coming from the PC have a CoS of 0.

If the end device is not a trusted device, the reclassification function (setting/zeroing the bits in the CoS and ToS fields) can be performed by the access layer switch if that device is capable of doing so. If the device is not capable, then the reclassification task falls to the distribution layer device. If reclassification cannot be performed at one of these two layers, a hardware and/or Cisco IOS software upgrade may be necessary.

Summary

This topic summarizes the key points discussed in this lesson.

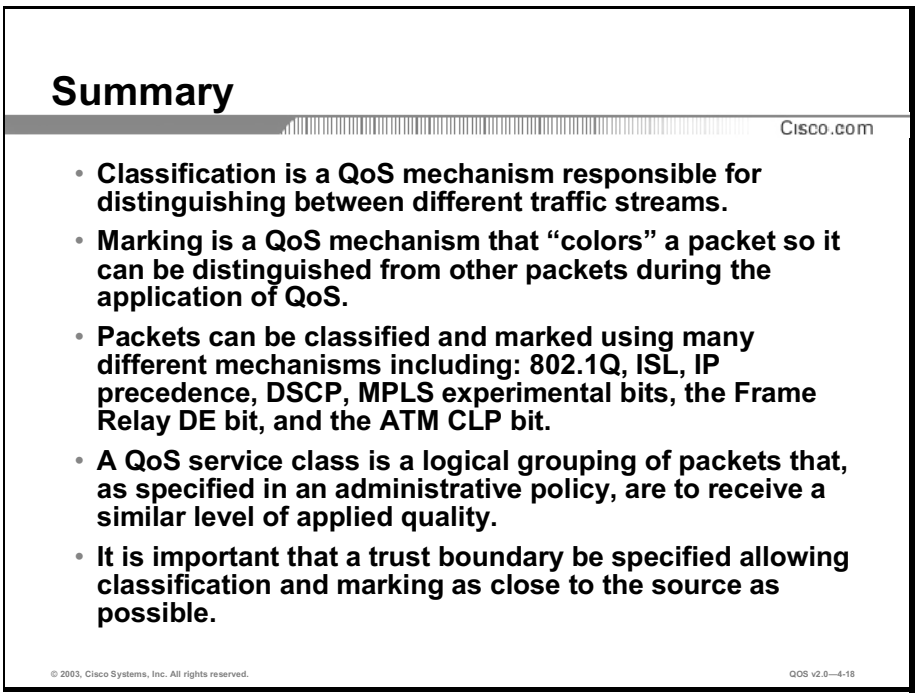

References

For additional information, refer to these resources:

- For an overview of classification, refer to "Classification Overview" at the following URL: http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122cgcr/fqos_c/fqcprt1/qcfclass.htm#wp1000872
- For additional information on 802.1p/Q marking, refer to "Bridging Between IEEE 802.1Q VLANs" at the following URL: http://www.cisco.com/univercd/cc/td/doc/product/software/ios121/121newft/121t/121t3/dtbridge.htm#xtocid114535
- For additional information on ISL marking, refer to "Configuring Routing between VLANs with ISL Encapsulation" at the following URL: http://www.cisco.com/univercd/cc/td/doc/product/software/ios113ed/113ed_cr/switch_c/xcisl.htm
- For additional information on ISL marking, refer to "Configuring QoS: Understanding How QoS Works" at the following URL: http://www.cisco.com/univercd/cc/td/doc/product/lan/cat6000/sw_5_5/cnfg_gd/qos.htm
- For additional information on DiffServ, refer to "DiffServ—The Scalable End-to-End QoS Model" at the following URL: http://www.cisco.com/warp/public/cc/pd/iosw/ioft/iofwft/prodlit/difse_wp.htm

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) What is the main purpose of the QoS classification mechanism?

A) to set fields in IP packets and identify that packet as belonging to a specific class of service

B) to signal network devices on which QoS mechanism should be employed to meet the requirements of a particular service class

C) to identify traffic as belonging to a specific class of service

D) to provide a mapping between link layer and network layer QoS

Q2) What are two main purposes of the QoS marking mechanism? (Choose two.)

A) to sort traffic into different service classes for QoS treatment

B) to signal network devices on which QoS mechanism should be employed to meet the requirements of a particular service class

C) to set fields in IP packets and identify that packet as belonging to a specific class of service

D) to allow edge devices to select the QoS level on an application-by-application basis

Q3) What are three QoS markers commonly used at the link layer? (Choose three.)

A) ISL

B) DSCP

C) 802.1p

D) IP precedence

E) Frame Relay DE bits

Q4) What are two QoS markers commonly used at the network layer? (Choose two.)

A) ISL

B) DSCP

C) MPLS experimental bits

D) IP precedence

Q5) When 802.1p CoS marking is used, what ensures end-to-end QoS? (Choose two.)

A) classification should be done as close to the edge as possible

B) the 802.1Q header travels with the packet from the source to the destination

C) core devices must be capable of performing the marking and policing functions

D) using marking at the network layer provides a more permanent marker that is preserved from source to destination

Q6) What is a QoS service class?

A) an applied per-hop behavior

B) a logical grouping of packets that are to receive similar QoS treatment

C) a mechanism for changing packet markings for trusted and non-trusted packets

D) a method of providing a mapping between link layer and network layer QoS

Q7) In which two scenarios would it be acceptable to place the trust boundary at the distribution layer? (Choose two.)

A) when the access layer device is not capable of performing this function

B) when the wiring closet switch does not trust the classification done by the end device

C) when the application in use is not of high importance and therefore marking is not required

D) when the network is very large and there are too many access layer switches to properly configure the marking capabilities

Quiz Answer Key

Q1) C

Relates to: Classification

Q2) B, C

Relates to: Marking

Q3) A, C, E

Relates to: Classification and Marking at the Link Layer

Q4) B, D

Relates to: Classification and Marking at the Network Layer

Q5) A, D

Relates to: Mapping CoS to Network Layer QoS

Q6) B

Relates to: QoS Service Class Defined

Q7) A, B

Relates to: Trust Boundaries

Case Study: Classification and Marking

Overview

This case study activity provides information regarding the QoS administrative policy requirements of a large, multisite network. Your task is to work with a partner to evaluate the QoS requirements, and based on these requirements, identify where QoS classification and marking mechanisms should be applied. You will discuss your solution with the instructor and other classmates, and the instructor will present a solution for the case study to the class.

Relevance

The ability to properly sort traffic into service classes is an important step in correctly implementing an administrative QoS policy.

Objectives

In this activity, you will define a QoS policy that assigns network traffic to service classes and identify where classification and marking should be applied to the network. Upon completing this case study, you will be able to meet these objectives:

- Review customer QoS requirements
- Identify QoS service class requirements
- Identify network locations where classification and marking should be applied
- Present a solution to the case study

Learner Skills and Knowledge

To benefit fully from this activity, you must have these prerequisite skills and knowledge:

- Basic knowledge of internetworking with TCP/IP concepts

Outline

The outline lists the topics included in this activity.

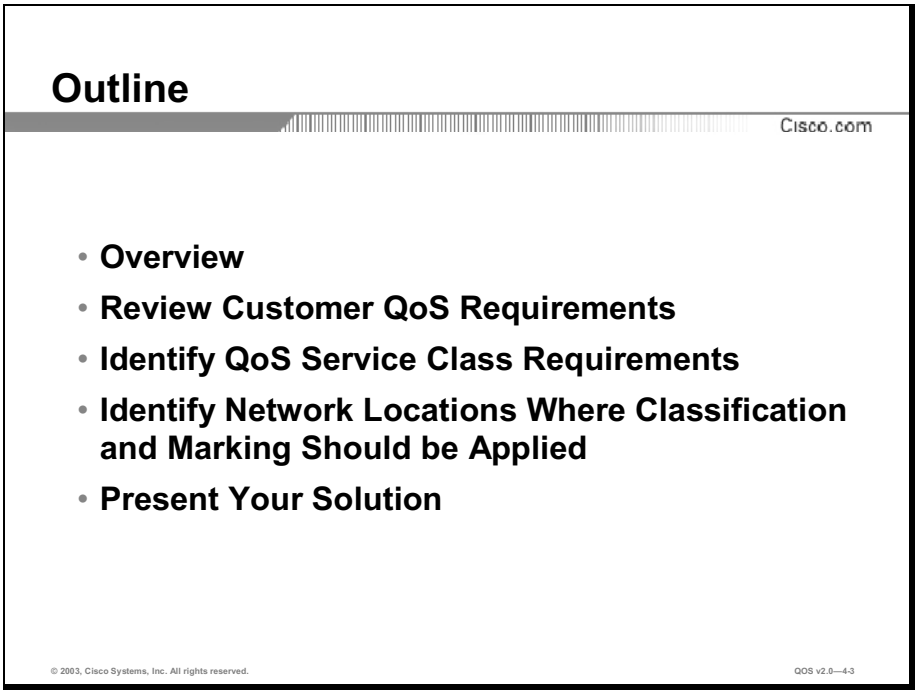

Required Resources

These are the resources required to complete this exercise:

- Case Study Activity: Classification and Marking with QoS Service Classes
- A workgroup consisting of two learners

Job Aids

No job aids are required to complete this case study.

Case Study Tasks

The activity includes these tasks:

Step 1 **Review customer QoS requirements:** Completely read the customer requirements provided.

Step 2 **Identify QoS service class requirements:** With the aid of your partner, identify the service classes required to implement the administrative QoS policy based on customer requirements.

Step 3 **Identify network locations where classification and marking should be applied:** Identify locations in the network where the QoS classification and marking mechanisms should be applied to properly implement the administrative QoS policy.

Step 4 **Present your solution:** After the instructor presents a solution to the case study, present your solution to the class with your partner.

Case Study Verification

You have completed this activity when your case study solution has been presented to the class and you have justified any major deviations from the case study solution supplied by the instructor.

Review Customer QoS Requirements

This case study involves analyzing an administrative QoS policy of the JC Whitney Corporation, a fictitious manufacturer of medical equipment. The company has provided you with a brief description of their requirements. It is your task to provide the network engineers from JC Whitney a QoS solution to meet their requirements.

Read the customer requirements and discuss them with your partner. Identify the different types of traffic in use in the JC Whitney network and the different service classes required to implement their administrative QoS policy.

Company Background

JC Whitney Corporation is a leading manufacturer of medical equipment used in outpatient surgical centers throughout the United States. The company headquarters are located in Eugene, Oregon. The JC Whitney corporate network is shown in the figure.

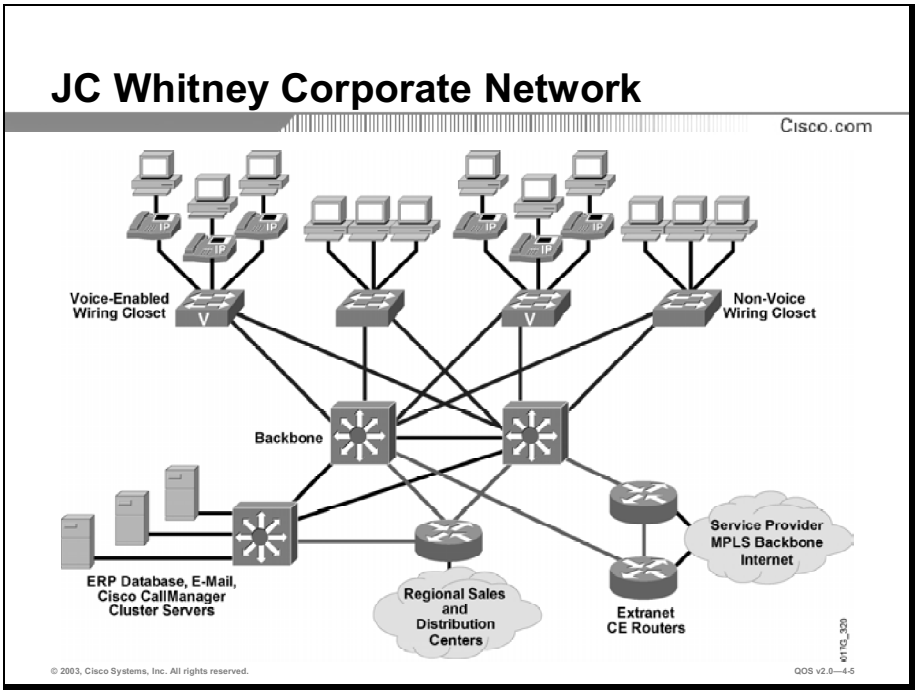

In addition to the headquarters facility, JC Whitney consists of 5 manufacturing facilities and 120 regional sales and distribution centers. The network at each of the manufacturing facilities is similar to the JC Whitney corporate network. The regional sales and distribution centers are very low-cost, low-overhead sites.

The regional sales and distribution center networks are shown in the figure.

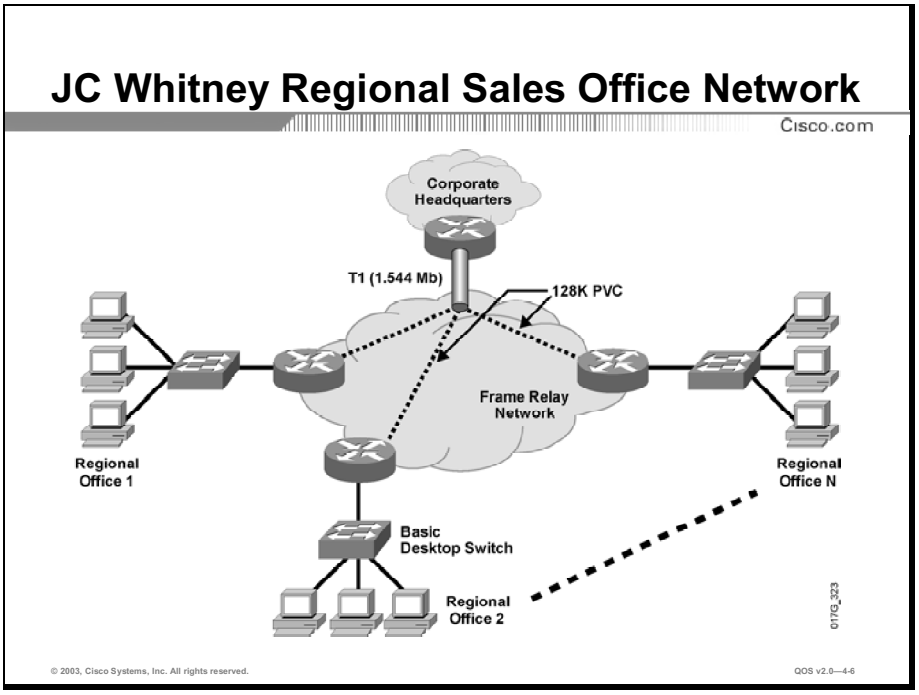

The manufacturing strategy of JC Whitney is to leverage the expertise of contract manufacturers through its extensive extranet of partners. Currently, the JC Whitney extranet consists of nine contract manufacturers and suppliers that are all connected using a national service provider backbone.

The provider currently utilizes MPLS on its backbone as shown in the figure.

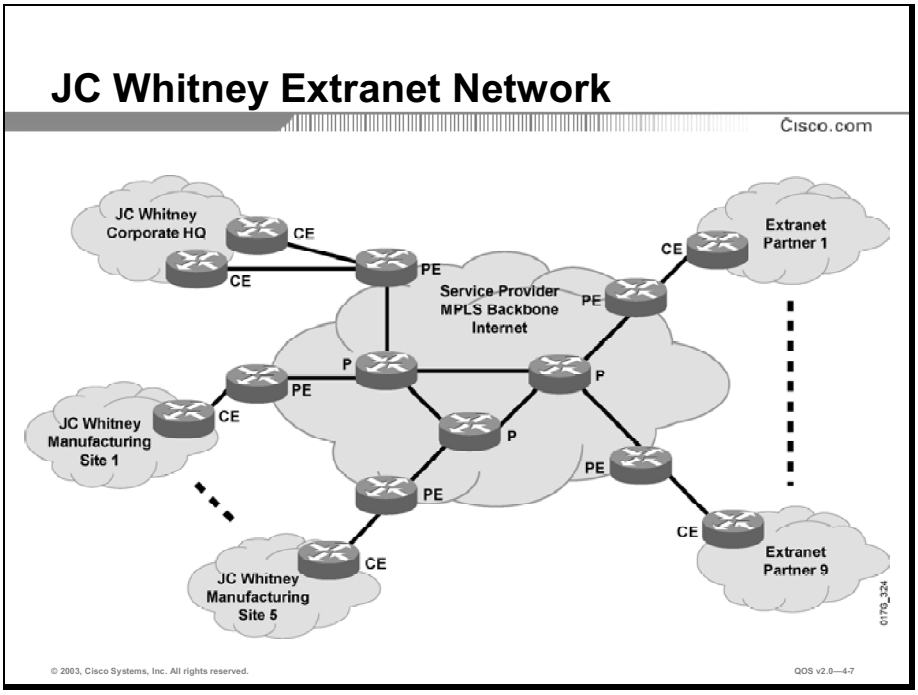

Customer Situation

JC Whitney has recently opened up Internet access to its regional manufacturing facilities and to its regional sales and distribution centers. As a result, access times to many of the company mission-critical applications such as sales and manufacturing databases have increased dramatically. In addition, response time between the corporate headquarters and JC Whitney extranet partners has increased, causing database queries to time out in some instances. No new applications have been added to the network other than enabling corporate-wide Internet access.

The JC Whitney network engineering staff explains their network applications in the following manner:

- JC Whitney has standardized on Open Shortest Path First (OSPF) as its routing protocol and therefore uses it on all of its routers company wide.
- The corporate headquarters and the five manufacturing facilities use VoIP for all intra-site and inter-site communications.
- The entire enterprise resource planning (ERP) database for the company is located at the corporate site. All sites (manufacturing, regional sales and distribution centers, extranet partners), use this centralized database for inventory control, sales data, invoicing, etc. Without complete access and reachability to the ERP database and its applications, JC Whitney could not manufacture product, ship inventory, or bill for its services.
- E-mail is another application that is used heavily at JC Whitney. The exchange servers and mail gateways are all located in the server farm at the corporate headquarters location.
- Internet services have recently been introduced company wide. One of its largest uses has been messaging between regional sales and distribution centers and between corporate staff and manufacturing. No internal messaging service currently exists at JC Whitney. As a result, the productivity gains realized by this Internet service have become somewhat important to the company. No other business applications currently exist on the Internet.
- Although the JC Whitney manufacturing facilities operate 24/7, the evening shifts have a reduced staff and line output. As a result, database synchronization and server backups are performed during the evening hours. A TCP-based backup application manages file transfers between manufacturing sites and the corporate headquarters using an automated version of FTP. Database synchronization is also TCP-based and has no critical bandwidth or latency requirements.

Working with the network engineering staff at JC Whitney and the service provider, you have been enlisted to assist JC Whitney by defining QoS requirements for their network. Their first priority is to determine what service classes to use and to identify where QoS classification and marking mechanisms should be configured in the network to enable JC Whitney administrative QoS policy, resolving the response time issues they are experiencing.

Identify QoS Service Class Requirements

Identify the different service classes required to implement the JC Whitney administrative QoS policy. Use the QoS Service Classes table to help you with your answer choices. Write your answers in the table below.

JC Whitney Service Classes

Customer Traffic	Service Class

QoS Service Classes

PHB	DSCP	DSCP Value	Intended Protocols and Applications	Service Class	Service Class and Configuration
EF	EF	101110	Interactive Voice	Voice Bearer	Admission Control = RSVP Queuing = Priority
AF1	AF11 AF12 AF13	001010 001100 001110	Intranet, General Data Service	Bulk Data	Queuing = Rate Based Active Queue Mgt = WRED minth AF13 < maxth AF13 <= minth AF12 < maxth AF12 <= minth AF11 < maxth AF11
AF2	AF21 AF22 AF23	010010 010100 010110	Database access, transaction services, interactive traffic, preferred data service	Transactional	Queuing = Rate Based Active Queue Mgt = WRED minth AF23 < maxth AF23 <= minth AF22 < maxth AF22 <= minth AF21 < maxth AF21
AF3	AF31 AF32 AF33	011010 011100 011110	Locally defined mission-critical applications	Mission-Critical	Queuing = Rate Based Active Queue Mgt = WRED minth AF33 < maxth AF33 <= minth AF32 < maxth AF32 <= minth AF31 < maxth AF31
AF4	AF41 AF42 AF43	100010 100100 100110	Interactive video and associated voice	Interactive Video	Admission Control = RSVP Queuing = Rate Based Active Queue Mgt = WRED minth AF43 < maxth AF43 <= minth AF42 < maxth AF42 <= minth AF41 < maxth AF41
CS6	Class 6	110000	BGP, OSPF, etc	Routing (Reserved)	Queuing = Rate Based Small guaranteed minimum rate Active Queue Mgt = RED minth < maxth, but minth is deep to minimize loss
CS4	Class 4	100000	Often proprietary	Streaming Video	Admission Control = RSVP Queuing = Rate Based Active Queue Mgt = RED minth < maxth

PHB	DSCP	DSCP Value	Intended Protocols and Applications	Service Class	Service Class and Configuration
CS3	Class 3	011000	SIP, H.323, etc.	Voice Signaling	Queuing = Rate Based Small guaranteed minimum rate Active Queue Mgt = RED minth < maxth, but minth is deep to minimize loss
CS1	Class 1	001000	User-selected service, Point-to-Point Applications	Less-than-Best Effort Data (Scavenger)	Queuing = Rate Based No bandwidth guarantee Active Queue Mgt = RED minth < maxth
Default	Default (Best-Effort) Class 0	000000	Unspecified traffic, Email	Best-Effort	Queuing = Rate Based Minimal bandwidth guarantee Active Queue Mgt or Per-flow fair queuing Active Queue Mgt = RED minth < maxth

In order to provide end-to-end QoS, multiple markers may be required. For each service class required for the JC Whitney network, complete the table below with the appropriate value of each specified marker.

JC Whitney QoS Service Class Requirements

Service Class	L 3 Classification			L 2 Classification	
	DSCP PHB	DSCP	IP Precedence	CoS	MPLS EXP

Identify Network Locations Where Classification and Marking Should be Applied

Using the information provided in the review of customer QoS requirements for this case study, use the diagrams of the JC Whitney network below to indicate trust boundaries, where classification and marking should be applied, markers in use, and locations where QoS markers change to ensure end-to-end QoS. Below is a sample network showing trust boundaries, where classification and marking should be applied, and markers in use. Use this sample to assist you in completing this activity. When completing this activity, indicate the following on each network diagram provided:

- Trust boundaries
- QoS markers in use
- Network locations where classification and marking should be used
- Locations where QoS markers change

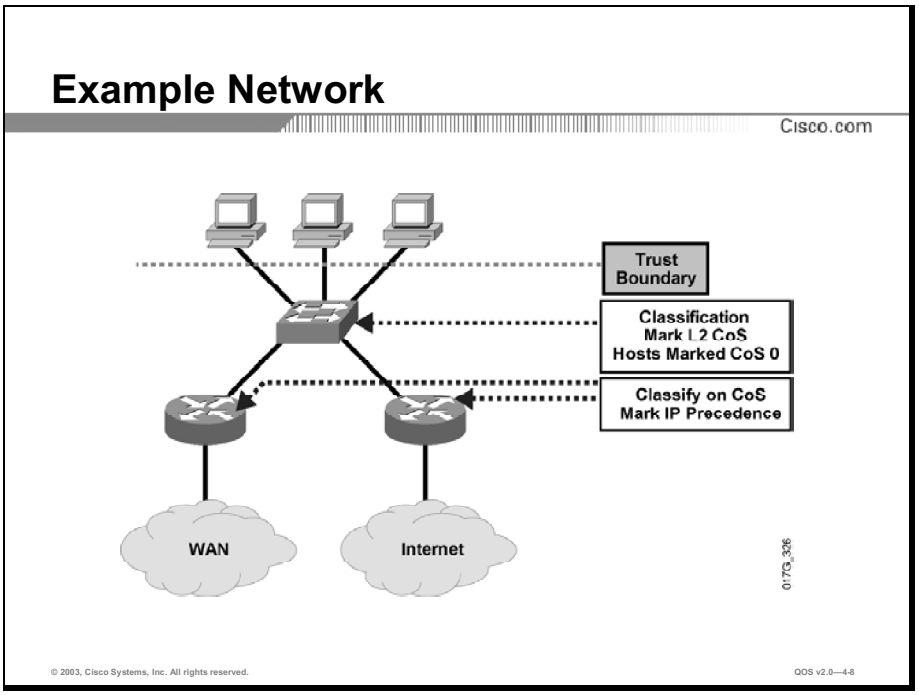

Sample network illustration of what should be marked for this section of the case study.

JC Whitney Corporate Headquarters Network

The JC Whitney network consists of a converged voice and data network. Because voice is a business-critical application, all voice traffic should be treated appropriately. The user community at JC Whitney ranges from novice data-entry clerks, to advanced systems programmers. As a result, security measures require that user workstations should *not* be allowed to set packet priorities.

Use the network diagram of the JC Whitney corporate network below to indicate the following:

- Trust boundaries
- QoS markers in use
- Network locations where classification and marking should be used
- Locations where QoS markers change

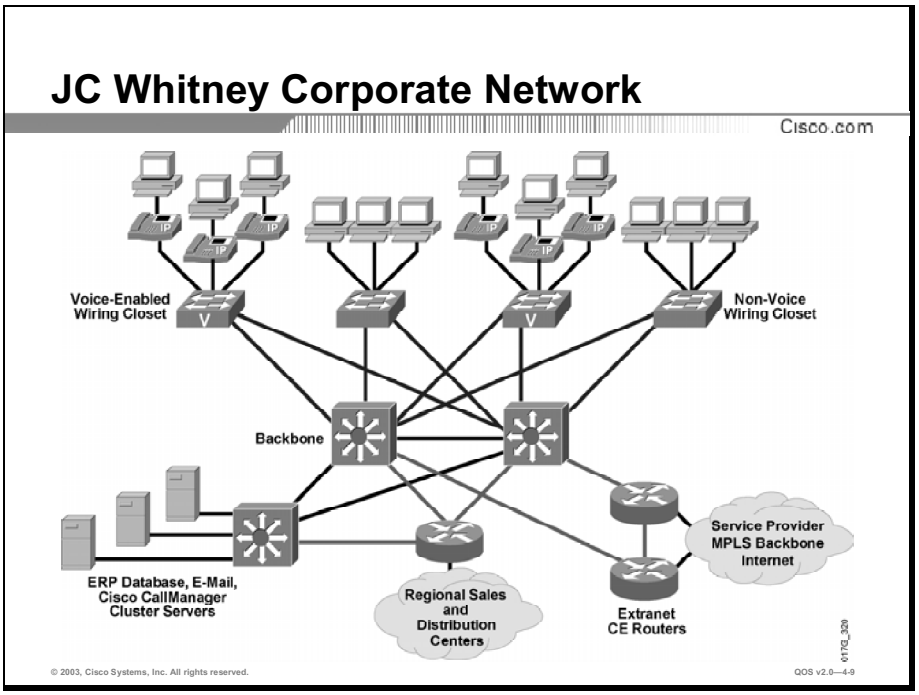

JC Whitney Regional Sales and Distribution Center Networks

The JC Whitney regional sales and distribution center networks are very low-overhead operations. Each sales office is staffed with three to nine employees. Distribution centers are similar to sales offices, but can be supported by up to 20 employees. The network at each center consists of a basic 10/100 Mbps desktop switch that is used to connect the office workstations to the corporate headquarters or a regional manufacturing facility via a Frame Relay connected low-end router.

Use the network diagram of the JC Whitney corporate network below to indicate the following:

- Trust boundaries
- QoS markers in use
- Network locations where classification and marking should be used
- Locations where QoS markers change

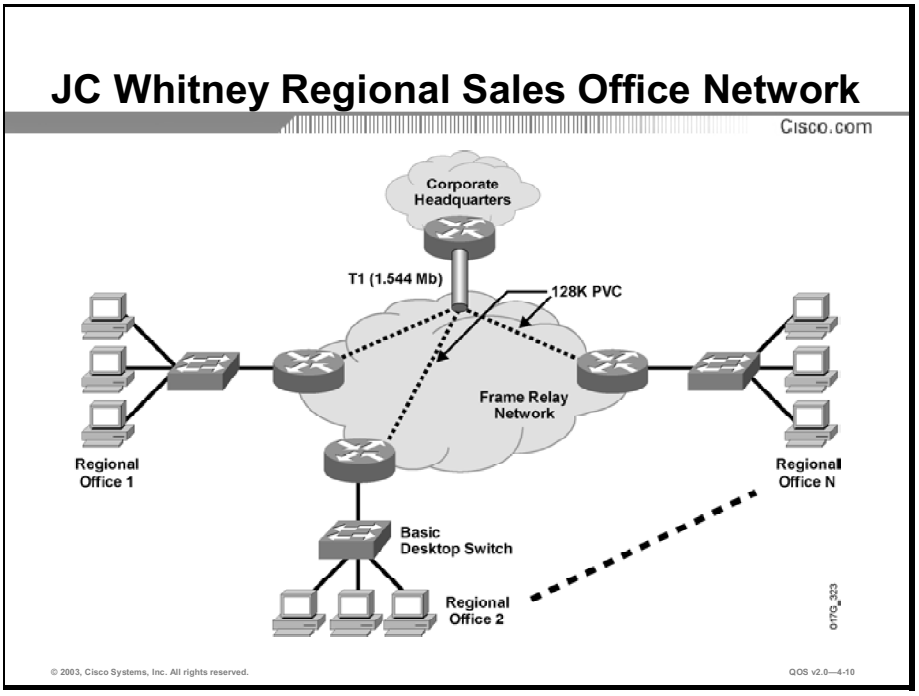

Present Your Solution

Together with your partner, present your solution to the class. Include the following information:

- Customer service class requirements
- Network diagrams indicating where classification and marking should be applied
- Justification for differences from the solution presented by the instructor

Case Study Answer Key

Identify QoS Service Class Requirements

JC Whitney Service Classes

Customer Traffic	Service Class
OSPF Routing Protocol	Reserved
Voice over IP	Voice Bearer
Voice Signaling (Skinny, SIP, etc.)	Voice Signaling
ERP (Transactional Database)	Transactional Data
E-mail	Best-Effort Data
Internet (Browsing, Messaging)	Bulk Data
Backup, Synch (FTP Bulk transfer)	Bulk Data

JC Whitney QoS Service Class Requirements

Service Class	L 3 Classification			L 2 Classification	
	DSCP PHB	DSCP	IP Precedence	CoS	MPLS EXP
Reserved	CS 6	48 (110 000)	6	6	6
Voice Bearer	EF	46 (101 110)	5	5	5
Voice Signaling	AF31	26 (011 010)	3	3	3
Transactional Data	AF21	18 (010 010)	2	2	2
Bulk Data	AF11	10 (001 010)	1	1	1
Best-Effort Data	Default	0 (000 000)	0	0	0

Identify Network Locations Where Classification and Marking Should be Applied

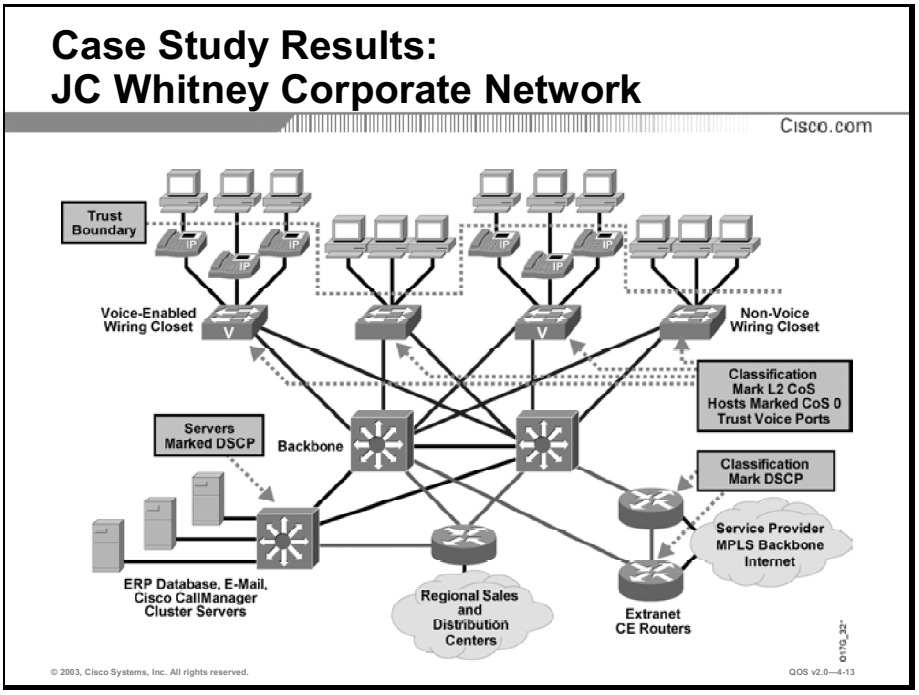

JC Whitney Corporate Headquarters Network

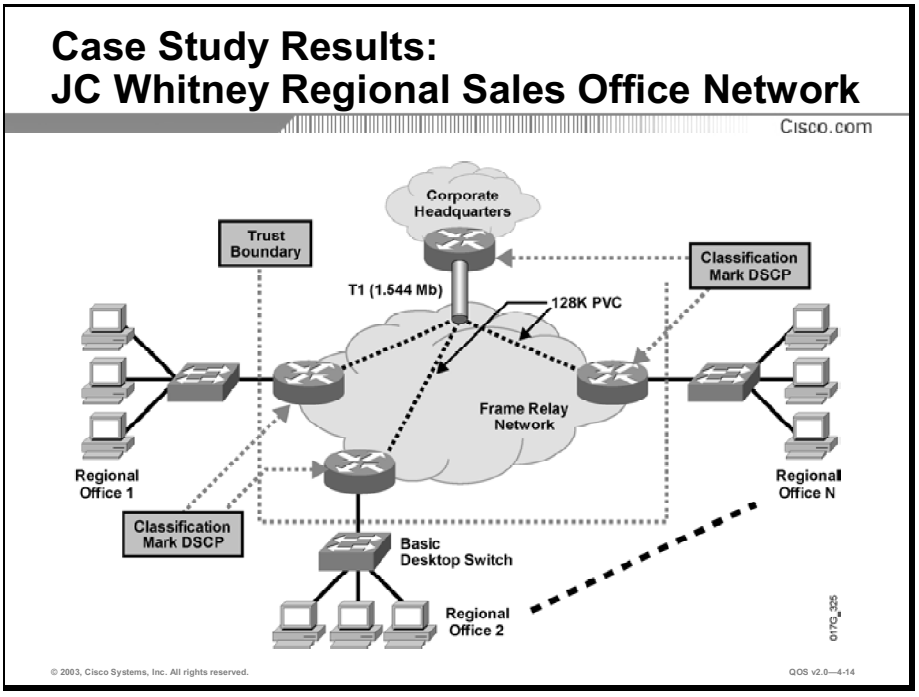

JC Whitney Regional Sales and Distribution Center Networks

Using MQC for Classification

Overview

The application of QoS requires that traffic be separated into service classes upon which differentiated levels of service are applied. Separation of traffic into different service classes requires QoS classification mechanisms. The MQC is one such mechanism for classifying network traffic.

This lesson describes the packet classification features of the MQC including input interface, access control lists (ACLs), CoS, IP precedence, and DSCP. This lesson also describes how MQC class maps can be configured to classify network traffic.

Relevance

Classification is a fundamental requirement for any network deployment of QoS. As such, it is of key importance to understand what traffic can be classified, how different classification mechanisms function, and how classification mechanisms are configured on Cisco IOS devices when implementing QoS.

Objectives

Upon completing this lesson, you will be able to use MQC CLI commands to classify packets. This includes being able to meet these objectives:

- Describe the different IP packet classification options in the MQC
- Identify the Cisco IOS commands used to configure classification of packets with MQC
- Identify the Cisco IOS commands required to classify IP packets using input interface with MQC
- Identify the Cisco IOS commands required to classify IP packets using CoS with MQC
- Identify the Cisco IOS commands required to classify IP packets using access lists with MQC
- Identify the Cisco IOS commands required to classify IP packets using IP precedence with MQC
- Identify the Cisco IOS commands required to classify IP packets using DSCP with MQC
- Identify the Cisco IOS commands required to classify IP packets using RTP (UDP port range) with MQC
- Identify the Cisco IOS commands used to monitor classification with MQC

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Basic knowledge of internetworking with TCP/IP concepts
- Basic knowledge of the Cisco IOS command-line interface

Outline

The outline lists the topics included in this lesson.

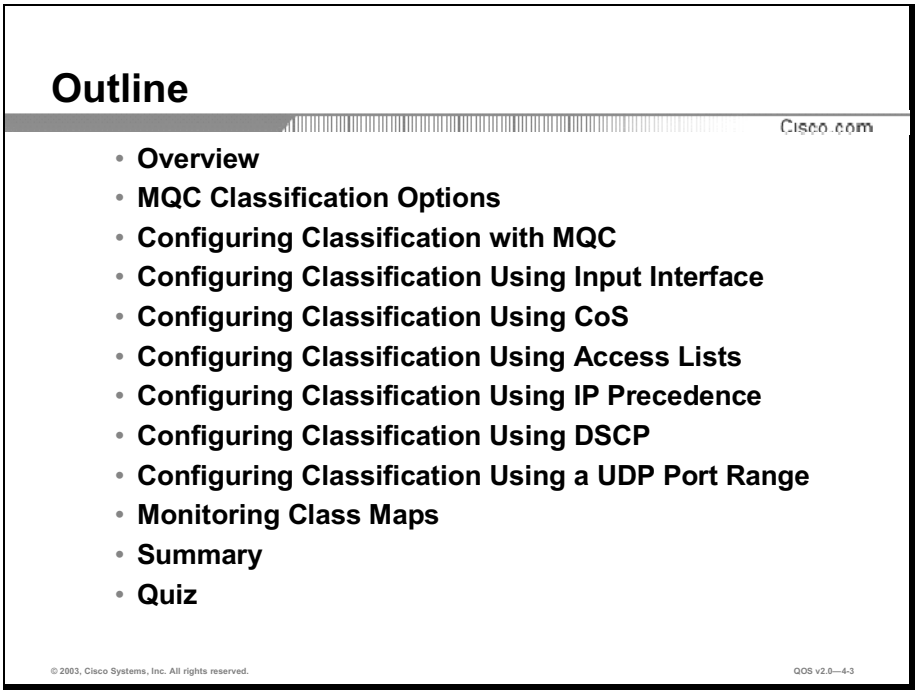

MQC Classification Options

This topic describes the different IP packet classification options available in MQC.

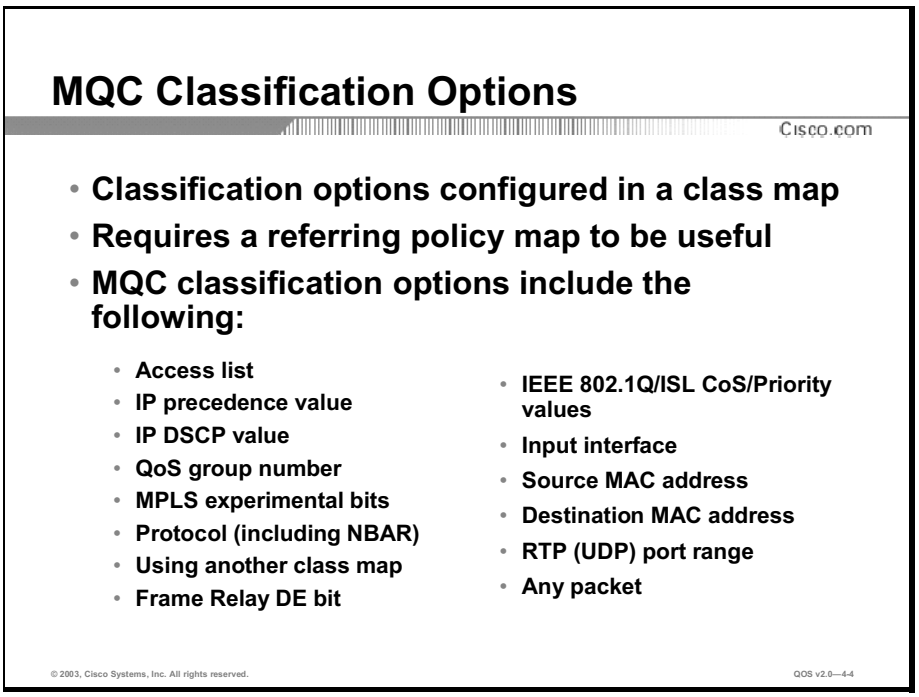

Classification using MQC is accomplished by specifying a traffic match criteria within a configured class map for each different service class. In order for QoS mechanisms to utilize the class map, it must be referenced through the use of a policy map, which is subsequently applied to an inbound our outbound interface as a service policy.

In older Cisco IOS software releases, the router classified a packet against every individual QoS feature. This resulted in additional processing overhead, inaccurate packet counters, and double accounting issues. Common classification is a feature that was introduced in Cisco IOS 12.2, and is enabled by default whenever classification is invoked within a policy map. With common classification, a packet is classified only once per service policy and matches a single class in the policy. Because matching terminates at the first matching class it is important to ensure that the classes are configured in the right sequence within a policy. After a packet is classified against a particular class, it is subjected to all the QoS features configured within that class.

MQC classification with class maps is extremely flexible and can classify packets by using the following classification tools:

- **Access list:** Access lists for any protocol can be used within the class map configuration mode. The MQC can be used for other protocols, not only IP.
- **IP precedence:** IP packets can be classified directly by specifying IP precedence values.
- **DSCP:** IP packets can be classified directly by specifying IP DSCP values. DiffServ enabled networks can have up to 64 classes if DSCP is used to mark packets.

- **QoS group:** A QoS group parameter can be used to classify packets in situations where up to 100 classes are needed or the QoS group parameter is used as an intermediary marker; for example, MPLS to QoS group translation on input and QoS group to DSCP translation on output. QoS group markings are local to a single router.
- **MPLS experimental bits:** Packets can be matched based on the value in the experimental bits of the MPLS header of labeled packets.
- **Protocol:** Classification is possible by identifying Layer 3 or Layer 4 protocols. Advanced classification is also available by using the NBAR tool where dynamic protocols are identified by inspecting higher-layer information.
- **Class map hierarchy:** Another class map can be used to implement template-based configurations.
- **Frame Relay DE bit:** Packets can be matched based on the value of the underlying Frame Relay DE bit.
- **CoS:** Packets can be matched based on the information contained in the three CoS bits (when using IEEE 802.1Q encapsulation) or priority bits (when using the ISL encapsulation).
- **Input interface:** Packets can be classified based on the interface from which they enter the Cisco IOS device.
- **MAC address:** Packets can be matched based on their source or destination MAC addresses.
- **User Datagram Protocol (UDP) port range:** Real-Time Transport Protocol (RTP) packets can be matched based on a range of UDP port numbers.
- **All packets:** MQC can also be used to implement a QoS mechanism for all traffic in which case classification will put all packets into one class.

Configuring Classification with MQC

This topic identifies the Cisco IOS commands used to configure classification of packets with MQC.

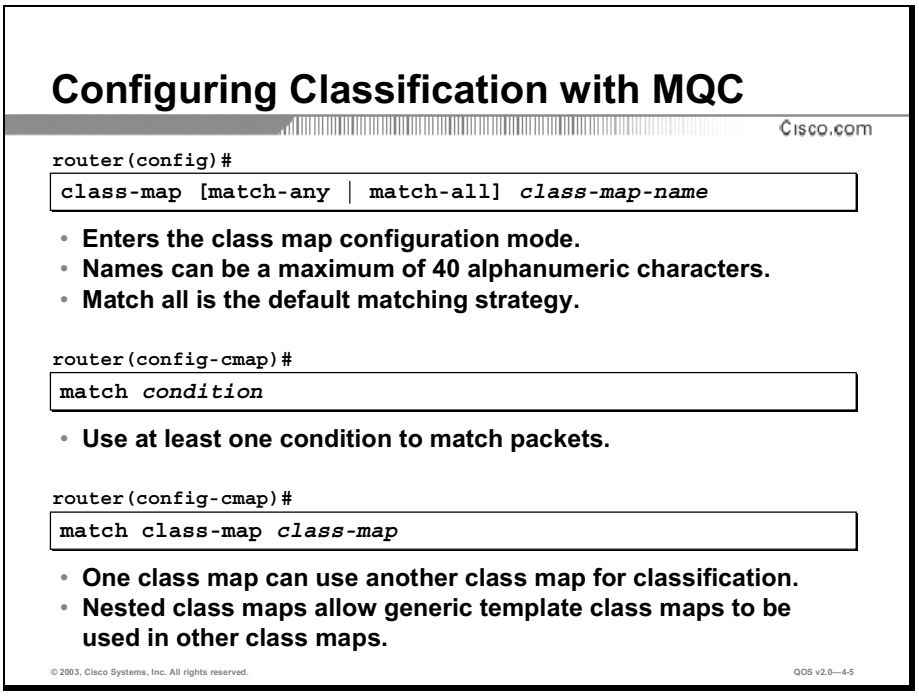

The **class-map** global configuration command is used to create a class map and enter the class map configuration mode. A class map is identified by a case-sensitive name; therefore, all subsequent references to the class map must use exactly the same name.

The **match** command is used to specify the classification criteria when in class map configuration mode. Multiple match commands can be used within a class map. At least one match command should be used within the class map configuration mode. (Match none is the default.)

It is also possible to nest class maps in MQC configurations. Nesting class maps is accomplished using the **match class-map** command within the class map configuration. By nesting class maps, the creation of generic classification templates and more sophisticated classification are possible.

class-map [match-any | match-all] *class-map-name*

Syntax Description

Parameter	Description
class-map-name	Name of the class for the class map. The class name is used for both the class map and to configure policy for the class in the policy map.
match-all | **match-any**	Name of the class for the class map. The name can be a maximum of 40 alphanumeric characters. The class name is used for both the class map and to configure policy for the class in the policy map.

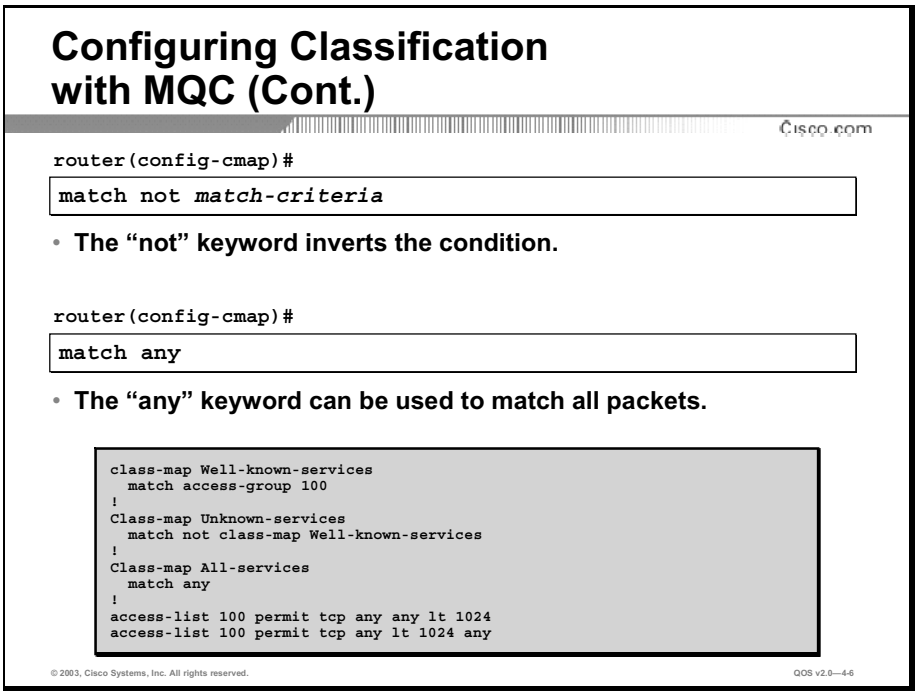

There are some additional options that give extra power to class maps:

- Any condition can be negated by inserting the keyword "**not**".
- A class map can use another class map to match packets.
- The **any** keyword can be used to match all packets.

The example shows three class maps:

- Class map **Well-known-services** uses an access list to match all the packets with the source or destination port number lower than 1024.
- Class map **Unknown-services** uses the first class map but negates the result. The same could be achieved by using the same access list with a negation.
- Class map **All-services** actually matches all the packets.
- **match not** match-criteria.

Syntax Description

Parameter	Description
match-criteria	(Required) Specifies the match criterion value that is an unsuccessful match criterion. All other values of the specified match criterion will be considered successful match criteria.

Configuring Classification Using Input Interface

This topic identifies the Cisco IOS commands that are required to classify IP packets using input interface with MQC.

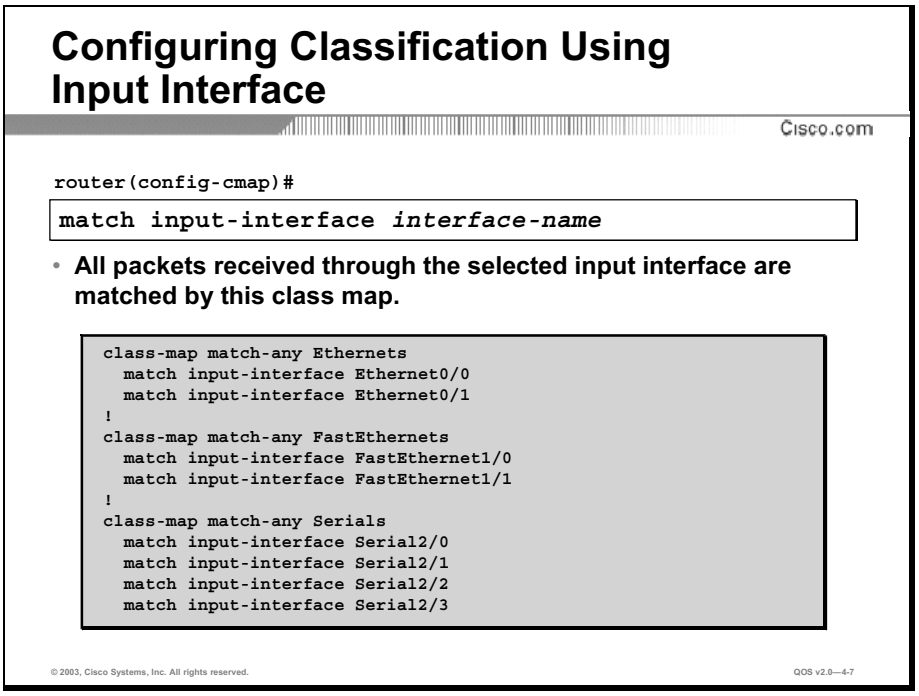

As shown in the example, a packet can also be classified based on the input interface. In the first class map example, called Ethernets, the **match input-interface** will match any packet that arrives on either the E0/0 or E0/1 interfaces.

In the second class map, FastEthernets, any packet arriving on either the FastEthernet 1/0 or FastEthernet 1/1 interface will be matched.

And in the last class map example, Serials, incoming packets arriving on any of S2/0, S2/1, S2/2 or S2/3 will be matched.

match input-interface *interface-name*

Syntax Description

Parameter	Description
interface-name	Name of the input interface to be used as match criteria.

Configuring Classification Using CoS

This topic identifies the Cisco IOS commands that are required to classify IP packets using CoS with MQC.

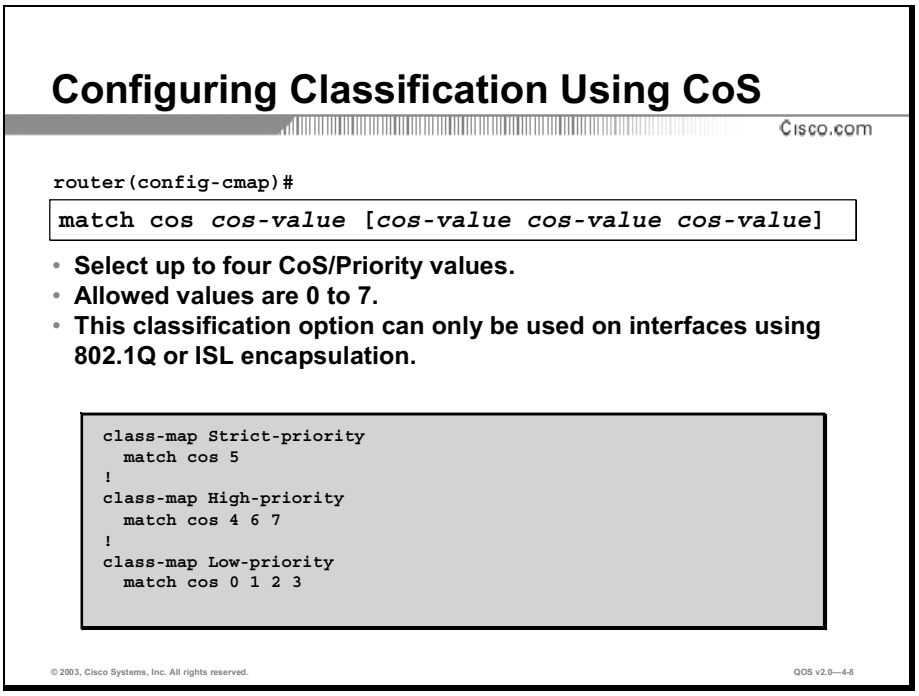

Routers can also match on the three CoS bits in 802.1Q headers or priority bits in the ISL header. These bits can be used in a LAN-switched environment to provide differentiated quality of service.

This is demonstrated in the example. In the first class map, Strict-priority, packets will be matched if they have a CoS value of 5.

In the second class map example High-priority, packets will be matched if they have a CoS value of either 4, 6, or 7.

And in the last class map example Low-priority, packets will be matched if they have a CoS value of any of 0, 1, 2, or 3.

match cos *cos-value* **[***cos-value cos-value cos-value***]**

Syntax Description

Parameter	Description
cos-value	(Optional) Specific IEEE 802.1Q/ISL CoS value. The *cos-value* is from 0 to 7; up to four CoS values can be specified in one **match cos** statement.

Configuring Classification Using Access Lists

This topic identifies the Cisco IOS commands that are required to classify IP packets using access lists with MQC.

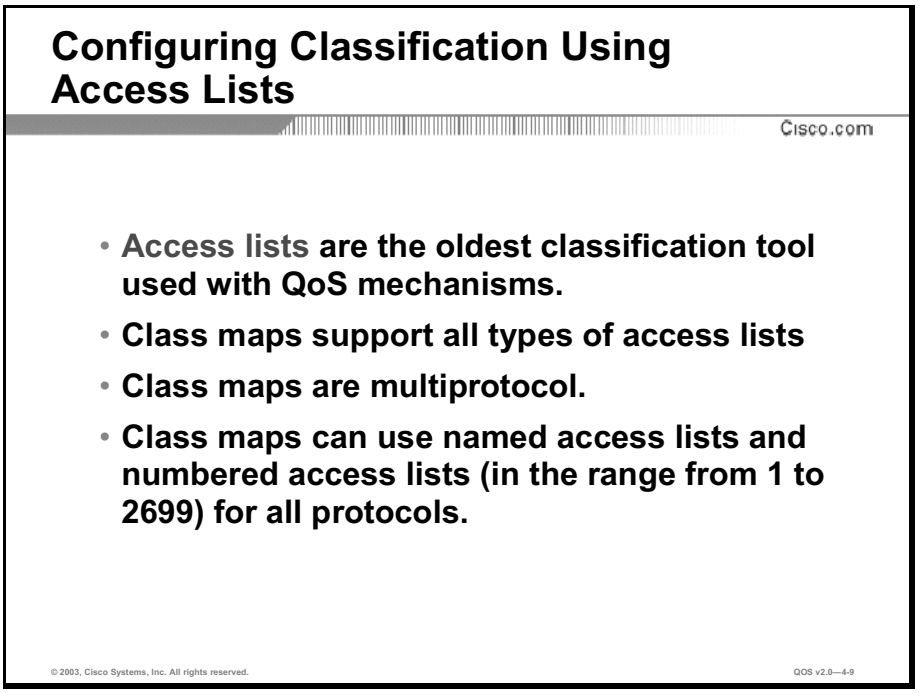

Access lists were originally used for filtering of inbound or outbound packets on interfaces. They were later reused for filtering of routing updates and also for classification with early QoS tools; for example, priority queuing (PQ) custom queuing, and traffic shaping.

Access lists are still one of the most powerful classification tools. Class maps can use any type of access list (not only IP access lists).

Access lists also have a drawback. Compared to other classification tools they are one of the most CPU-intensive. For this reason access lists should not be used for classification on high-speed links where they could severely impact performance of routers. Access lists are typically used on low-speed links at network edges where packets are classified and marked (for example, with IP precedence). Classification in the core is done based on the IP precedence value.

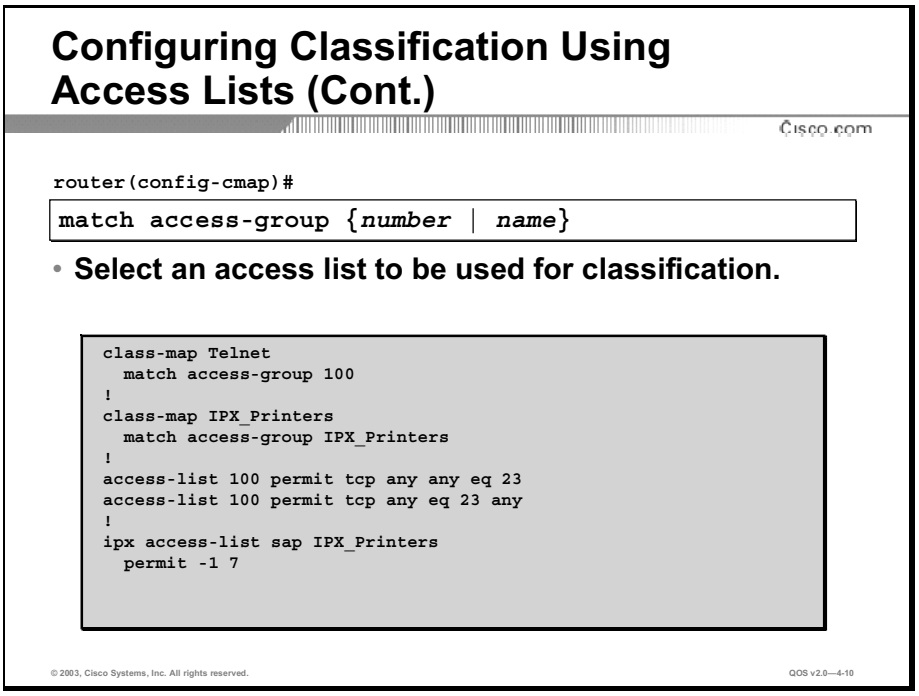

Use the **match access-group** command to attach an access list to a class map.

The example in the figure shows how a numbered or named access list can be used for classification.

In the first example, class map Telnet, packets will matched according to those allowed by the access-group 100. When exploring what will be allowed, access-list 100 will permit port 23.

In the second example, packets will be allowed if they matched according to those allowed by the access-list IPX_Printers.

match access-group {*number* | *name*}

Syntax Description

Parameter	Description
access-group	A numbered ACL whose contents are used as the match criteria against which packets are checked to determine if they belong to this class.
name access-group-name	A named ACL whose contents are used as the match criteria against which packets are checked to determine if they belong to this class.

Configuring Classification Using IP Precedence

This topic identifies the Cisco IOS commands required to classify IP packets using IP precedence with MQC.

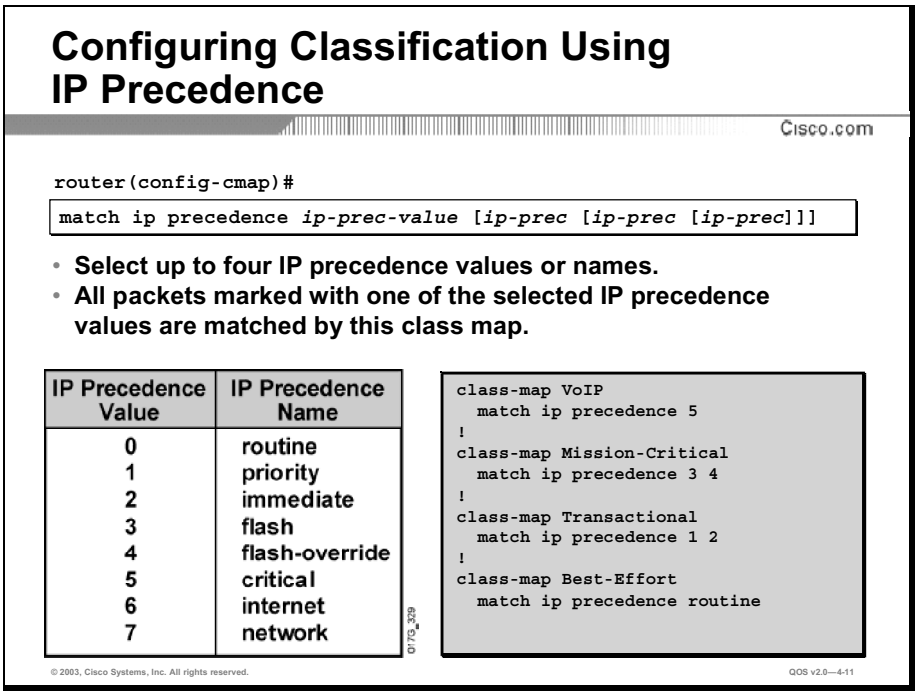

A much faster method of classification is by matching the IP precedence. Up to four IP precedence values or names can be used to classify packets based on the IP precedence field in the IP header.

The figure contains a mapping between IP precedence values and names. The running configuration, however, only shows IP precedence values (not names).

match ip precedence *ip-prec-value* **[***ip-prec* **[***ip-prec* **[***ip-prec***]]]**

Syntax Description

Parameter	Description
ip-precedence-value	Specifies the exact value from 0 to 7 used to identify an IP precedence value.

Configuring Classification Using DSCP

This topic identifies the Cisco IOS commands that are required to classify IP packets using DSCP with MQC.

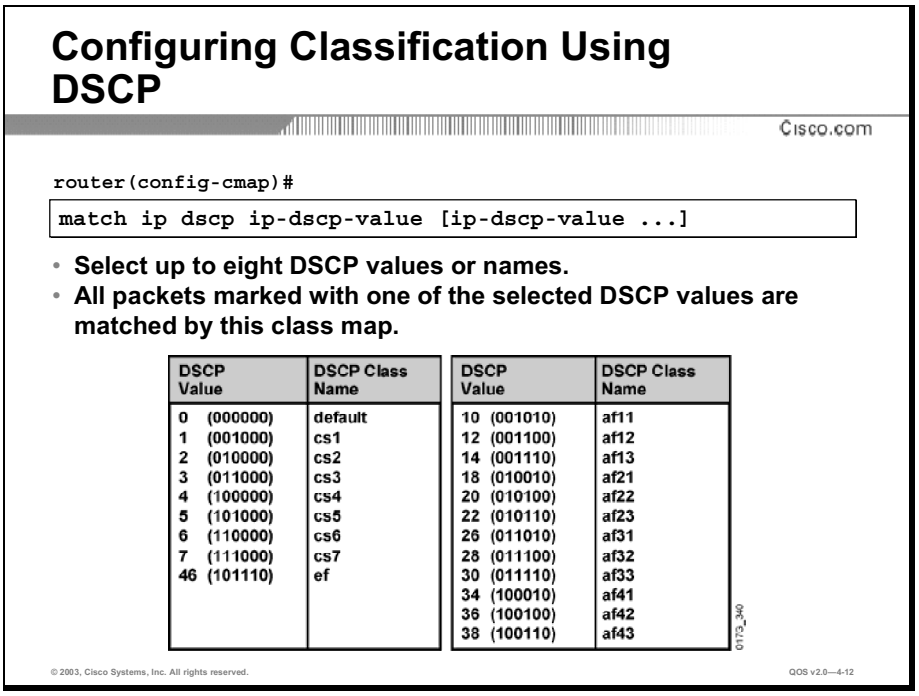

IP packets can also be classified based on the IP DSCP field. A QoS design can be based on IP precedence marking or DSCP marking. DSCP marking can include backward compatibility with IP precedence by using the Class Selector (CS) values (most significant three bits of the DSCP value).

match ip dscp *ip-dscp-value* **[***ip-dscp-value* **...]**

Syntax Description

Parameter	Description
ip	(Optional) Specifies that the match is for IPv4 packets only. If not used, the match is on both IPv4 and IPv6 packets.
dscp-value	Specifies the exact value from 0 to 63 used to identify an IP DSCP value.

A sample design that includes backward compatibility would use the following values to mark packets belonging to class Gold, which is guaranteed Assured Forwarding (AF) per-hop behavior (PHB):

- af11 marks low-drop packets
- af12 marks medium-drop packets
- af13 marks high-drop packets
- cs4 marks low-drop packets (for backward compatibility with IP precedence 4)
- cs3 marks high-drop packets (for backward compatibility with IP precedence 3)

Configuring Classification Using DSCP (Cont.)

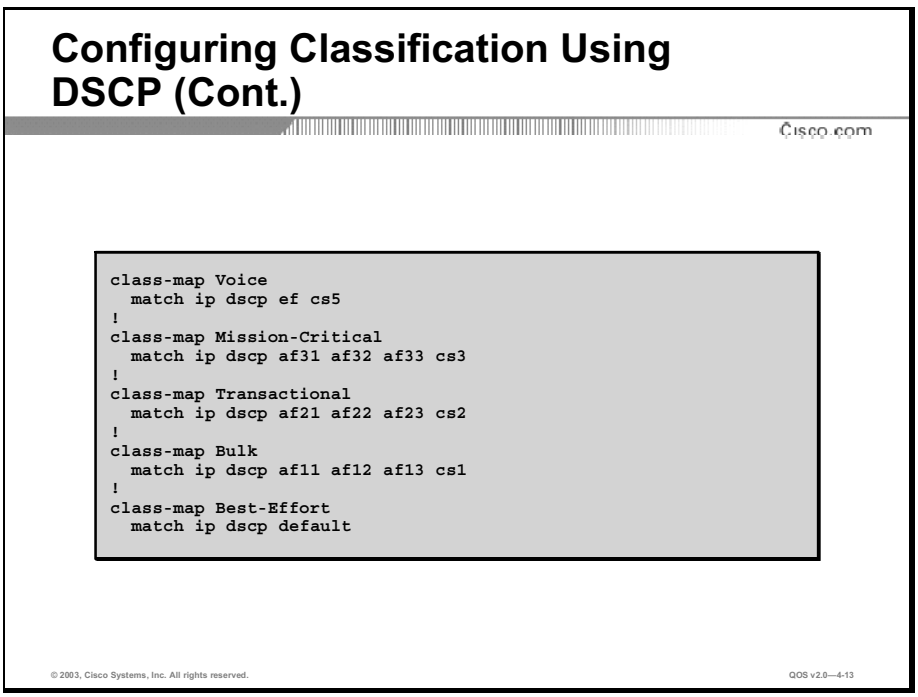

The figure illustrates implementation of a design with five classes:

- **Voice:** Identified by DSCP value ef, which looks like IP precedence value 5 in non-DSCP compliant devices.
- **Mission-critical:** Identified by DSCP values af31, af32 and af33. The class is also identified by IP precedence 3.
- **Transactional:** Identified by DSCP values af21, af22 and af23. The class is also identified by IP precedence 2.
- **Bulk:** Identified by DSCP values af11, af12 and af13. The class is also identified by IP precedence 1.
- **Best-effort:** Identified by the default DSCP value that is equal to the default IP precedence value (0).

From a non-DSCP compliant device the design looks slightly different:

- **Voice:** IP precedence 5
- **Mission-critical:** IP precedence 3
- **Transactional:** IP precedence 2
- **Bulk:** IP precedence 1
- **Best-effort:** IP precedence 0

A DSCP-compliant device treats packets marked by a non-DSCP-compliant device according to the design. A non-DSCP-compliant device does not treat packets marked by a DSCP-compliant device correctly due to values of drop ratings:

- AF1 (001xx0) looks like IP precedence 1. Therefore, class bulk incorrectly appears as class mission-critical in a non-DSCP-compliant device.

- AF2 (010xx0) looks like IP precedence 2. Therefore, class transactional correctly appears as class transactional in a non-DSCP-compliant device.
- AF3 (011xx0) looks like IP precedence 3. Therefore, class mission-critical appears as class bulk in a non-DSCP compliant device.
- EF (101110) looks like IP precedence 5, which is also used for voice in a non-DSCP compliant device.

As seen from the example it is very important to understand the impact of DSCP on non-DSCP-compliant devices. A DiffServ-based QoS design should include the impact of DSCP on parts of the networks where all routers are not DSCP-compliant.

The example shows that a network core, if upgraded to support DSCP, can correctly handle packets classified by edge devices that have not yet been upgraded.

Configuring Classification Using a UDP Port Range

This topic identifies the Cisco IOS commands that are required to classify IP packets using RTP (UDP port range) with MQC.

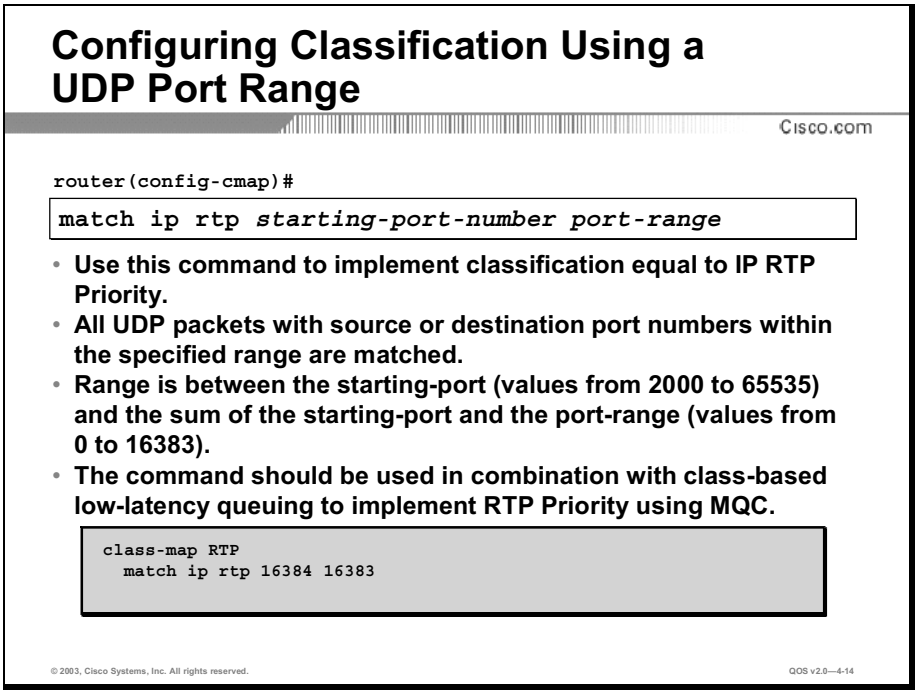

IP RTP Priority was introduced to provide low-latency queuing (LLQ) in combination with weighted fair queuing (WFQ). The **match ip rtp** command can be used to match packets in the same way as with IP RTP Priority. It should also be combined with LLQ to generate a similar result as IP RTP Priority.

match ip rtp *starting-port-number port-range*

Syntax Description

Parameter	Description
starting-port-number	The starting RTP port number. Values range from 2000 to 65535.
port-range	The RTP port number range. Values range from 0 to 16383.

Monitoring Class Maps

This topic identifies the Cisco IOS commands that are used to monitor classification with MQC.

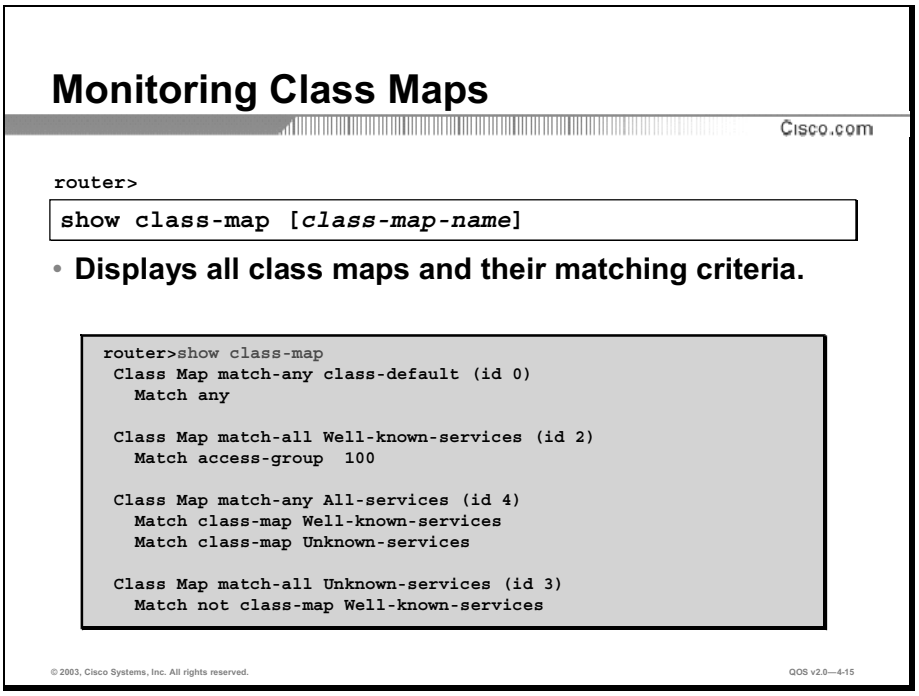

The **show class-map** command lists all class maps with their match statements. This command can be issued from the EXEC or Privileged EXEC mode.

The **show class-map** command with a name of a class map displays the configuration of the selected class map.

In the figure, the **show class map** Cisco IOS command shows all the class maps that have been configured and what match statements are contained in the maps.

The first class map listed is the default class map. The default class map contains only a single match statement; match-any.

The second class map listed, Well-known-services, has one match statement that will compare packets against the configured access-group 100.

The third class map displayed, All-services, contains two match statements that compare packets against two other configured class maps, well-known-services and unknown-services.

show class-map [*class-map-name***]**

Syntax Description

Parameter	Description
`class-map-name`	(Optional) Name of the class map.

Summary

This topic summarizes the key points discussed in this lesson.

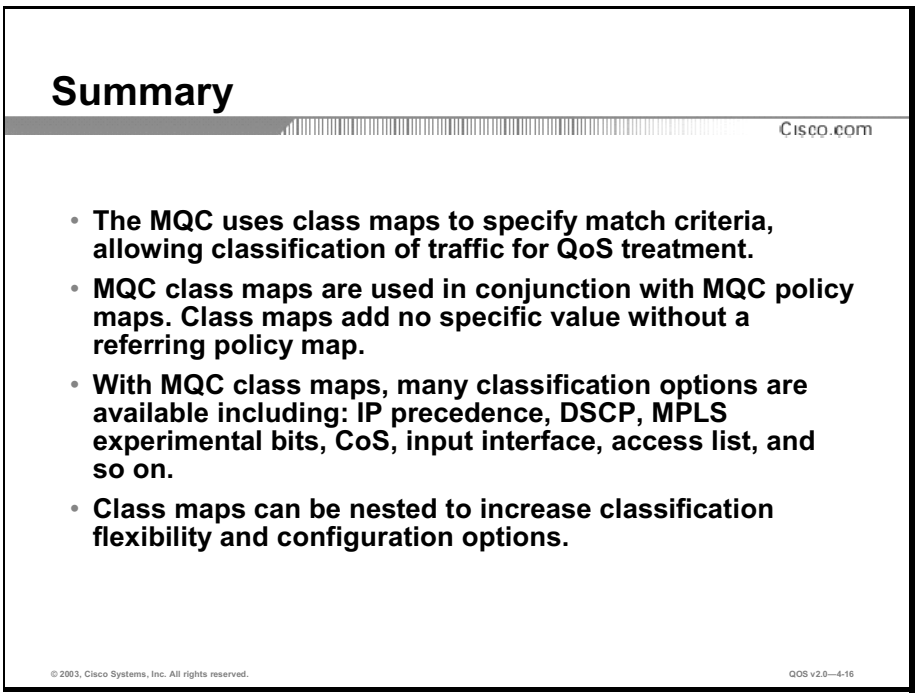

References

For additional information, refer to these resources:

- For more information on classification using MQC, refer to “Cisco Modular Quality of Service Command Line Interface” at the following URL: http://www.cisco.com/warp/public/cc/pd/iosw/prodlit/moqcs_wp.pdf
- For more information on classification using MQC, refer to Configuring the Modular Quality of Service Command-Line Interface” at the following URL: http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122cgcr/fqos_c/fqcprt8/qcfmcli2.pdf

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) MQC classification is obtained by specifying a traffic match criteria within a configured _____.

A) class map

B) policy map

C) route map

D) service policy

Q2) What command within a class map is used to specify the classification criteria?

A) **all**

B) **match**

C) **set**

D) none of the above

Q3) Which of the following commands is correct and will classify packets based on an interface?

A) **match *interface name* interface**

B) **match interface *interface name***

C) **match *interface name* input-interface**

D) **match input-interface *interface name***

Q4) What is the maximum number of CoS values that can be specified when using the match cos cos-value?

A) 1

B) 2

C) 4

D) 8

Q5) Class maps can use which of the following?

A) IP

B) IPX

C) named only

D) any access list

Q6) When classifying packets using IP precedence which of the following is correct?

A) only 1 value can be specified

B) up to 2 values can be specified

C) up to 4 values can be specified

D) only the IP precedence name can be used, not the numerical value

Q7) When classifying packets using DSCP, which of the following is correct?

A) up to 8 DSCP values can be specified at the same time

B) up to 8 DSCP class names can be specified at the same time

C) only DSCP class names, class cs1–cs7 and ef, can be matched

D) packets cannot be matched based on DSCP if packets arrive on a serial interface

Q8) The **match ip rtp** command is equal to which of the following?

A) ip rtp priority

B) ip tcp priority

C) ip idp priority

D) none of the above

Q9) Which of the Cisco IOS commands is correct to display all configured class maps?

A) Router>show class map

B) Router>show class-map

C) router(config)# show class map

D) router(config)# show class-map

Quiz Answer Key

Q1) A

Relates to: MQC Classification Options

Q2) B

Relates to: Configuring Classification with MQC

Q3) D

Relates to: Configuring Classification Using Input Interface

Q4) C

Relates to: Configuring Classification Using CoS

Q5) D

Relates to: Configuring Classification Using Access Lists

Q6) C

Relates to: Configuring Classification Using IP Precedence

Q7) A

Relates to: Configuring Classification Using DSCP

Q8) A

Relates to: Configuring Classification Using a UDP Port Range

Q9) B

Relates to: Monitoring Class Maps

Using MQC for Class-Based Marking

Overview

The process of packet classification can be both complex and CPU-intensive. Therefore, it is desirable to classify packets as close to the source as possible—at the edges of the network. Performing classification in the core is undesirable because it would necessarily add delay in transiting the core. To provide differential levels of treatment to service classes, traffic must be identified as “belonging” to a specific class. Instead of classifying traffic at each hop in the network as the packet traverses the network to its ultimate destination, QoS marking mechanisms are used. Marking allows specific fields in a frame or packet to be set that identifies that frame or packet as belonging to a specific service class. The MQC provides one such mechanism for marking network traffic.

This lesson describes the class-based marking capability of the Cisco IOS MQC and how policy maps can be configured to mark network traffic. MQC marking features covered in this lesson include CoS, IP precedence, and DSCP.

Relevance

Marking is a fundamental requirement for any network deployment of QoS. As such, it is of key importance to understand what markers can be set after traffic has been classified, how different marking mechanisms function, and how marking mechanisms are configured on Cisco IOS devices when implementing QoS.

Objectives

Upon completing this lesson, you will be able to use class-based marking to assign packets to a specific service class. This includes being able to meet these objectives:

- Describe class-based marking
- Describe the different IP packet marking options available in class-based marking
- Identify the Cisco IOS commands required to configure class-based marking
- Identify the Cisco IOS commands required to mark IP packets using CoS with class-based marking
- Identify the Cisco IOS commands required to mark IP packets using IP precedence with class-based marking
- Identify the Cisco IOS commands required to mark IP packets using DSCP with class-based marking
- Identify the Cisco IOS commands used to monitor class-based marking

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Basic knowledge of internetworking with TCP/IP concepts
- Basic knowledge of the Cisco IOS command-line interface

Outline

The outline lists the topics included in this lesson.

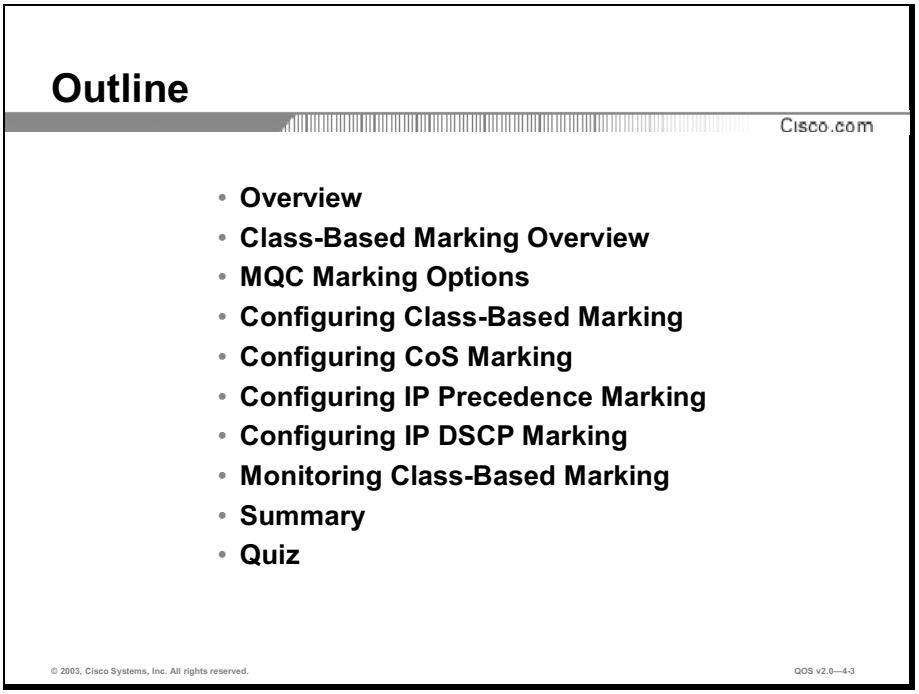

Class-Based Marking Overview

This topic describes the MQC class-based marking mechanism.

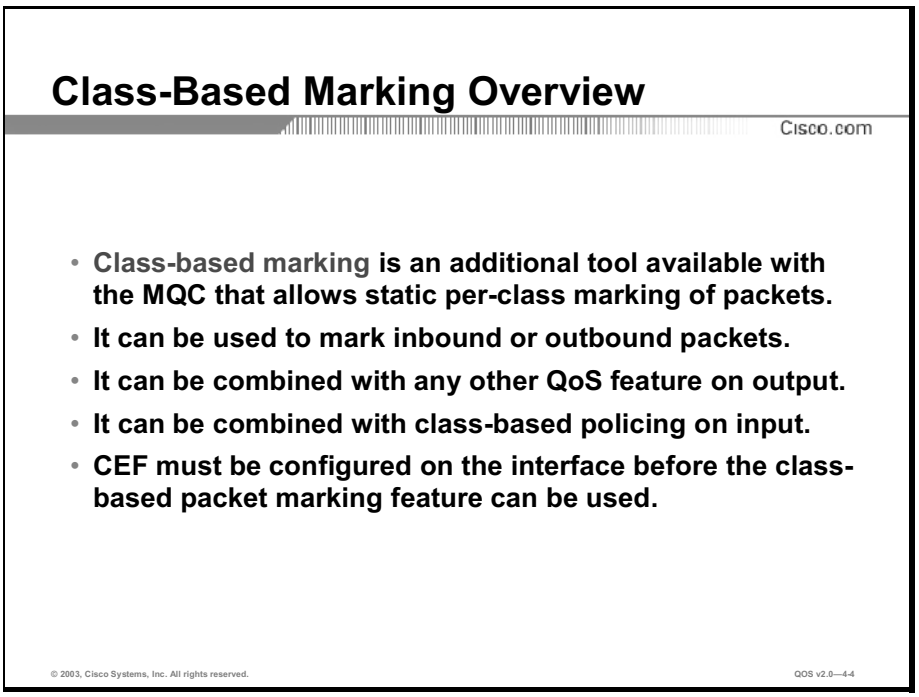

Marking packets or frames lets you set information in the Layer 2 and Layer 3 headers of a packet, so the packet or frame can be identified and distinguished from other packets or frames.

The class-based weighted fair queuing (CBWFQ) provides packet-marking capabilities using class-based marking, which is configured within the Cisco IOS MQC feature. It is the most flexible Cisco IOS marking tool, extending the marking functionality of committed access rate (CAR) and policy routing.

Class-based marking can be used on input or output of interfaces as part of a defined input or an output service policy. On input, class-based marking can be combined with class-based policing, and on output, with any other CBWFQ QoS feature.

MQC Marking Options

This topic describes the different IP packet marking options that are available in class-based marking.

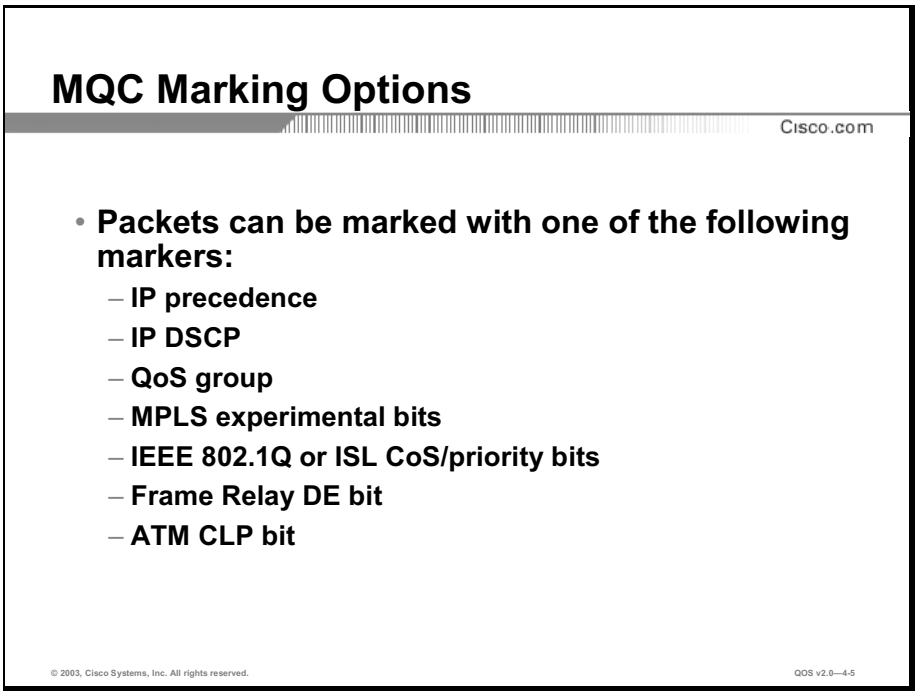

Class-based marking supports the following markers:

- IP precedence
- IP DSCP value
- QoS group
- MPLS experimental bits
- IEEE 802.1Q or ISL CoS/priority bits
- Frame Relay DE bit
- ATM CLP bit

Class-based marking can be combined with other mechanisms available in the MQC.

Configuring Class-Based Marking

This topic identifies the Cisco IOS commands that are required to configure class-based marking.

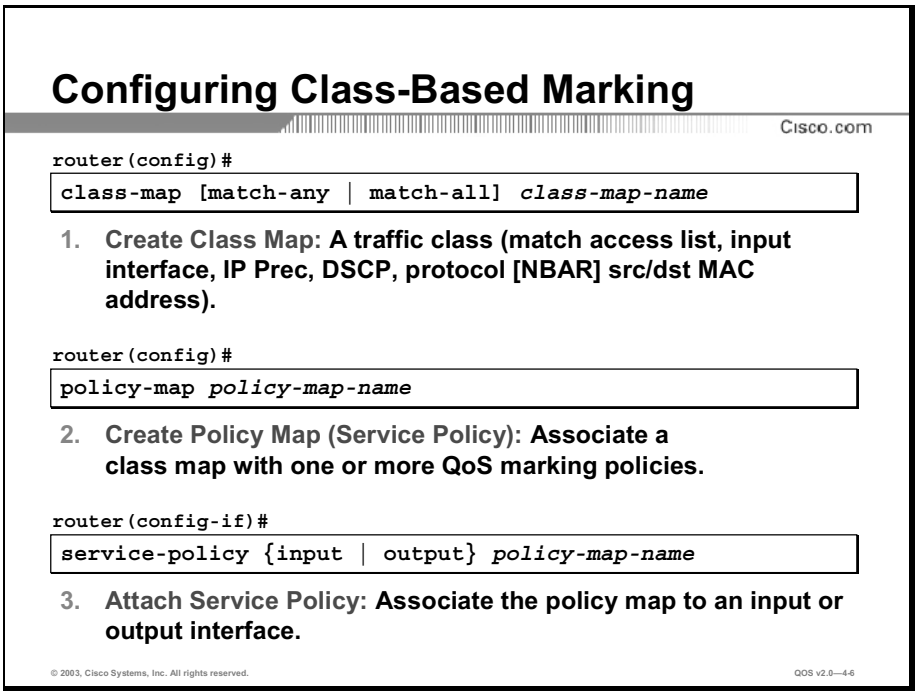

When configuring class-based marking, three configuration steps need to be completed:

- Create a class-map
- Create a policy-map
- Attach the policy-map to an interface by using the **service-policy** Cisco IOS command.

class-map *class-map-name*

Syntax Description

Parameter	Description
class-map-name	Name of the class for the class map. The class name is used for both the class map and to configure policy for the class in the policy map.
match-all | **match-any**	Determines how packets are evaluated when multiple match criteria exist. Packets must either meet all of the match criteria (**match-all**) or one of the match criteria (**match-any**) in order to be considered a member of the class.

policy-map *policy-map-name*

Syntax Description

Parameter	Description
policy-map-name	Name of the policy map.

service-policy policy-map-name

Syntax Description

Parameter	Description
policy-map-name	Specifies the name of the predefined policy map to be used as a QoS policy. The name can be a maximum of 40 alphanumeric characters.

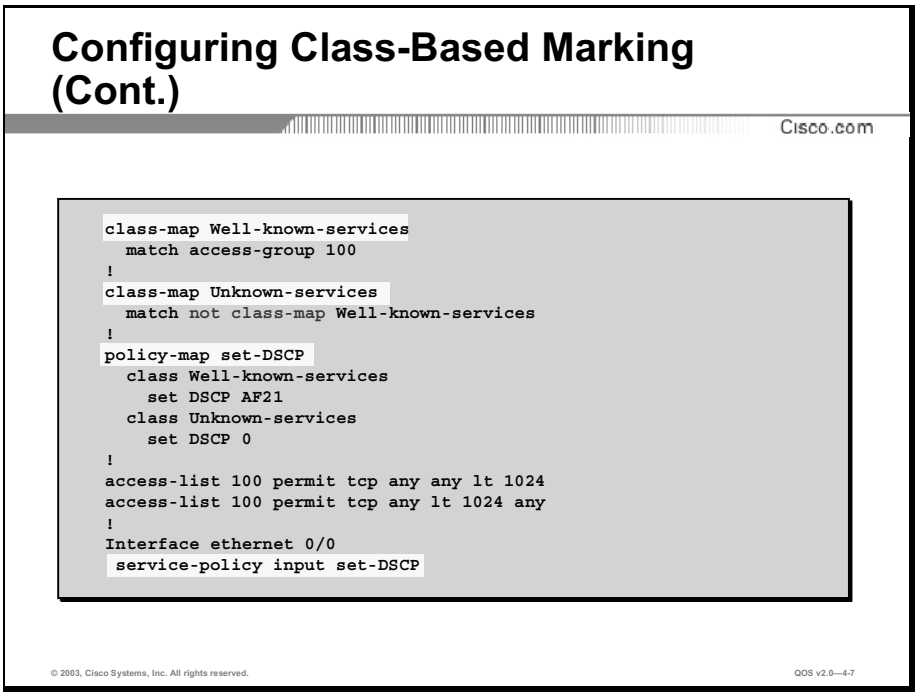

In the figure, two class-maps have been configured, *Well-known-services* and *Unknown-services.* The match criterion is specified in access-list 100.

The policy-map *set-DSCP* has been created to associate the class-maps *Well-known-services* and *Unknown-services* with it.

For packets allowed by class-map *Well-known-services* the IP DSCP value will be set AF21. Those matching class-map *Unknown-services* will have the IP DSCP value set to 0.

The policy-map is attached to E0/0 for incoming packets by the **service-policy** command.

Configuring CoS Marking

This topic identifies the Cisco IOS commands that are required to mark IP packets using CoS with class-based marking.

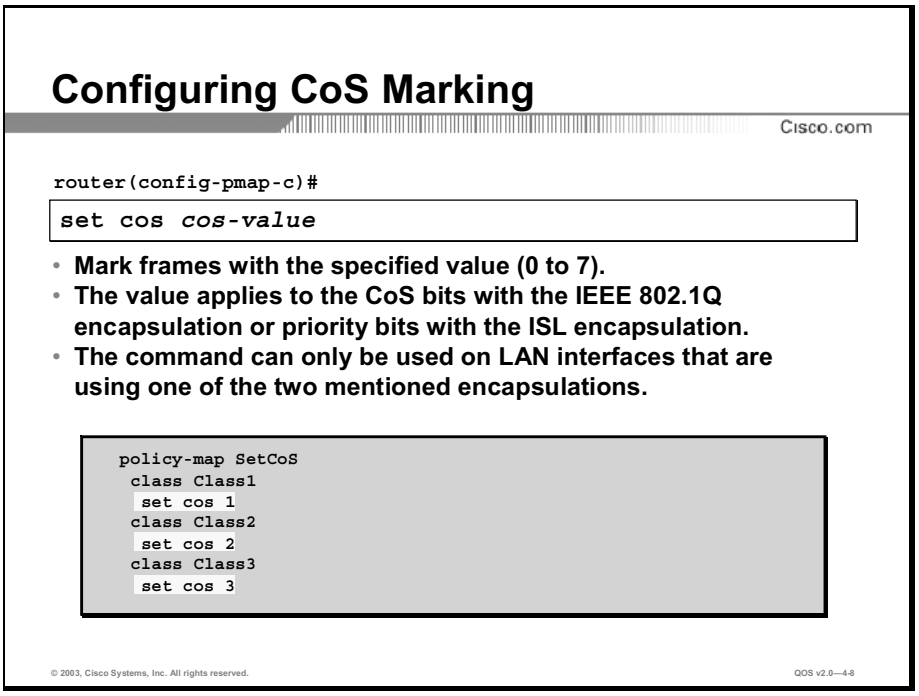

The IEEE 802.1p standard specifies a standard for delivering QoS in LANs. Packets are marked with three CoS bits where CoS values range from zero for low priority to seven for high priority. CoS can only be applied on trunks because VLAN trunking encapsulations designate fields with available space to carry CoS bits. There are currently two widely deployed trunking protocols that can transport CoS markings as follows:

- ISL frame headers have a 1-byte user field that carries the CoS value in the three least significant bits.
- IEEE 802.1p and 802.1q frame headers have a 2-byte TCI field that carries the CoS value in the three most significant bits that are called the user priority bits.

Note: Other frame types (untagged) cannot carry CoS values.

In general, Layer 2 switches can examine, use, or alter MAC layer markings (but not IP precedence or DSCP settings), because IP precedence and DSCP are Layer 3. Layer 2 markings are generally applied on egress trunk ports.

set cos *cos-value*

Syntax Description

Parameter	Description
cos-value	Specific IEEE 802.1Q CoS value from 0 to 7.

Configuring IP Precedence Marking

This topic identifies the Cisco IOS commands that are required to mark IP packets using IP precedence with class-based marking.

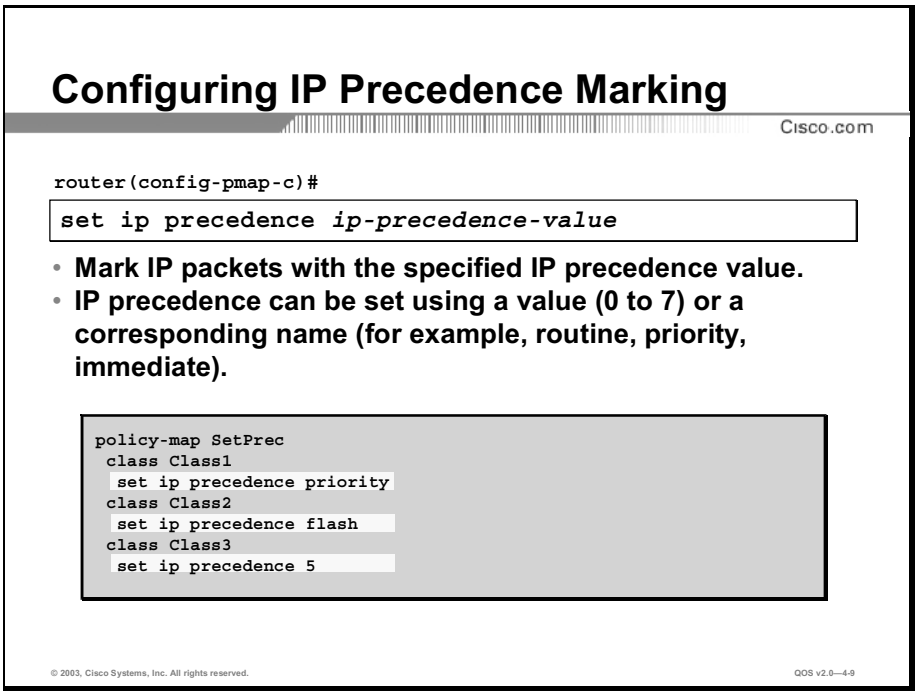

IP precedence is encoded into the three high-order bits of the ToS field in the IP header. It supports eight classes: two of these classes, IP precedence 6 and 7, are reserved and should not be used for user-defined classes. IP precedence 0 is the default value and is usually used for the best-effort class.

To set the precedence value in the IP header, use the **set ip precedence** QoS policy-map configuration command. To leave the precedence value at the current setting, use the **no** form of this command.

set ip precedence *ip-precedence-value*

Syntax Description

Parameter	Description
ip-precedence-value	A number from 0 to 7 that sets the precedence bit in the IP header.

Configuring IP DSCP Marking

This topic identifies the Cisco IOS commands that are required to mark IP packets using DSCP with class-based marking.

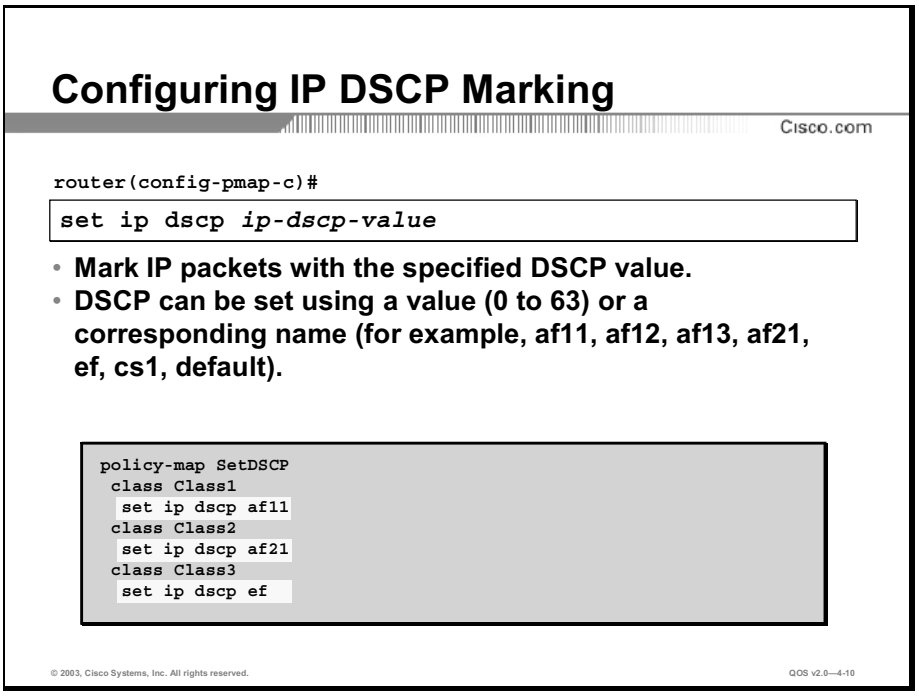

Differentiated Services (DiffServ) is a new model that supercedes—and is backward compatible with—IP precedence. DiffServ uses 6 prioritization bits that permit classification of up to 64 values (0 to 63). A DiffServ value is called a DSCP. The **set ip dscp** command is used to mark packets of a class with a DSCP value.

To mark a packet by setting the IP DSCP in the ToS byte, use the **set ip dscp** QoS policy-map configuration command. To remove a previously set IP DSCP, use the **no** form of this command.

set ip dscp *ip-dscp-value*

Syntax Description

Parameter	Description
ip-dscp-value	A number from 0 to 63 that sets the IP DSCP value. Reserved keywords **EF** (expedited forwarding), **AF11** (assured forwarding class AF11), and **AF12** (assured forwarding class AF12) can be specified instead of numeric values.

Monitoring Class-Based Marking

This topic identifies the Cisco IOS commands that are used to monitor class-based marking.

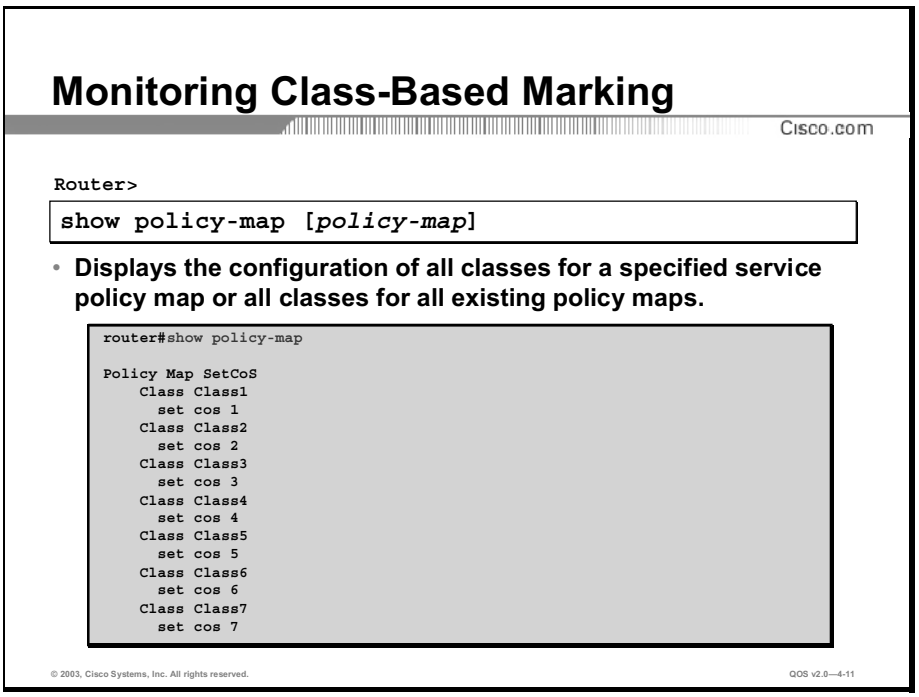

The show policy-map command displays all classes for service policy that is specified in the command line.

To display the configuration of all classes for a specified service policy map or all classes for all existing policy maps, use the **show policy-map** EXEC or privileged EXEC command.

show policy-map [*policy-map*]

Syntax Description

Parameter	Description
policy-map	(Optional) The name of the service policy map whose complete configuration is to be displayed.

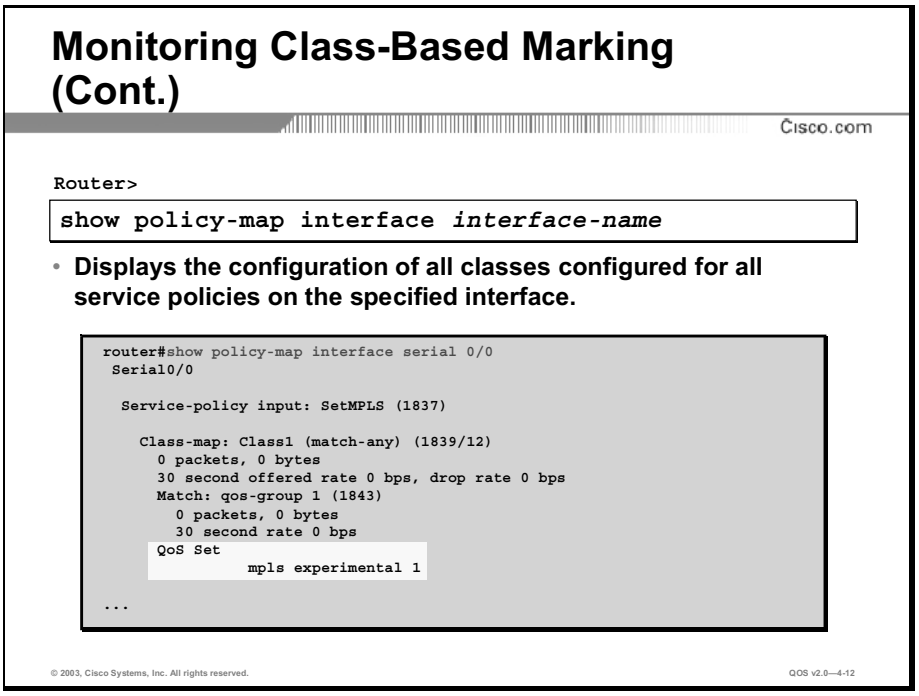

The **show policy-map interface** command displays all service policies applied to the interface. Among the settings, marking parameters and statistics are displayed.

To display the configuration of all classes configured for all service policies on the specified interface or to display the classes for the service policy for a specific permanent virtual circuit (PVC) on the interface, use the **show policy-map interface** EXEC or privileged EXEC command.

show policy-map interface *interface-name* [**vc** [***vpi***/] ***vci***][**dlci** ***dlci***] [**input** | **output**]

Syntax Description

Parameter	Description
interface-name	Name of the interface or subinterface whose policy configuration is to be displayed.
vpi/	(Optional) ATM network virtual path identifier (VPI) for this PVC. The absence of the "/" and a *vpi* value defaults the *vpi* value to 0. On the Cisco 7200 and 7500 series routers, this value ranges from 0 to 255. The *vpi* and *vci* arguments cannot both be set to 0; if one is 0, the other cannot be 0. If this value is omitted, information for all virtual circuits (VCs) on the specified ATM interface or subinterface is displayed.
vci	(Optional) ATM network virtual channel identifier (VCI) for this PVC. This value ranges from 0 to 1 less than the maximum value set for this interface by the **atm vc-per-vp** command. Typically, lower values 0 to 31 are reserved for specific traffic (F4 Operation, Administration, and Maintenance [OAM], switched virtual circuit [SVC] signalling, Integrated Local Management Interface [ILMI], and so on) and should not be used. The VCI is a 16-bit field in the header of the ATM cell. The VCI value is unique only on a single link, not throughout the ATM network, because it has local significance only. The *vpi* and *vci* arguments cannot both be set to 0; if one is 0, the other cannot be 0.
dlci	(Optional) Indicates that a specific PVC for which policy configuration will be displayed.
dlci	(Optional) A specific data-link connection identifier (DLCI) number used on the interface. Policy configuration for the corresponding PVC will be displayed when a DLCI is specified.
input	(Optional) Indicates that the statistics for the attached input policy will be displayed.
output	(Optional) Indicates that the statistics for the attached output policy will be displayed.

Summary

This topic summarizes the key points discussed in this lesson.

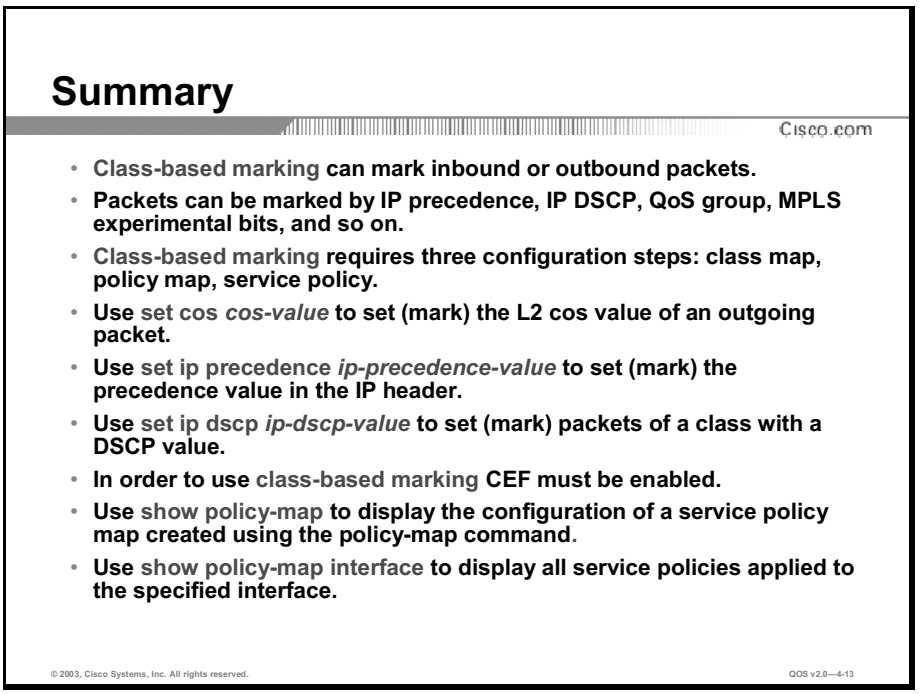

References

For additional information, refer to these resources:

- For more information on Class-Based Marking, refer to “Class-Based Marking” at the following URL: http://www.cisco.com/univercd/cc/td/doc/product/software/ios121/121newft/121t/121t5/cbpmark2.pdf

Next Steps

For the associated lab exercise, refer to the following section of the course Lab Guide:

- Lab Exercise 4-1: Classification and Marking Using MQC

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) Which of the following cannot be accomplished by marking packets or frames?

A) set information in Layer 1 header

B) set information in Layer 2 header

C) set information in Layer 3 header

D) set information in Layer 4 header

Q2) Which three of the following markers are supported by class-based marking? (Choose three.)

A) ATM DE bit

B) Frame Relay group

C) MPLS experimental bits

D) IEEE 802.1Q or ISL CoS/priority bits

E) IP precedence

Q3) Which three of the following are steps in configuring class-based marking? (Choose three.)

A) create a class map

B) create a policy map

C) apply a policy to a class map

D) attach a service policy to an interface

Q4) How many bits is the CoS field?

A) 2 bits

B) 3 bits

C) 4 bits

D) 6 bits

Q5) IP precedence is encoded as how many bits and in which field in the IP header?

A) 3 low order bits of the CoS field

B) 3 low order bits of the ToS field

C) 3 high order bits of the CoS field

D) 3 high order bits of the ToS field

Q6) DSCP uses how may prioritization bits?

A) 2

B) 3

C) 4

D) 6

Q7) Which of the following commands will display the configuration of all classes configured for all service policies on the specified interface?

A) **show policy map s0/0**

B) **show policy-map s0/0 policy default**

C) **show policy-map interface s0/0**

D) **show policy-map interface s0/0 class default**

Quiz Answer Key

Q1) A

Relates to: Class-Based Marking Overview

Q2) C, D, E

Relates to: MQC Marking Options

Q3) A, B, D

Relates to: Configuring Class-Based Marking

Q4) B

Relates to: Configuring CoS Marking

Q5) D

Relates to: Configuring IP Precedence Marking

Q6) D

Relates to: Configuring IP DSCP Marking

Q7) C

Relates to: Monitoring Class-Based Marking

Using NBAR for Classification

Overview

NBAR, a feature in Cisco IOS software, provides intelligent network classification for your infrastructure. NBAR is a classification engine that can recognize a wide variety of applications, including Web-based applications and client and server applications that dynamically assign TCP or UDP port numbers. After the application is recognized, the network can invoke specific services for that particular application. NBAR currently works with QoS features to ensure that the network bandwidth is best used to fulfill company objectives. These features include the ability to guarantee bandwidth to critical applications, limit bandwidth to other applications, drop selective packets to avoid congestion, and mark packets appropriately so that the your network and the service provider network can provide QoS from end to end.

This lesson describes NBAR, a Cisco IOS protocol discovery and classification mechanism. NBAR features covered in this lesson include applications that NBAR can support, Packet Description Language Modules (PDLMs), and NBAR protocol discovery.

Relevance

Classification is a fundamental requirement for any network deployment of QoS. As such, it is of key importance to understand the different ways that traffic can be classified.

Objectives

Upon completing this lesson, you will be able to use NBAR to discover network protocols and to classify packets. This includes being able to meet these objectives:

- Describe the function of NBAR
- Identify the types of applications recognized by NBAR
- Explain the purpose of PDLMs in NBAR
- Describe NBAR protocol discovery and the NBAR Protocol Discovery MIB
- Identify the Cisco IOS commands required to configure and monitor NBAR protocol discovery
- Identify the Cisco IOS commands required to configure NBAR to recognize static port protocols
- Identify the Cisco IOS commands required to configure NBAR to recognize TCP and UDP stateful protocols

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Basic knowledge of internetworking with TCP/IP concepts
- Basic knowledge of the Cisco IOS command-line interface

Outline

The outline lists the topics included in this lesson.

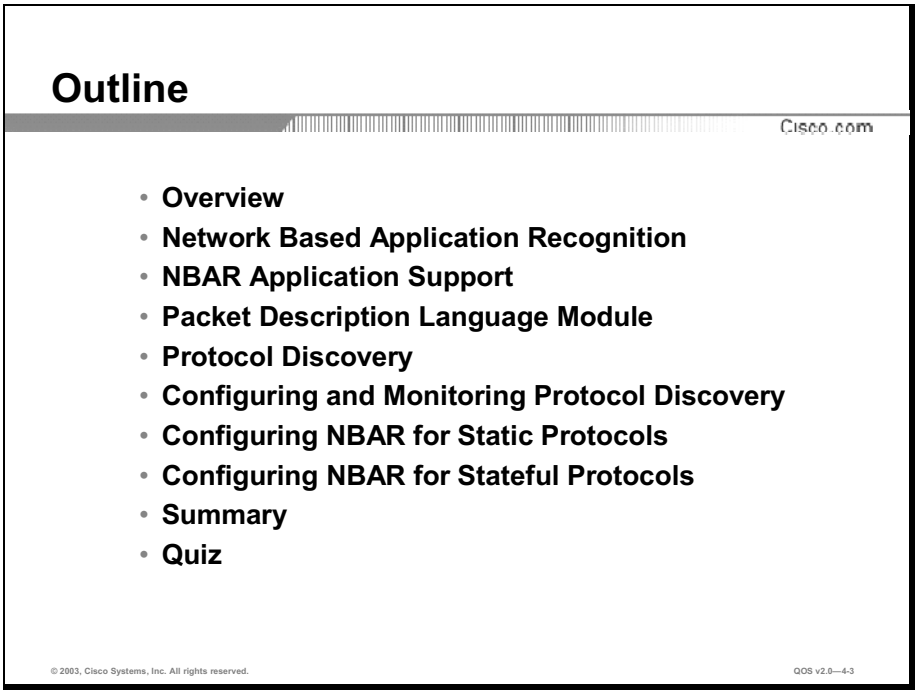

Network Based Application Recognition

This topic describes the NBAR, a Cisco IOS protocol discovery and classification mechanism.

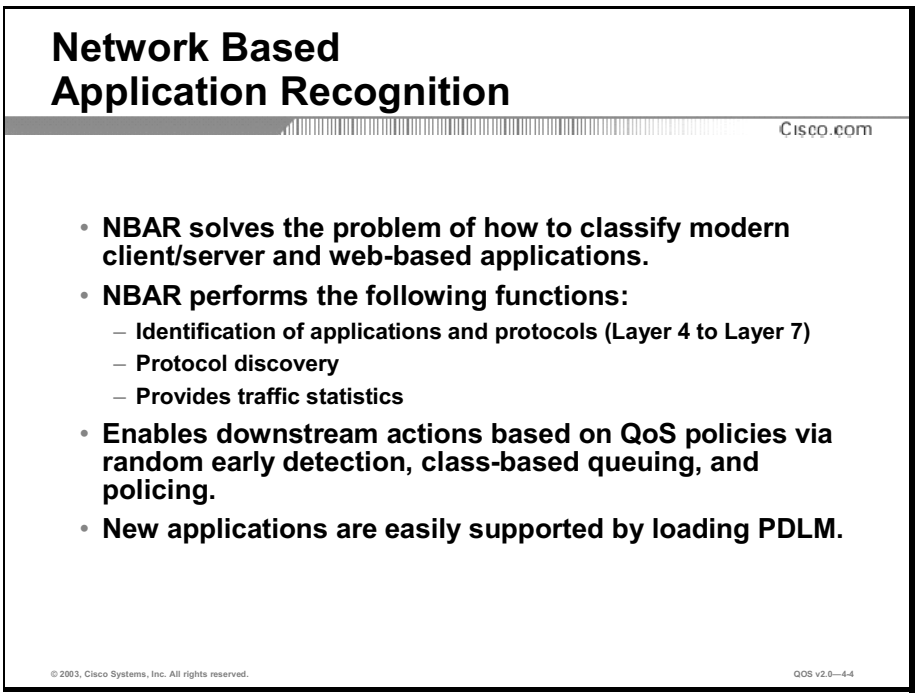

NBAR is an MQC-enabled classification and protocol discovery feature. NBAR can determine the mix of traffic on the network, which is important in isolating congestion problems.

NBAR can classify application traffic by looking beyond the TCP/UDP port numbers of a packet. This is subport classification. NBAR looks into the TCP/UDP payload itself and classifies packets based on the content within the payload, such as transaction identifier, message type, or other similar data.

Classification of HTTP, by URL or Multipurpose Internet Mail Extensions (MIME) type is an example of subport classification. NBAR classifies HTTP traffic by text within the URL, using regular expression matching. NBAR uses the UNIX filename specification as the basis for the URL specification format. The NBAR engine then converts the specification format into a regular expression.

The NBAR protocol discovery feature provides an easy way to discover application protocols that are transiting an interface. The protocol discovery feature discovers any protocol traffic supported by NBAR. Protocol discovery can be applied to interfaces and can be used to monitor both input and output traffic. Protocol discovery maintains the following per-protocol statistics for enabled interfaces: total number of input and output packets and bytes, and input and output bit rates.

An external PDLM can be loaded at run time to extend the NBAR list of recognized protocols. PDLMs can also be used to enhance an existing protocol recognition capability. PDLMs allow NBAR to recognize new protocols without requiring a new Cisco IOS image or a router reload.

NBAR is not supported on the following logical interfaces:

- Fast EtherChannel
- Interfaces configured to use tunneling or encryption

NBAR does not support the following:

- More than 24 concurrent URLs, hosts, or MIME-type matches
- Matching beyond the first 400 bytes in a packet payload
- Multicast and switching modes other than Cisco Express Forwarding (CEF)
- Fragmented packets
- URL/host/MIME classification with secure HTTP
- Packets originating from or destined to the router running NBAR

Note: NBAR cannot be used to classify output traffic on a WAN link where tunneling or encryption is used. Therefore, NBAR should be configured on other interfaces on the router (such as a LAN link) to perform input classification before the traffic is switched to the WAN link for output. However, NBAR protocol discovery is supported on interfaces where tunneling or encryption is used. You can enable protocol discovery directly on the tunnel or on the interface where encryption is performed to gather key statistics on the various applications that are traversing the interface. The input statistics also show the total number of encrypted/tunneled packets received in addition to the per-protocol breakdowns.

To run distributed NBAR on a Cisco 7500 series router, you must be using a processor that has 64 MB of DRAM or more. At the time of this publication, the following processors met this requirement:

- Versatile interface processor (VIP)2-50, VIP4-50, VIP4-80, and VIP6-80
- GigabitEthernet Interface Processor (GEIP) and GEIP+
- Spatial Reuse Protocol Interface Processor (SRPIP)

Note: For the latest information regarding NBAR use restrictions, please refer to the Cisco IOS documentation for your specific software release.

NBAR Application Support

This topic describes how NBAR supports various applications.

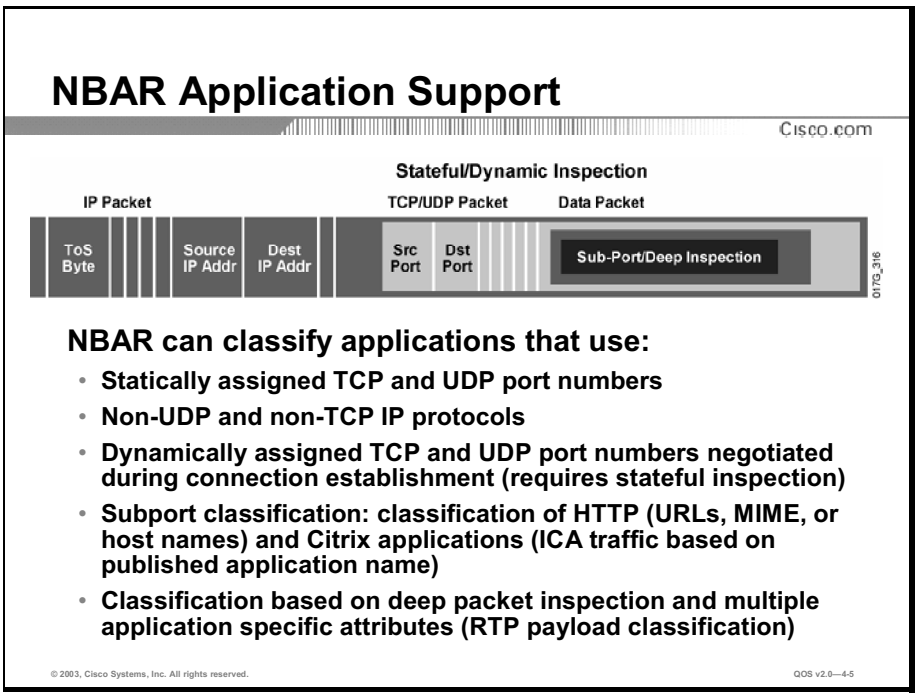

NBAR supports simpler configuration coupled with stateful recognition of flows. The simpler configuration means you do not have to examine a protocol analyzer capture to calculate ports and details. Stateful recognition means smarter, deeper packet recognition.

NBAR can be used to recognize packets belonging to different types of applications:

- Static applications establish sessions to well-known TCP or UDP destination port numbers. Such applications were classified by using access lists.
- Dynamic applications use multiple sessions that use dynamic TCP or UDP port numbers. Typically, there is a control session to a well-known port number and the other sessions are established to destination port numbers negotiated through the control sessions. NBAR inspects the port number exchange through the control session.
- Some non-IP protocols can also be recognized by NBAR.
- NBAR also has the capability to inspect some applications for other information and classify based on that information. For example, NBAR can classify HTTP sessions based on the requested URL, including MIME type or host name.

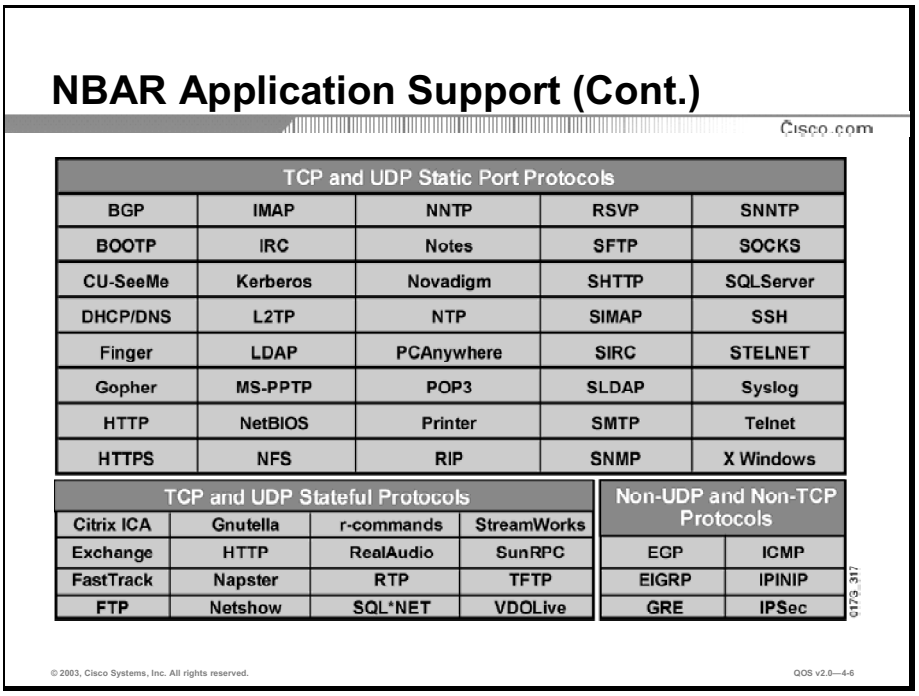

The following table lists the non-TCP and non-UDP protocols supported by NBAR.

Non-TCP and Non-UDP NBAR Supported Protocols

Protocol	Network Protocol	Protocol ID	Description
EGP	IP	8	Exterior Gateway Protocol
GRE	IP	47	Generic Routing Encapsulation
ICMP	IP	1	Internet Control Message Protocol
IPINIP	IP	4	IP in IP
IPSec	IP	50, 51	IP Encapsulating Security Payload/Authentication Header
EIGRP	IP	88	Enhanced Interior Gateway Routing Protocol

Although access lists can also be used to classify applications based on static port numbers, NBAR is easier to configure and can provide classification statistics that are not available when using access lists.

The following table contains the static IP protocols supported by NBAR.

Static TCP and UDP NBAR Supported Protocols

Protocol	Network Protocol	Protocol ID	Description
BGP	TCP/UDP	179	Border Gateway Protocol
CU-SeeMe	TCP/UDP	7648, 7649	Desktop videoconferencing
CU-SeeMe	UDP	24032	Desktop video conferencing
DHCP/ BOOTP	UDP	67, 68	Dynamic Host Configuration Protocol/ Bootstrap Protocol
DNS	TCP/UDP	53	Domain Name System
Finger	TCP	79	Finger user information protocol
Gopher	TCP/UDP	70	Internet Gopher Protocol
HTTP	TCP	80	Hypertext Transfer Protocol
HTTPS	TCP	443	Secured HTTP
IMAP	TCP/UDP	143, 220	Internet Message Access Protocol
IRC	TCP/UDP	194	Internet Relay Chat
Kerberos	TCP/UDP	88, 749	Kerberos Network Authentication Service
L2TP	UDP	1701	L2F/L2TP tunnel
LDAP	TCP/UDP	389	Lightweight Directory Access Protocol
MS-PPTP	TCP	1723	Microsoft Point-to-Point Tunneling Protocol for VPN
MS-SQLServer	TCP	1433	Microsoft SQL Server Desktop Videoconferencing
NetBIOS	TCP	137, 139	NetBIOS over IP (MS Windows)
NetBIOS	UDP	137, 138	NetBIOS over IP (MS Windows)
NFS	TCP/UDP	2049	Network File System
NNTP	TCP/UDP	119	Network News Transfer Protocol
Notes	TCP/UDP	1352	Lotus Notes
Novadigm	TCP/UDP	3460-3465	Novadigm Enterprise Desktop Manager (EDM)
NTP	TCP/UDP	123	Network Time Protocol
PCAnywhere	TCP	5631, 65301	Symantec PCAnywhere
PCAnywhere	UDP	22, 5632	Symantec PCAnywhere
POP3	TCP/UDP	110	Post Office Protocol
Printer	TCP/UDP	515	Printer
RIP	UDP	520	Routing Information Protocol
RSVP	UDP	1698,17	Resource Reservation Protocol
SFTP	TCP	990	Secure FTP

Protocol	Network Protocol	Protocol ID	Description
SHTTP	TCP	443	Secure HTTP
SIMAP	TCP/UDP	585, 993	Secure IMAP
SIRC	TCP/UDP	994	Secure IRC
SLDAP	TCP/UDP	636	Secure LDAP
SNNTP	TCP/UDP	563	Secure NNTP
SMTP	TCP	25	Simple Mail Transfer Protocol
SNMP	TCP/UDP	161, 162	Simple Network Management Protocol
SOCKS	TCP	1080	Firewall security protocol
SPOP3	TCP/UDP	995	Secure POP3
SSH	TCP	22	Secured Shell
STELNET	TCP	992	Secure Telnet
Syslog	UDP	514	System Logging Utility
Telnet	TCP	23	Telnet Protocol
X Windows	TCP	6000-6003	X11, X Windows

The following table lists the dynamic (or stateful) protocols supported by NBAR.

Stateful NBAR Supported Protocols

Stateful Protocol	Transport Protocol	Description
FTP	TCP	File Transfer Protocol
Exchange	TCP	MS-RPC for Exchange
HTTP	TCP	HTTP with URL, MIME, or Host classification
Netshow	TCP/UDP	Microsoft Netshow
Realaudio	TCP/UDP	RealAudio Streaming Protocol
r-commands	TCP	rsh, rlogin, rexec
StreamWorks	UDP	Xing Technology Stream Works audio and video
SQL*NET	TCP/UDP	SQL*NET for Oracle
SunRPC	TCP/UDP	Sun Remote Procedure Call
TFTP	UDP	Trivial File Transfer Protocol
VDOLive	TCP/UDP	VDOLive Streaming Video

Packet Description Language Module

This topic describes PDLM.

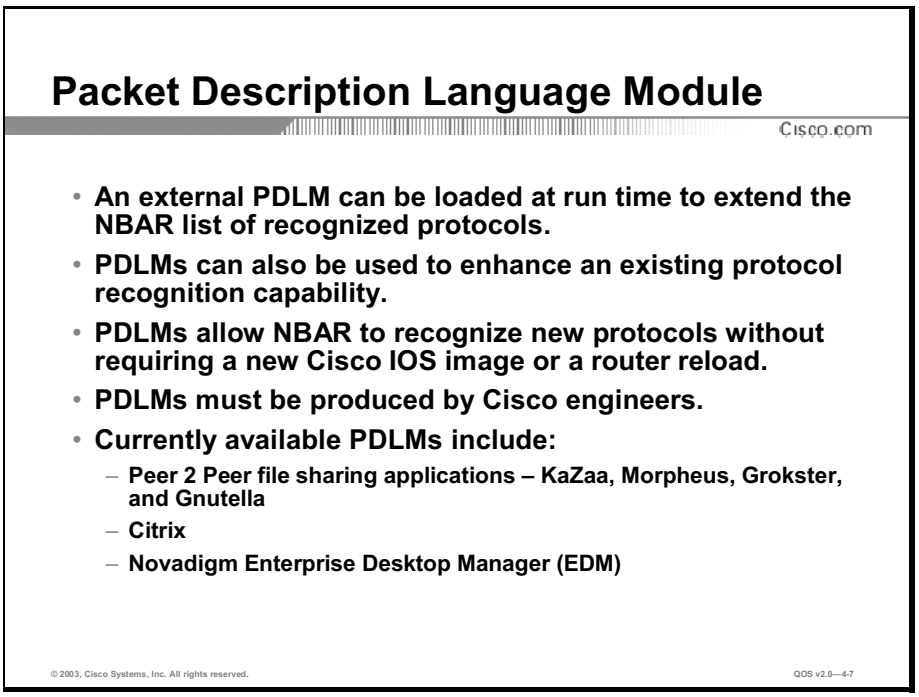

New features are usually added to new versions of the Cisco IOS software. NBAR is the first mechanism that supports dynamic upgrades without having to change the Cisco IOS version or restart a router.

PDLMs contain the rules that are used by NBAR to recognize an application and can be used to bring new or changed functionality to NBAR.

An external PDLM can be loaded at run time to extend the NBAR list of recognized protocols. PDLMs can be used to enhance an existing protocol recognition capability. PDLMs allow NBAR to recognize new protocols without requiring a new Cisco IOS image or a router reload.

Note: New PDLMs are released only by Cisco and are available from local Cisco representatives. They can be loaded from flash memory. Registered users can find them at: http://www.cisco.com/cgi-bin/tablebuild.pl/pdlm

To extend or enhance the list of protocols recognized by NBAR through a PDLM provided by Cisco, use the **ip nbar pdlm** configuration command. Use the **no** form of this command to unload a PDLM if it was previously loaded.

Use the **show ip nbar port-map** command to display the current protocol-to-port mappings in use by NBAR.

Protocol Discovery

This topic describes the NBAR protocol discovery feature.

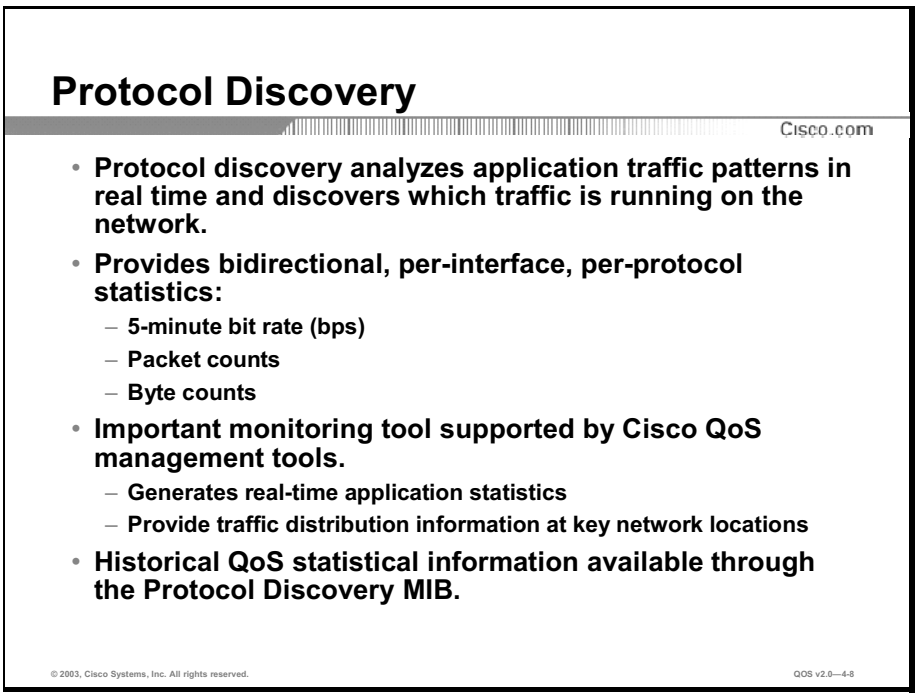

To develop and apply QoS policies, NBAR includes a protocol discovery feature that provides an easy way to discover application protocols that are transiting an interface. The protocol discovery feature discovers any protocol traffic that is supported by NBAR.

NBAR protocol discovery captures key statistics associated with each protocol in a network. These statistics can be used to define traffic classes and QoS policies for each traffic class.

Protocol discovery can be applied to interfaces and can be used to monitor both input and output traffic. In addition, protocol discovery shows the mix of applications currently running on the network. This helps in defining QoS classes and polices, such as how much bandwidth to provide to mission-critical applications and to determine which protocols should be policed. The following per-protocol, bidirectional statistics are available:

- Packet and byte counts
- Bit rates

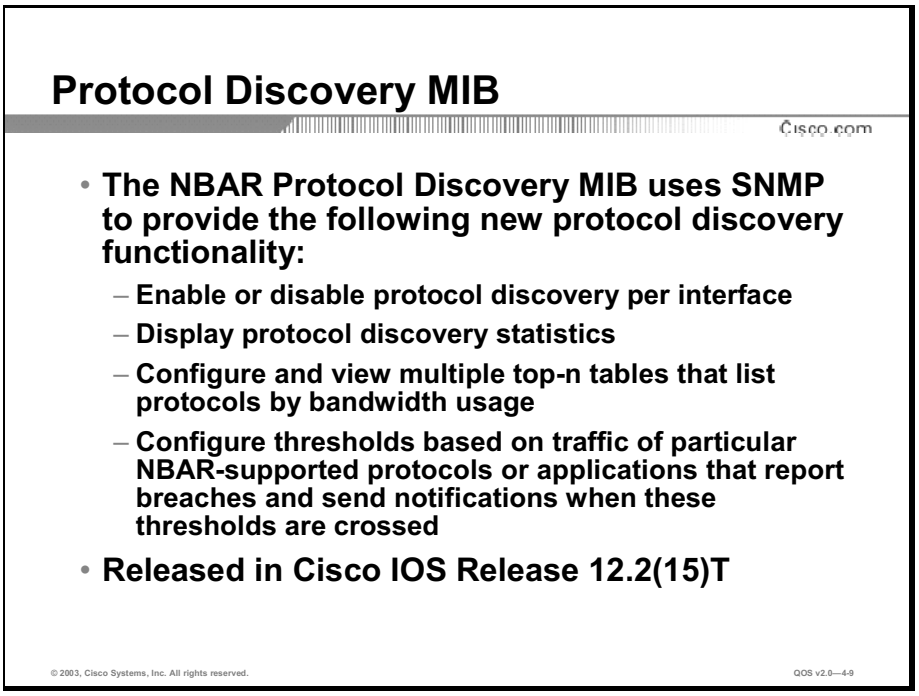

The Protocol Discovery Management Information Base (MIB) allows users to:

- Enable NBAR protocol discovery on multiple interfaces across multiple routers in a network
- Gather statistics
- Set traps and threshold alarms on protocols
- Study historical trending for a whole network.

Configuring and Monitoring Protocol Discovery

This topic identifies the Cisco IOS commands that are required to configure and monitor NBAR protocol discovery.

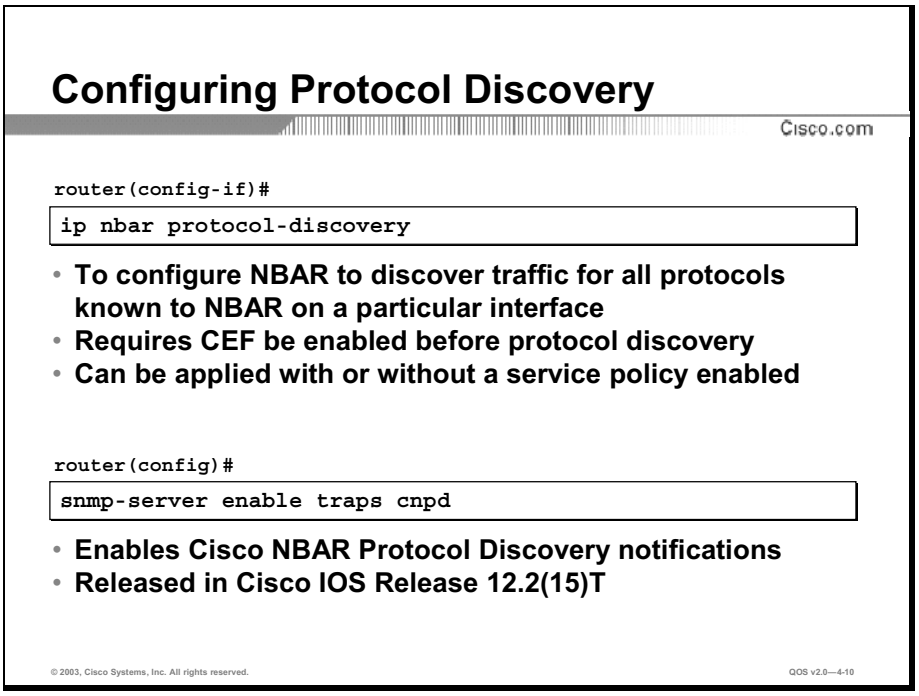

The NBAR feature has two components:

- One component monitors applications traversing a network.
- The other component classifies traffic by protocol.

In order to monitor applications traversing a network, protocol discovery must be enabled. The ability to classify traffic by protocol using NBAR and then applying QoS to the classified traffic is configured using the MQC.

Use the **ip nbar protocol-discovery** command to configure NBAR to keep traffic statistics for all protocols known to NBAR. Protocol discovery provides an easy way to discover application protocols transiting an interface so that QoS policies can be developed and applied. The protocol discovery feature discovers any protocol traffic supported by NBAR. Protocol discovery can be used to monitor both input and output traffic and can be applied with or without a service policy enabled.

Note: You must enable CEF before you configure NBAR. For more information on CEF, refer to Cisco Express Forwarding Overview at the following URL: http://www.cisco.com/univercd/cc/td/doc/product/software/ios121/121cgcr/switch_c/xcprt2/xcdcef.htm.

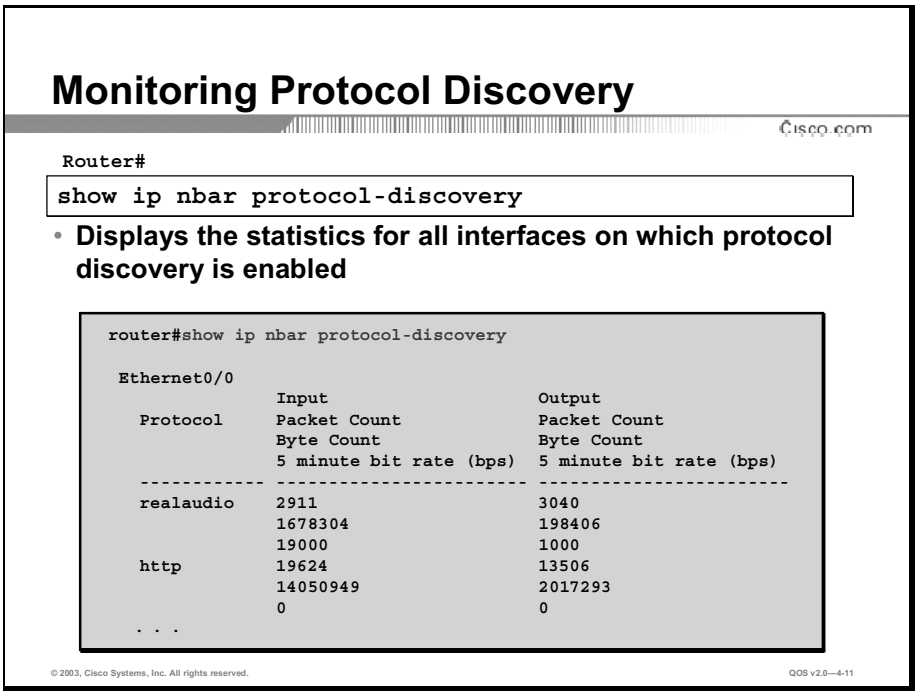

Use the **show ip nbar protocol-discovery** command to display statistics gathered by the NBAR protocol discovery feature. This command, by default, displays statistics for all interfaces on which protocol discovery is currently enabled. The default output of this command includes—in the following order—input bit rate (bps), input byte count, input packet count, and protocol name. Output statistics include packet count, byte count, and the output bit rate in bps.

Protocol discovery can be used to monitor both input and output traffic and can be applied with or without a service policy enabled. NBAR protocol discovery gathers statistics for packets switched to output interfaces. These statistics are not necessarily for packets that exited the router on the output interfaces because packets might have been dropped after switching for various reasons (policing at the output interface, access lists, or queue drops). The example displays partial output of the **show ip nbar protocol-discovery** command for an Ethernet interface.

Configuring NBAR for Static Protocols

This topic identifies the Cisco IOS commands that are required to configure NBAR for static protocols.

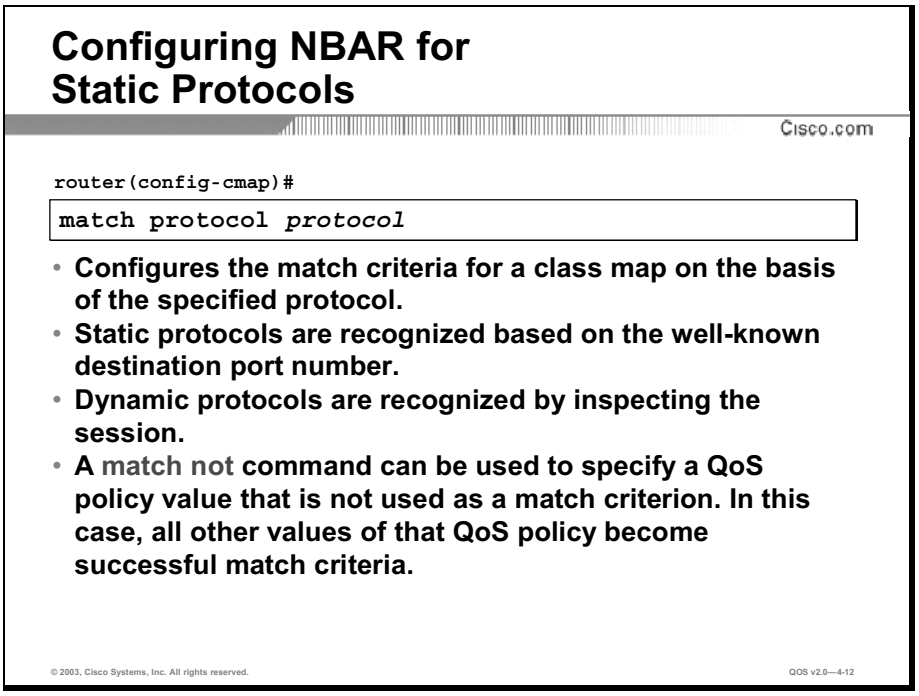

When configuring NBAR the administrator does not need to understand the way a certain protocol works. The configuration simply requires the administrator to enter the name of the protocol (static or stateful).

match protocol protocol-name

Syntax Description

Parameter	Description
`protocol-name`	Name of the protocol used as a matching criterion. Supported protocols include the following (some protocols omitted. Refer to Cisco IOS documentation for complete details): **aarp**—AppleTalk Address Resolution Protocol **arp**—IP Address Resolution Protocol (ARP) **bridge**—bridging **cdp**—Cisco Discovery Protocol **compressedtcp**—compressed TCP **dlsw**—data-link switching **ip**—IP **ipx**—Novell IPX

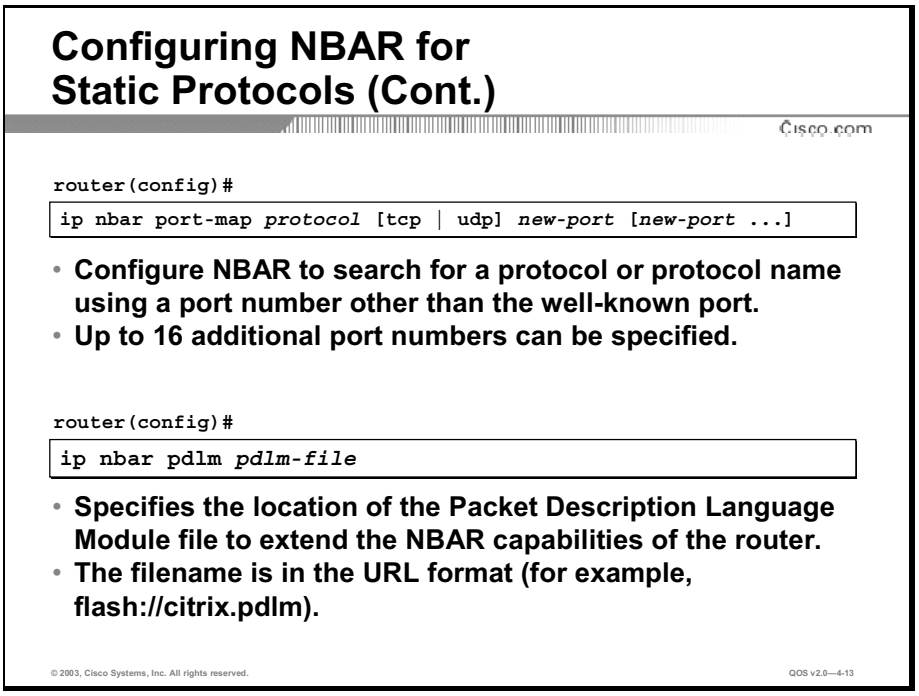

Some protocols (static or stateful) can use additional TCP or UDP ports. Use the **ip nbar port-map** command to extend the NBAR functionality for well-known protocols to new port numbers.

To extend or enhance the list of protocols recognized by NBAR through a Cisco PDLM, use the **ip nbar pdlm** global configuration command.

The *pdlm-file* parameter should be in the URL format and can point to the flash where the Cisco IOS software is stored (for example, **flash://citrix.pdlm**). The file can also be located on a TFTP server (for example, **tftp://10.1.1.1/nbar.pdlm**).To unload a PDLM if it was previously loaded, use the **no** form of this command.

ip nbar pdlm pdlm-name

Syntax Description

Parameter	Description
pdlm-name	The URL where the PDLM can be found on the Flash card.

Example: Configuring NBAR for Static Protocols

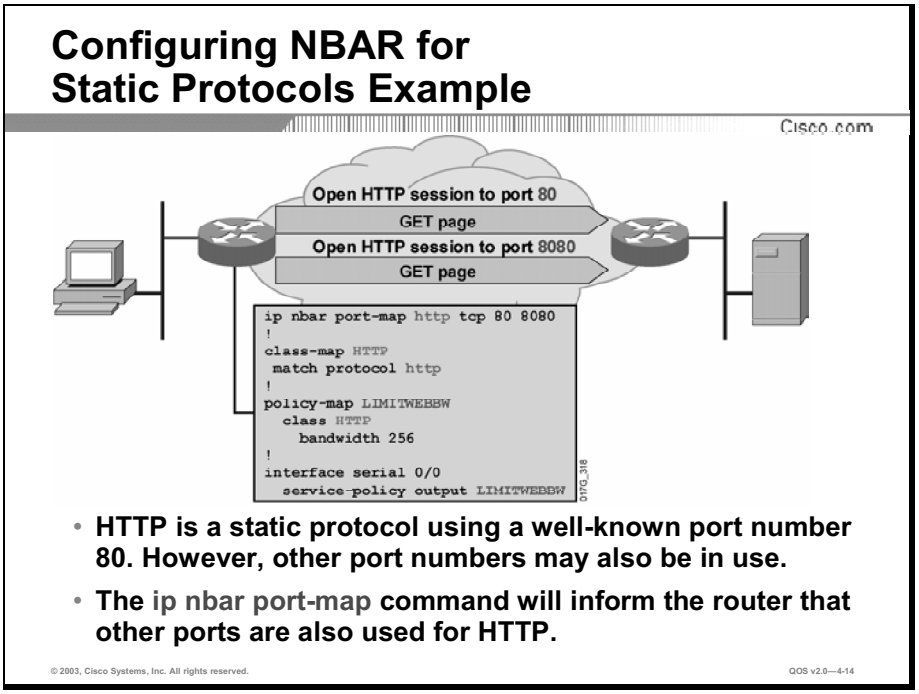

The example illustrates a simple classification of all HTTP sessions. HTTP sessions using the default well-known TCP port number 80 are simple to classify (it is a static protocol).

HTTP is often used on other port numbers. The example shows the usage of the **ip nbar port-map** command to also enable HTTP recognition on TCP port 8080.

The class map called HTTP is used to match the http protocol. The policy map LIMITWEBB will use the class map HTTP and set the bandwidth for HTTP traffic to 256.

The policy map is then applied as a service policy for outbound traffic on S0/0.

Configuring NBAR for Stateful Protocols

This topic identifies the Cisco IOS commands that are required to configure NBAR for stateful protocols.

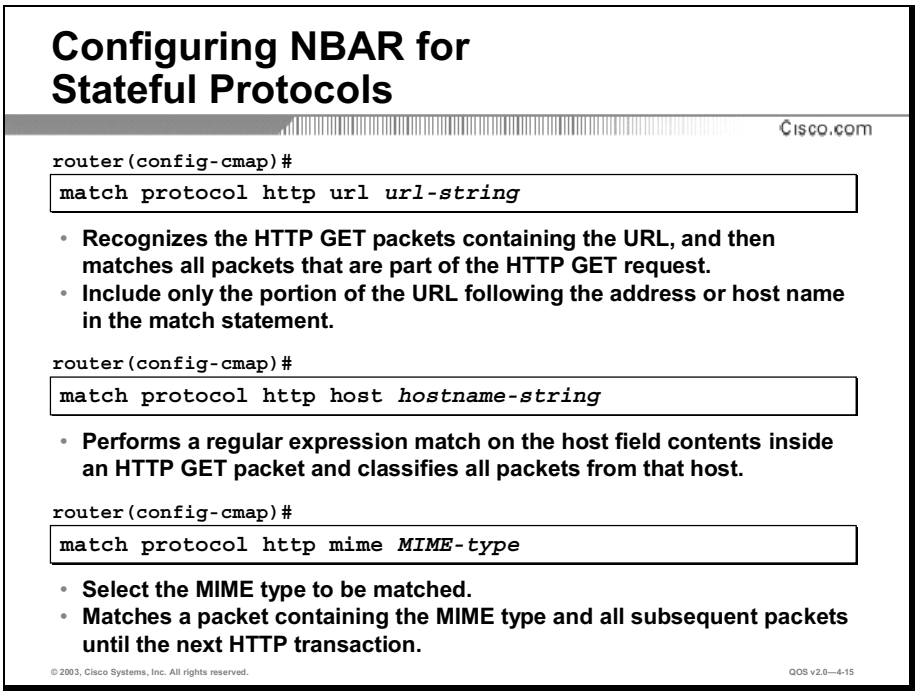

NBAR has enhanced classification capabilities for HTTP. It can classify packets belonging to HTTP flows based on:

- URL portion after the host name, which appears in the GET request of the HTTP session
- Host name specified in the GET request
- MIME type specifying the type of object in the HTTP response

Note: The **match protocol** command has been discussed previously in this lesson

The following example classifies, within the class map called “class1,” HTTP packets based on any URL containing the string “whatsnew/latest” followed by zero or more characters:

```
class-map class1
 match protocol http url whatsnew/latest*
```

The following example classifies, within the class map called “class2,” packets based on any host name containing the string “cisco” followed by zero or more characters:

```
class-map class2
 match protocol http host cisco*
```

The following example classifies, within the class map called “class3,” packets based on the Joint Photographics Expert Group (JPEG) MIME type:

```
class-map class3
 match protocol http mime "*jpeg"
```


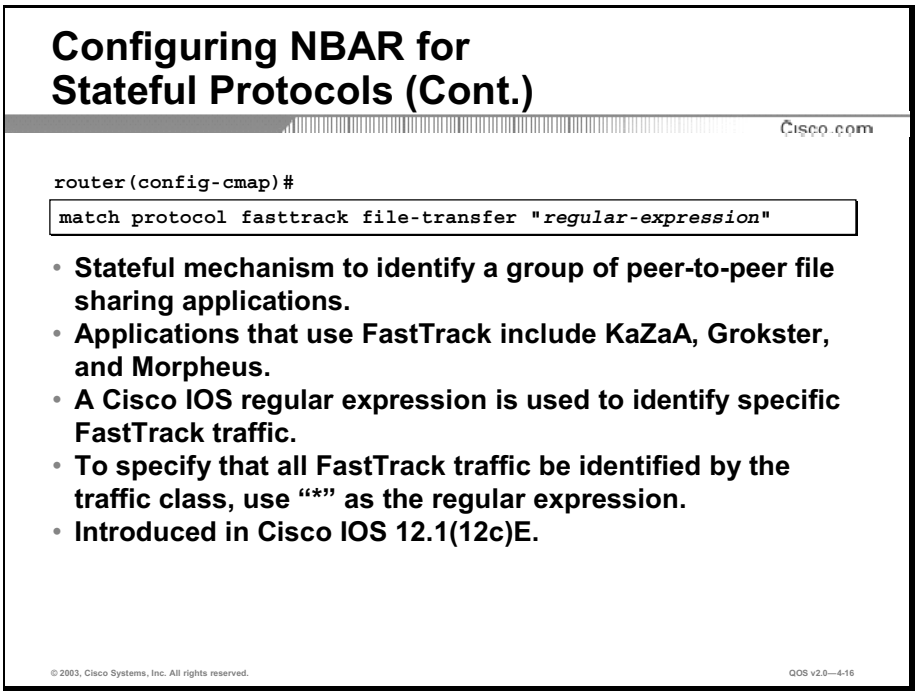

Applications that use FastTrack include KaZaA, Grokster, and Morpheus (although newer versions of Morpheus use Gnutella).

A regular expression is used to identify specific FastTrack traffic. For instance, entering "cisco" as the regular expression would classify the FastTrack traffic containing the string "cisco" as matches for the traffic policy.

To specify that all FastTrack traffic be identified by the traffic class, use "*" as the regular expression.

The following example configures NBAR to match all FastTrack traffic:

```
match protocol fasttrack file-transfer "*"
```

In the following example, all FastTrack files that have the ".mpeg" extension will be classified into class map nbar.

```
class-map match-all nbar
match protocol fasttrack file-transfer "*.mpeg"
```

The following example configures NBAR to match FastTrack traffic that contains the string "cisco":

```
match protocol fasttrack file-transfer "*cisco*"
```

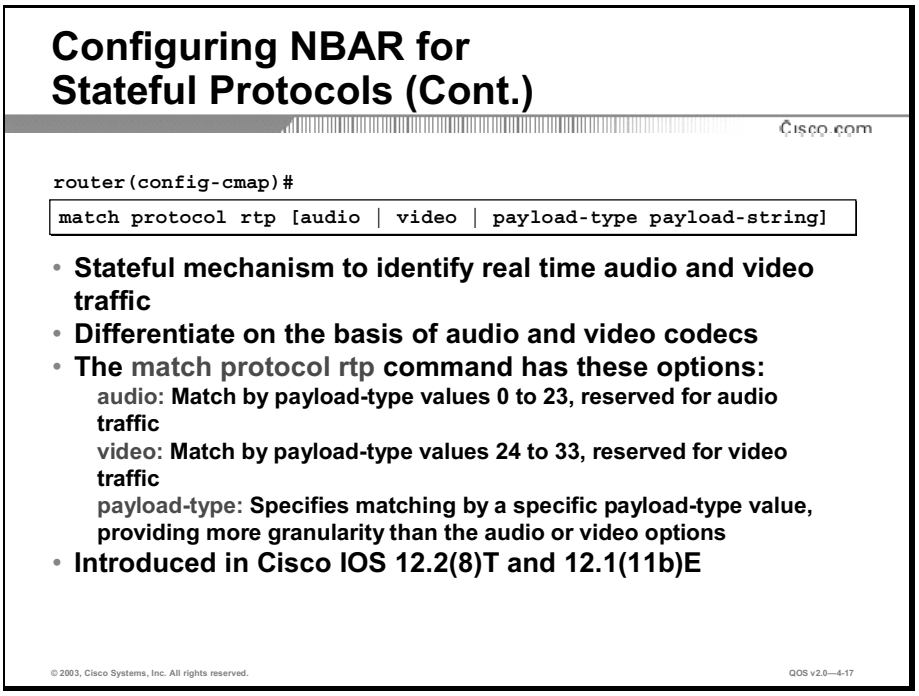

RTP consists of a data and a control part. The control part is called RTC Protocol (RTCP). It is important to note that the NBAR RTP payload classification feature does not identify RTCP packets, and that RTCP packets run on odd numbered ports while RTP packets run on even numbered ports.

The data part of RTP is a thin protocol providing support for applications with real-time properties (such as continuous media [audio and video]), which includes timing reconstruction, loss detection, and security and content identification. The RTP payload type is the data transported by RTP in a packet (for example, audio samples or compressed video data).

NBAR RTP payload classification not only allows one to statefully identify real-time audio and video traffic, but it also can differentiate on the basis of audio and video codecs to provide more granular QoS. The RTP payload classification feature, therefore, looks deep into the RTP header to classify RTP packets.

The payload string parameter can contain commas to separate payload-type values and hyphens to indicate a range of payload-type values. A *payload-string* can be specified in hexadecimal (prepend 0x to the value) and binary (prepend b to the value) notations in addition to standard number values.

NBAR RTP payload type classification was first introduced in Cisco IOS Release 12.2(8)T and is also available in Cisco IOS Release 12.1(11b)E.

Example: Configuring NBAR for Stateful Protocols

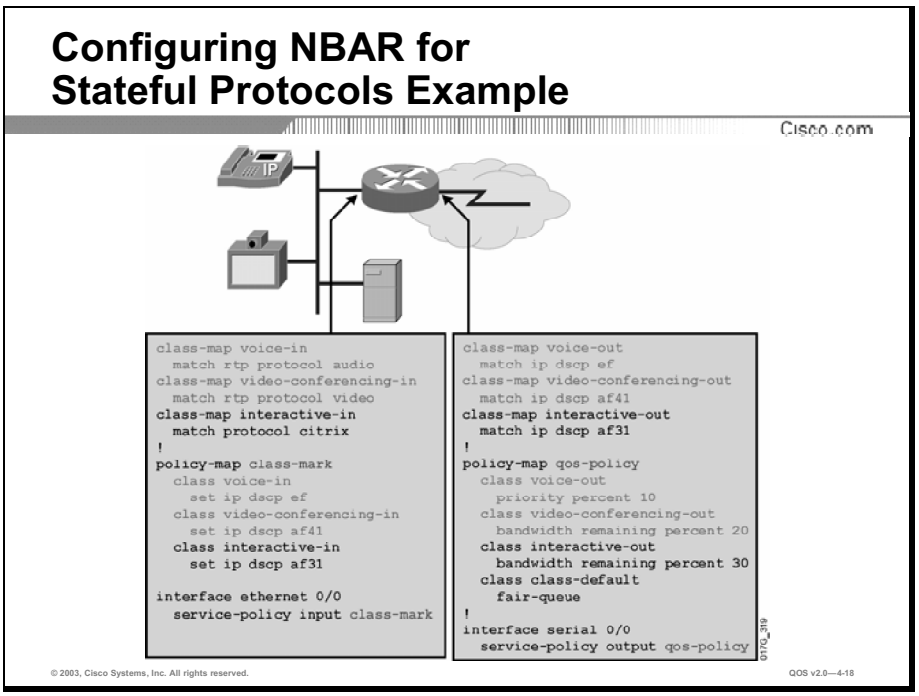

The example illustrates a simple classification of RTP sessions, both on the input interface and on the output interface of the router.

On the input interface three class maps have been created: voice-in, videoconferencing-in, and interactive-in. The voice-in class map will match the RTP audio protocol; the videoconferencing-in class map will match the RTP video protocol and the interactive-in class map will match the Citrix protocol.

The policy map class mark will then do the following:

- If the packet matches the voice-in class map, the packet DSCP field will be set to EF. If the packet matches the videoconferencing-in class map the packet DSCP field will be set to AF 41. If the packet matches the interactive-in class map, the DSCP field will be set to AF 31.
- The policy map class mark is applied to the input interface, E0/0.

On the output three class maps have been created, voice-out, videoconferencing-out, and interactive-out. The voice-out class map will match the DSCP field, EF. The videoconferencing-out class map will match the DSCP field, AF 41, and the interactive-out class map will match the DSCP field, AF 31.

In the figure, policy-map **qos-policy** will then do the following:

- If the packet matches the class map voice-out, the packet priority will be set to 10 percent of the bandwidth. If the packet matches the class map videoconferencing-out, the packet priority will be set to 20 percent of the bandwidth. If the packet matches the class map interactive-out, the packet priority will be set to 30 percent of the bandwidth. All other packets will be classified as class-default and fair-queuing will be performed on them.
- The policy map class mark is applied to the output interface, S0/0.

Summary

This topic summarizes the key points discussed in this lesson.

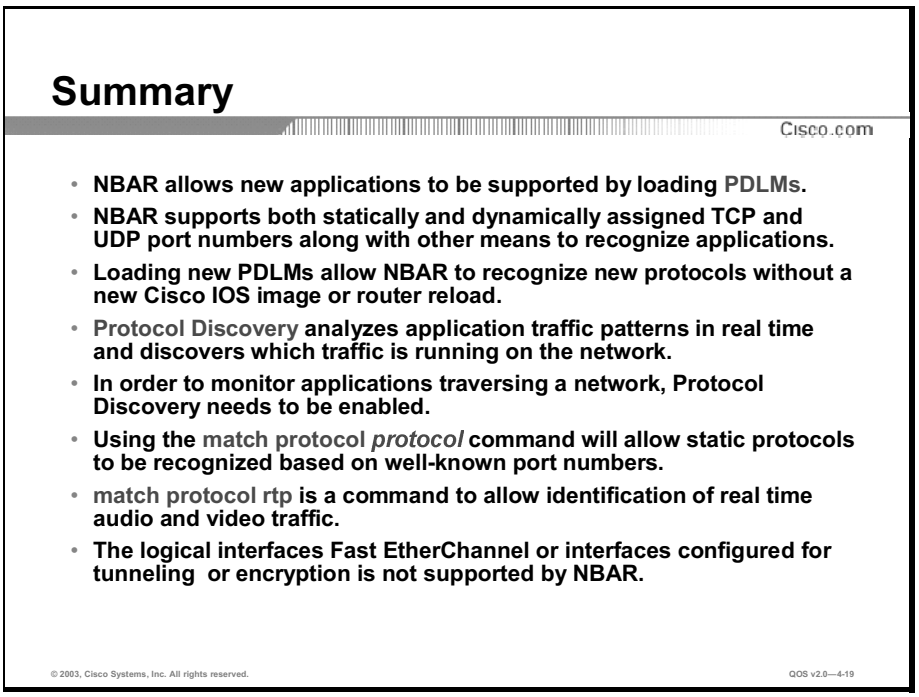

References

For additional information, refer to this resource:

- For a description of all NBAR features and commands, refer to "Network-Based Application Recognition" at the following URL: http://www.cisco.com/univercd/cc/td/doc/product/software/ios121/121newft/121t/121t5/dtnbar.htm

Next Steps

For the associated lab exercise, refer to the following section of the course Lab Guide:

- Lab Exercise 4-2: Classification using NBAR

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) NBAR classification of HTTP (by URL or MIME type) is an example of _____.

A) PDLM

B) port classification

C) subport classification

D) FastTrack

Q2) Which two of the following are NBAR supported applications? (Choose two.)

A) BGP

B) EIGRP

C) OSPF

D) RIP

Q3) Which of the following is correct regarding extending application support for NBAR?

A) FastTrack

B) QoS

C) PDLM

D) RTP

Q4) Which of the following is not a feature of Protocol Discovery?

A) bidirectional

B) per-interface

C) per-protocol statistics

D) FastTrack run time protocol

Q5) The **ip nbar protocol-discovery** configuration command is performed at the _____.

A) class map configuration level

B) global configuration level

C) interface configuration level

D) router configuration level

Q6) The match protocol *protocol* configuration command is performed at the _____.

A) class map configuration level

B) global configuration level

C) interface configuration level

D) router configuration level

Q7) What will be the result of the following configuration?

```
ip nbar port-map http tcp 80 8080
!
class-map HTTP
 match protocol http
!
policy-map LIMITWEBBW
 class HTTP
 bandwidth 256
!
interface serial 0/0
 service-policy output LIMITWEBBW
```

A) All HTTP traffic will be matched.

B) Only port 80 traffic will be matched.

C) Only port 8080 traffic will be matched.

D) Both port 80 and port 8080 traffic will be matched.

Quiz Answer Key

Q1) C

Relates to: Network Based Application Recognition

Q2) A, B

Relates to: NBAR Application Support

Q3) C

Relates to: Packet Description Language Module

Q4) D

Relates to: Protocol Discovery

Q5) C

Relates to: Configuring and Monitoring Protocol Discovery

Q6) A

Relates to: Configuring NBAR for Static Protocols

Q7) D

Relates to: Configuring NBAR for Stateful Protocols

Configuring QoS Pre-Classify

Overview

The QoS for virtual private networks (VPNs) feature (QoS pre-classify) provides a solution for ensuring Cisco IOS QoS services operate in conjunction with tunneling and encryption on an interface. Cisco IOS software can classify packets and apply the appropriate QoS service *before* the data is encrypted and tunneled. The QoS for VPN feature allows users to look inside a packet so that packet classification is based on original port numbers and on source and destination IP addresses. This allows service providers and enterprises to treat mission-critical or multiservice traffic with higher priority across their networks while using VPNs for secure transport.

The QoS pre-classify feature is designed for tunnel interfaces. When the new feature is enabled, the QoS features on the output interface classify packets before encryption, allowing traffic flows to be adjusted in congested environments. The result is more effective packet tunneling.

This lesson describes QoS pre-classify, using QoS policies on VPN interfaces and configuring and monitoring VPN QoS.

Relevance

Classification is a fundamental requirement for any network deployment of QoS. As such, it is of key importance to understand how traffic can be classified in a VPN network.

Objectives

Upon completing this lesson, you will be able to use the QoS pre-classify feature to classify GRE, IPSec, and L2F and L2TP encapsulated packets. This includes being able to meet these objectives:

- Describe the purpose of pre-classification to support QoS in various VPN (IPSec, GRE, L2TP) configurations
- Differentiate situations where pre-classification is appropriate from those where it is not
- Identify the different VPN applications (IPSec, GRE, L2TP) that support QoS pre-classification
- Identify the Cisco IOS commands required to support IPSec, GRE, and L2TP QoS pre-classification
- Identify the Cisco IOS commands used to monitor IPSec, GRE, and L2TP QoS pre-classification

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Basic knowledge of internetworking with TCP/IP concepts
- Basic knowledge of the Cisco IOS command-line interface

Outline

The outline lists the topics included in this lesson.

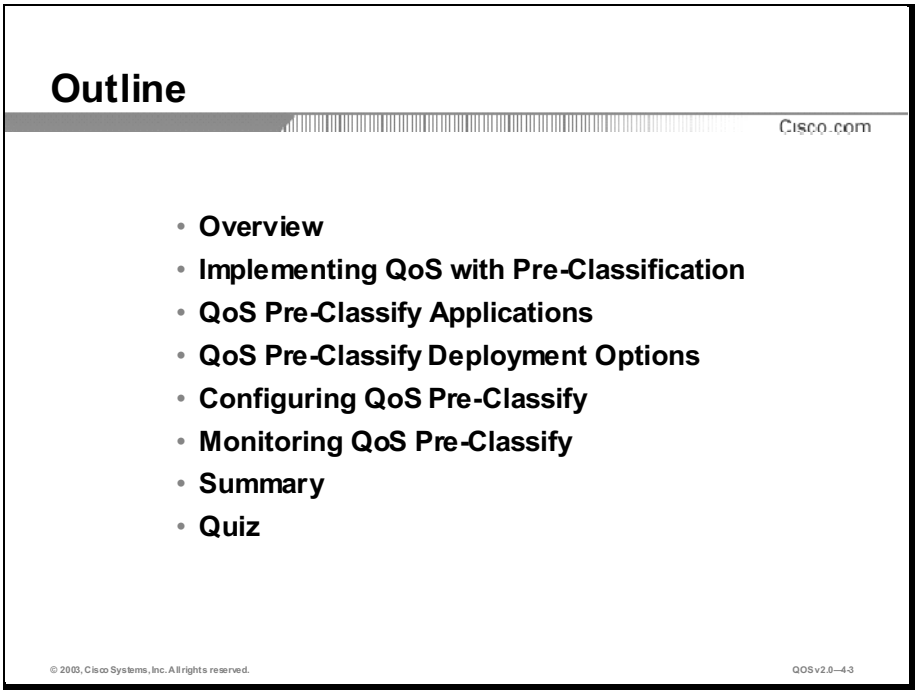

Implementing QoS with Pre-Classification

This topic presents an overview and the purpose for the VPN QoS feature.

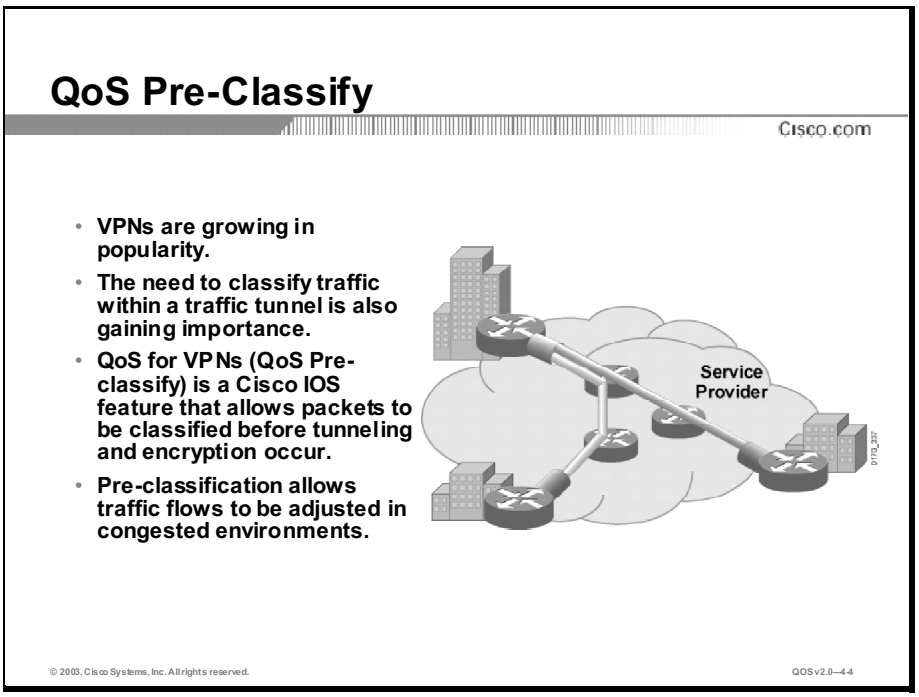

The QoS for virtual private networks (VPNs) feature (QoS pre-classify) is designed for tunnel interfaces. When the feature is enabled, the QoS features on the output interface classify packets *before* encryption, allowing traffic flows to be adjusted in congested environments. The result is more effective packet tunneling.

The QoS pre-classify feature provides a solution for making Cisco IOS QoS services operate in conjunction with tunneling and encryption on an interface. Cisco IOS software can classify packets and apply the appropriate QoS service before the data is encrypted and tunneled. The QoS for VPN feature allows users to look inside the packet so that packet classification can be done based on original port numbers and based on source and destination IP addresses. This allows the service provider to treat mission-critical or multiservice traffic with higher priority across its network.

QoS pre-classify is supported for Generic Routing Encapsulation (GRE), IP in IP (IPIP) tunnels, Layer 2 Tunneling Protocol (L2TP), Layer 2 Forwarding (L2F), Point-to-Point Tunneling Protocol (PPTP), and IPSec.

QoS Pre-Classify Applications

This topic describes some of the VPN applications that support QoS pre-classification.

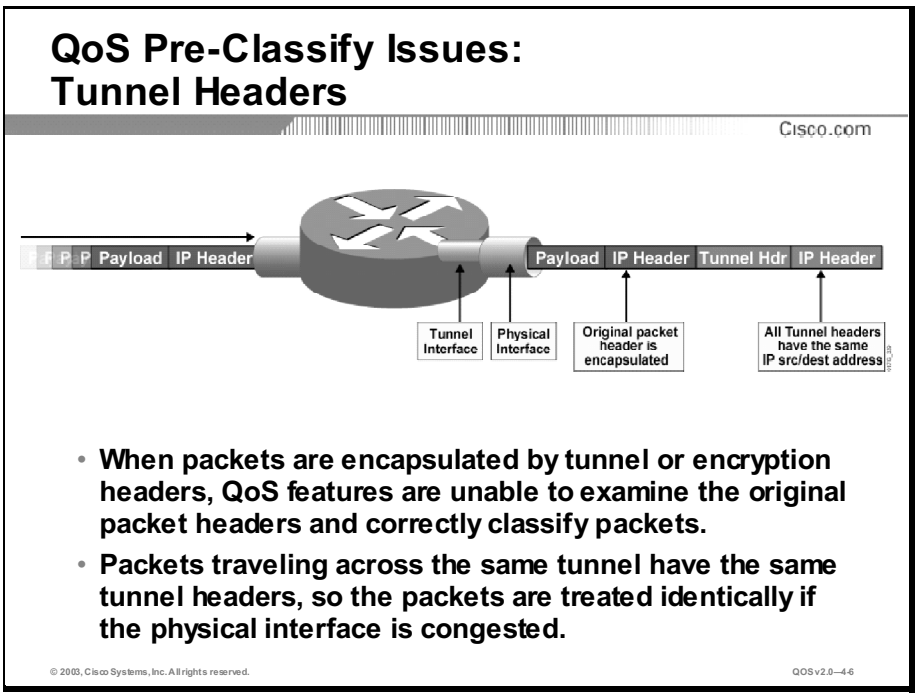

When packets are encapsulated by a tunneling or encryption protocol, the original packet header is no longer available for examination. From the QoS perspective, without the capability to examine the original packet header, providing differentiated levels of service becomes challenging. The main issue is that the QoS parameter normally found in the header of the IP packet should be reflected in the tunnel packet header, regardless of the type of tunnel in use.

Consider the four primary tunneling protocols relevant to VPNs:

- L2TP
- IPSec
- L2F
- GRE

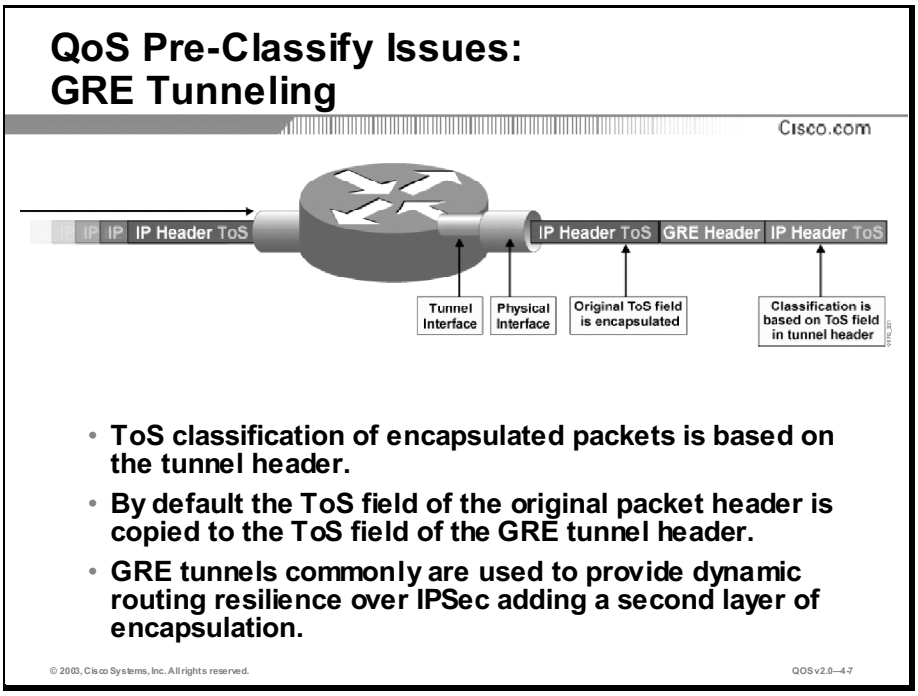

GRE tunnels that are based on RFC 1702 allow any protocol to be tunneled in an IP packet. Today, Cisco offers support for encapsulation of data using either IPSec or GRE. In either of these scenarios, Cisco IOS software offers the ability to copy the IP ToS values from the packet header into the tunnel header. This feature, which appears in Cisco IOS version 11.3T, allows the ToS bits to be copied to the tunnel header when the router encapsulates the packets.

It allows routers between GRE-based tunnel endpoints to adhere to precedence bits, thereby improving the routing of premium service packets. Now, Cisco IOS QoS technologies, such as policy routing, WFQ, and weighted random early detection (WRED), can operate on intermediate routers between GRE tunnel endpoints.

GRE tunnels are commonly used to provide dynamic routing resilience over IPSec. Normal IPSec configurations cannot transfer routing protocols, such as Enhanced Interior Gateway Routing Protocol (EIGRP) and OSPF, or non-IP traffic, such as Internetwork Packet Exchange (IPX) and AppleTalk.

QoS Pre-Classify Issues: IPSec Authentication Header (AH)

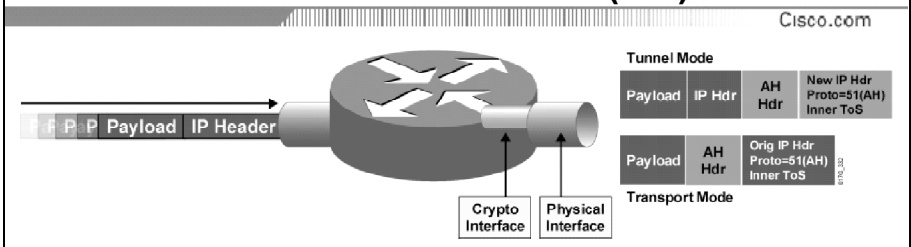

- **IPSec AH is for authentication only and does not perform encryption.**
- **With tunnel mode, the ToS byte value is copied automatically from the original IP header to the tunnel header.**
- **With transport mode, the original header is used and therefore the ToS byte is accessible.**

IPSec does not define the specific security algorithms to use, but rather it provides an open framework for implementing industry-standard algorithms.

Authentication Header (AH) provides strong integrity and authentication for IP datagrams using Secure Hash Algorithm (SHA) or MD5 hash algorithm. It also can provide non-repudiation. The Internet Assigned Numbers Authority (IANA) has assigned protocol number 51 to AH. Thus, in the presence of an AH header with both tunnel mode and transport mode, the IP header uses a value of 51 in the protocol field.

With tunnel mode, the ToS byte value is copied automatically from the original IP header to the tunnel header.

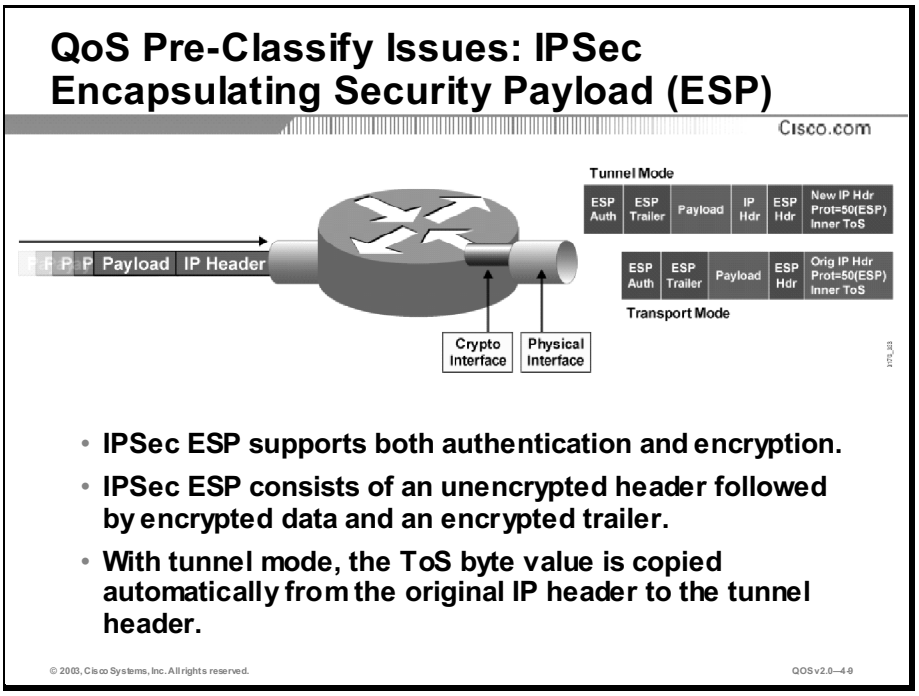

IPSec does not define the specific security algorithms to use, but rather it provides an open framework for implementing industry-standard algorithms.

Encapsulating Security Payload (ESP) consists of an unencrypted header followed by encrypted data and an encrypted trailer. ESP can provide both encryption and authentication.

As with AH, ESP supports SHA and MD5 hash algorithms for authentication. It supports Data Encryption Standard (DES) and 3DES as encryption protocols. The ESP header is at least 8 bytes. The IANA has assigned protocol number 50 to ESP. Thus, in the presence of (only) an ESP header with both tunnel mode and transport mode, the IP header uses a value of 50 in the protocol field.

With tunnel mode, the ToS byte value is copied automatically from the original IP header to the tunnel header.

QoS Pre-Classify Deployment Options

This topic describes situations where pre-classification is appropriate.

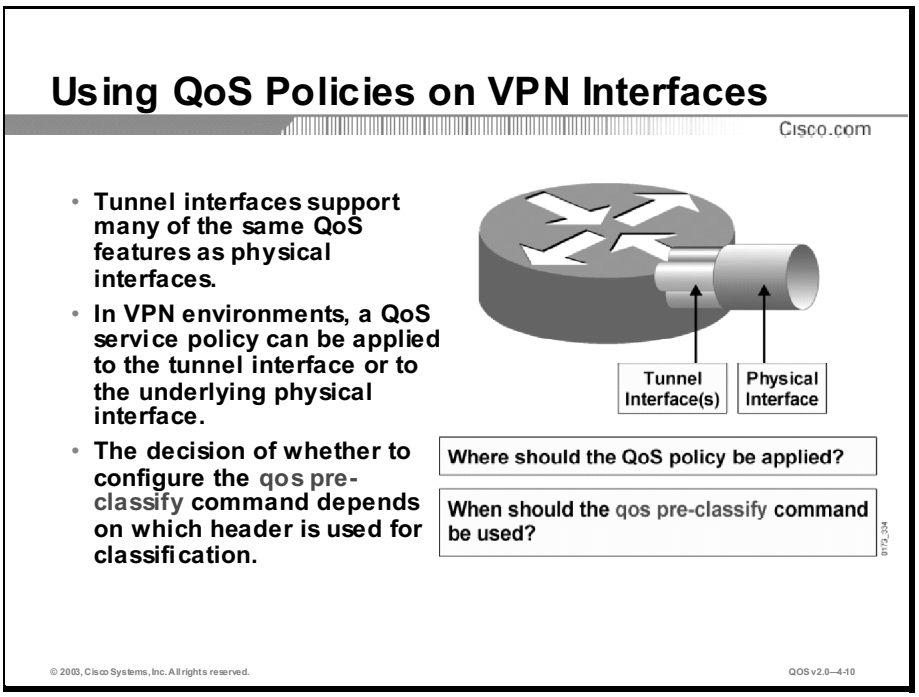

Classification defines the process of matching one or more fields in a packet header Layer 2, 3, or 4 and then placing that packet in a group or class of traffic. Using packet classification, you can partition network traffic into multiple priority levels or classes of service.

When configuring IPSec with GRE, the simplest classification approach is to match on IP precedence or DSCP values. Cisco IOS software release 11.3T introduced support for IPSec. Along with this support was the ToS byte preservation feature in which the router automatically copies the ToS header value from the original IP packet to the encapsulating IP header when using IPSec in tunnel mode.

ToS byte preservation also applies to AH. Also note that ESP in transport mode retains the original IP header and the original ToS value is transmitted even without ToS byte preservation. If packets arrive at the router without a set IP precedence or DSCP values, you can use class-based marking to re-mark the packet headers before encryption or encapsulation. When the packets reach the egress interface, the QoS output policy then can match and act on the re-marked values.

Alternately, you may need to classify traffic based on values other than IP precedence or DSCP. For example, you may need to classify packets based on IP flow or Layer 3 information, such as source and destination IP address. To do so, you must use the QoS for VPNs feature that you enable with the **qos pre-classify** command. This feature is available for Cisco 7100 series VPN routers and Cisco 7200 series routers (since 12.1(5)T) and for 2600 and 3600 series routers (since 12.2(2)T).

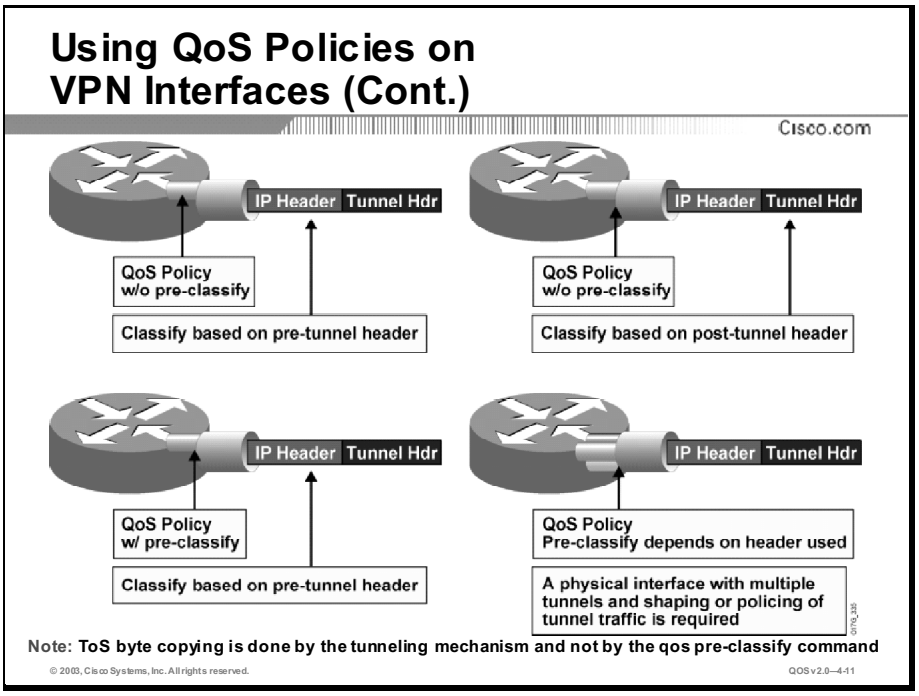

The **qos pre-classify** mechanism allows Cisco routers to make a copy of the inner IP header and to run a QoS classification before encryption, based on fields in the inner IP header. Without this feature the classification engine sees only a single encrypted and tunneled flow because all packets traversing across the same tunnel have the same tunnel header and thus, will receive the same treatment in the event of congestion.

If your classification policy matches on the ToS byte, you do not need to use the **qos pre-classify** command because the ToS value is copied to the outer header by default. In addition, you can create a simple QoS policy that sorts traffic into classes based on IP precedence. However, differentiating traffic within a class and separating it into multiple flow-based queues requires the **qos pre-classify** command.

Note:	ToS byte copying is done by the tunneling mechanism and not by the **qos pre-classify** command.

Configuring QoS Pre-Classify

This topic describes the Cisco IOS commands that are necessary to configure pre-classification.

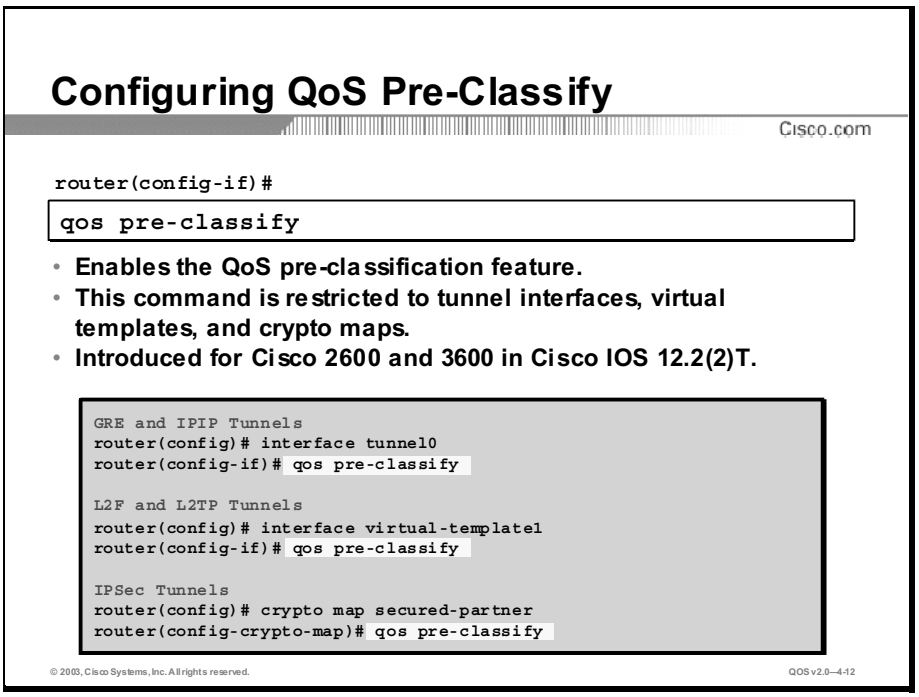

The **qos-pre-classify** Cisco IOS command enables the QoS pre-classification feature. The command can be applied to a tunnel interface, a virtual template interface, or a crypto map.

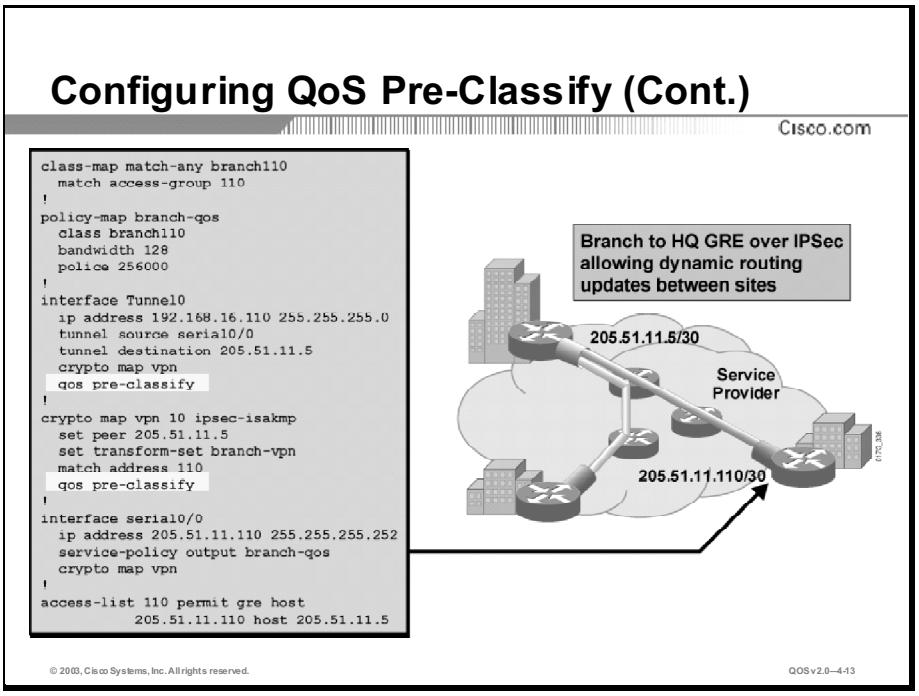

The figure shows the successful configuration of the **qos pre-classify** command.

The configuration of the branch router is shown as follows:

- On the S0/0 interface there is an outgoing service-policy that sets the bandwidth of the interface at 128 kbps and is policed at a rate of 256 kbps. This policy is applied to any match in the class map branch110.
- Also, a traffic tunnel has been built on interface S0/0 (whose destination is HQ for this branch ip address 205.51.11.5). It is on this traffic tunnel that qos pre-classification has been configured.

The example configuration also shows that qos pre-classify has been successfully enabled on the crypto map named vpn. This crypto map has also been applied to S0/0. If qos pre-classify is only enabled on the crypto map and not on the tunnel interface, the router will see one flow only, the GRE tunnel (protocol 47).

There are a few restrictions when configuring the QoS for VPNs feature:

- The QoS for VPNs feature can be enabled on IP packets only.
- If a packet is fragmented after encryption, only the first fragment is pre-classified. Subsequent fragments might receive different classifications. This behavior is consistent with QoS classification of non-tunneled fragments.
- Interfaces that run cascading QoS features, such as generic traffic shaping or custom queuing, are required to have QoS for VPNs enabled or disabled on all cascading features. If the QoS for VPN feature is enabled on one cascading feature, the QoS for VPN feature must be enabled on all cascading features. Similarly, if the QoS for VPN feature is disabled on one cascading feature, the QoS for VPN feature must be disabled on all cascading features.

- When configuring VPN QoS in conjunction with GRE or IPSec tunnel interfaces, the only congestion management (queuing) strategy that can be employed on the tunnel interface is FIFO because the device on the other end of the tunnel expects to receive packets in order. Any packet not arriving in order, because of queue management for example, will be discarded at the tunnel endpoint.

Monitoring QoS Pre-Classify

This topic describes the Cisco IOS commands that are necessary to monitor pre-classification.

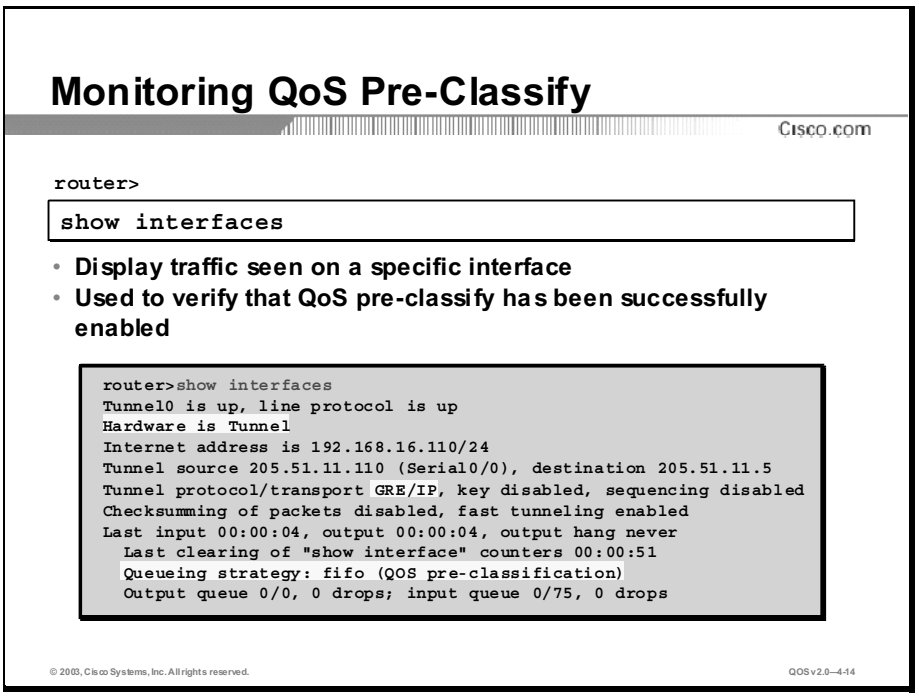

The **show interfaces** command is used to verify that the QoS for VPN feature has been enabled. Verified by examining the queuing strategy line in the above figure:

```
Queueing strategy: fifo (QOS pre-classification)
```


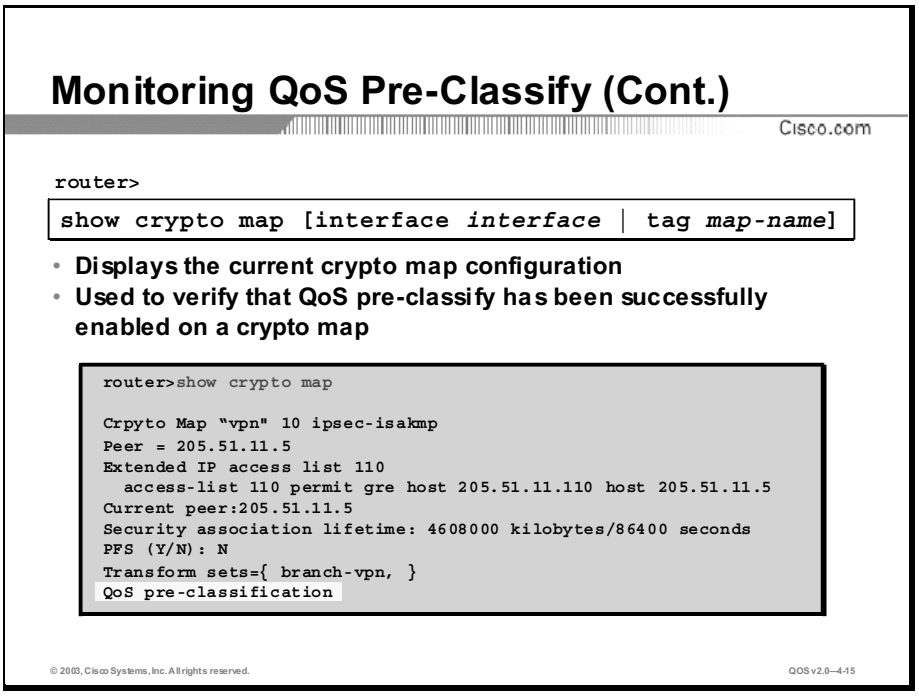

In the example the **show crypto map** command has been issued. This command will show current crypto map configuration and also whether the QoS for VPN feature has been successfully enabled on a crypto map.

show crypto map [**interface** *interface* | **tag** *map-name*]

Syntax Description

Parameter	Description
interface *interface*	(Optional) Displays only the crypto map set applied to the specified interface.
tag *map-name*	(Optional) Displays only the crypto map set with the specified *map-name*.

Summary

This topic summarizes the key points discussed in this lesson.

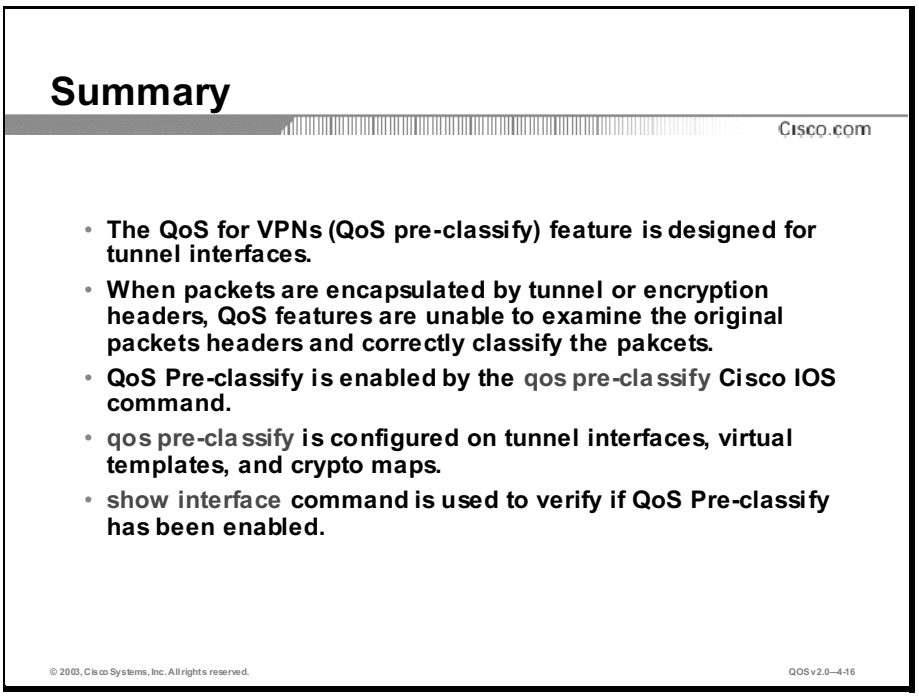

References

For additional information, refer to these resources:

- For more information on QoS Pre-Classify, refer to “Reference Guide to Implementing Crypto and QoS” at the following URL: http://www.cisco.com/warp/customer/105/crypto_qos.pdf
- For more information on QoS Pre-Classify, refer to “Configuring QoS for Virtual Private Networks” at the following URL: http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122cgcr/fqos_c/fqcprt1/qcfvpn.pdf

Next Steps

For the associated lab exercise, refer to the following section of the course Lab Guide:

- Lab Exercise 4-3: Configuring QoS Pre-Classify

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) When QoS pre-classify is enabled, the QoS features on the output interface are _____ before encryption?

A) marked

B) policed

C) shaped

D) classified

Q2) In GRE tunneling, by default the ToS field of the original packet header is copied to the _____ field of the ToS field for GRE tunnel header.

A) COS

B) DSCP

C) ToS

D) IP precedence

Q3) Which of the following Cisco IOS commands enables the QoS pre-classify feature?

A) **crytpo map vpn**

B) **qos pre-classify**

C) **qos-pre-classify**

D) **crytpo map vpn ipsec-isakmp**

Q4) At what configuration level is the QoS pre-classify feature enabled?

A) global

B) router

C) interface

D) privileged

Q5) Which of the following Cisco IOS commands is used to verify whether QoS pre-classify is enabled?

A) **qos pre-classify**

B) **show interface**

C) **show crypto map**

D) **crytpo map vpn ipsec-isakmp**

Quiz Answer Key

Q1) D

Relates to: Implementing QoS with Pre-Classification

Q2) C

Relates to: QoS Pre-Classify Applications

Q3) B

Relates to: QoS Pre-Classify Deployment Options

Q4) C

Relates to: Configuring QoS Pre-Classify

Q5) B

Relates to: Monitoring QoS Pre-Classify

Configuring QoS Policy Propagation Through BGP

Overview

The QoS Policy Propagation in Border Gateway Protocol (BGP [QPPB]) feature allows you to classify packets based on access lists, BGP community lists, and BGP autonomous system (AS) paths. The supported classification policies include IP precedence setting and the ability to tag the packet with a QoS class identifier internal to the router. After a packet has been classified, you can use other QoS features such as CAR and WRED to specify and enforce business policies to fit your business model.

This lesson describes the QPPB classification mechanism. QPPB features covered in this lesson include a review of CEF and the tasks and Cisco IOS commands that are required to configure QPPB on Cisco routers.

Relevance

Classification is a fundamental requirement for any network deployment of QoS. As such, it is of key importance to understand the different ways that traffic can be classified.

Objectives

Upon completing this lesson, you will be able to explain how to implement classification and marking in an inter-domain network using QPPB. This includes being able to meet these objectives:

- Describe the QPPB mechanism
- Describe the interaction between IP QoS and the Border Gateway Protocol
- Describe the operation of CEF
- List the steps required to configure QPPB on Cisco routers
- Identify the Cisco IOS commands required to configure QPPB on Cisco routers

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Basic knowledge of internetworking with TCP/IP concepts
- Basic knowledge of the Cisco IOS command-line interface

Outline

The outline lists the topics included in this lesson.

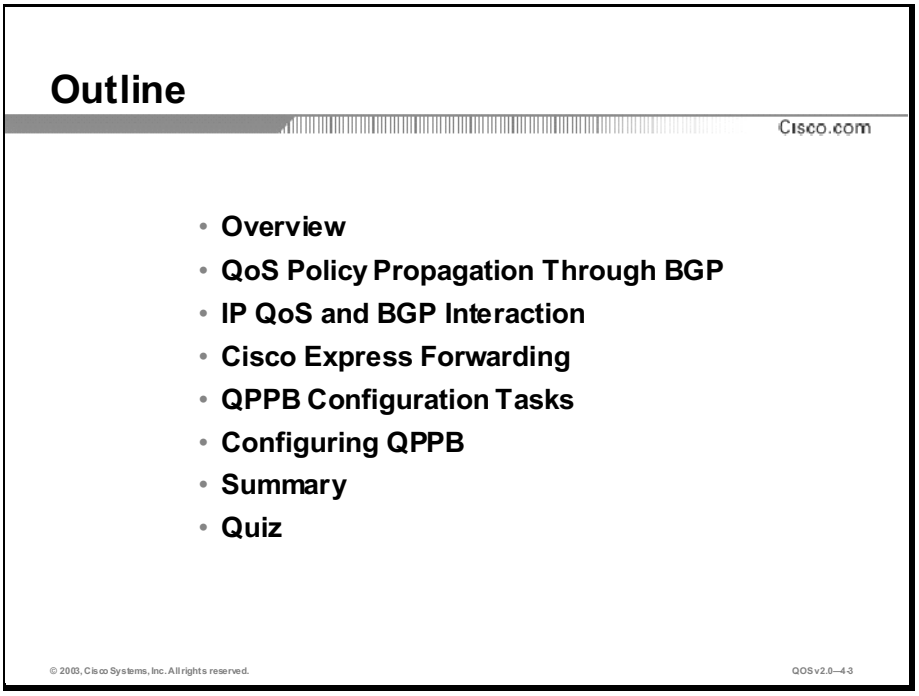

QoS Policy Propagation Through BGP

This topic describes the QPPB feature, which propagates QoS policy via BGP.

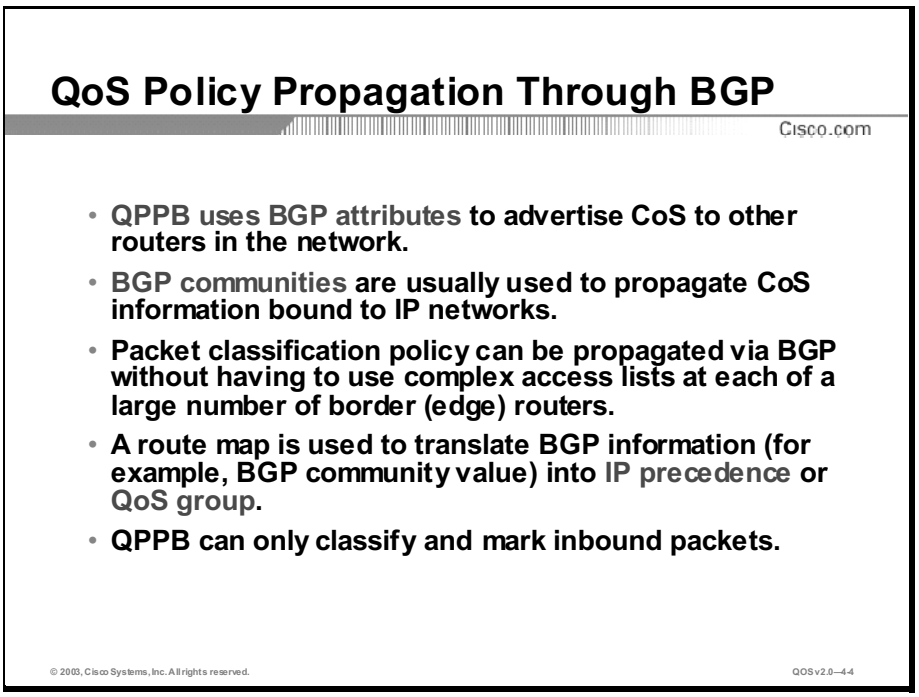

BGP is an inter-domain routing protocol that exchanges reachability information with other BGP systems. The QoS policy propagation via the BGP feature allows classifying packets based on access lists, BGP community lists, and BGP AS paths.

The supported classification policies include IP precedence setting and the ability to tag the packet with a QoS class identifier internal to the router. After a packet has been classified, one can use other QoS features such as policing, WRED, and traffic shaping to specify and enforce business policies to fit the business model.

The QoS policy propagation via BGP feature has the following enhancements:

- **QoS group ID:** You can set an internal QoS group ID that can be used later to perform policing or WFQ, based on the QoS group ID.
- **Source and destination address lookup:** You can specify whether the IP precedence level or QoS group ID used is obtained from the source (input) address or destination (output) address entry in the route table.

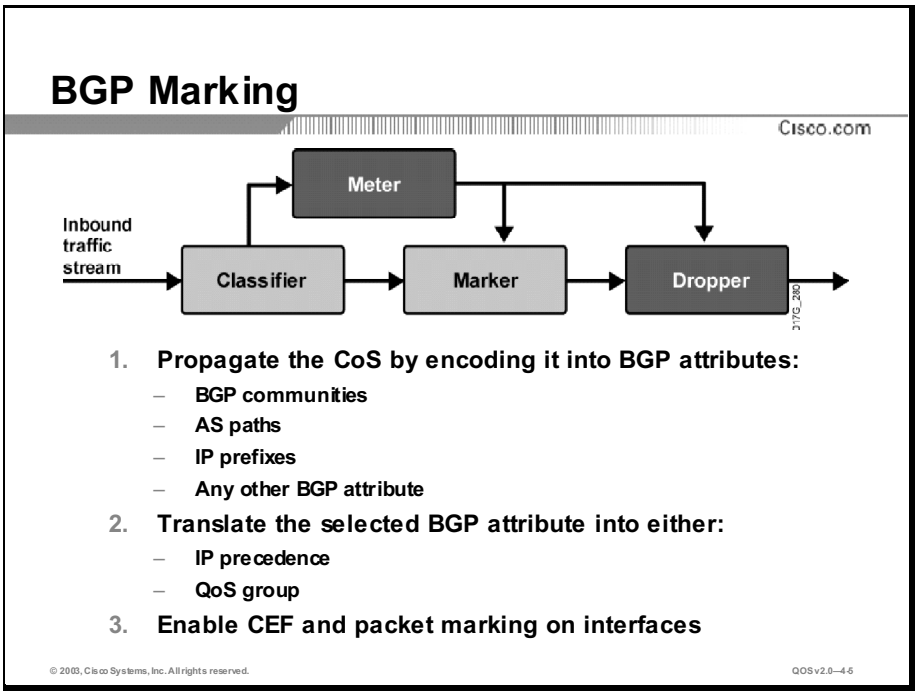

BGP propagates the CoS by encoding it into the following BGP attributes:

- BGP communities attribute
- AS path attribute
- IP prefix attribute
- Or any other BGP attribute

BGP can translate the selected BGP attribute into either:

- IP precedence
- QoS group

The QPPB feature requires that CEF and packet marking is enabled on interfaces.

IP QoS and BGP Interaction

This topic describes the interaction between QoS and BGP.

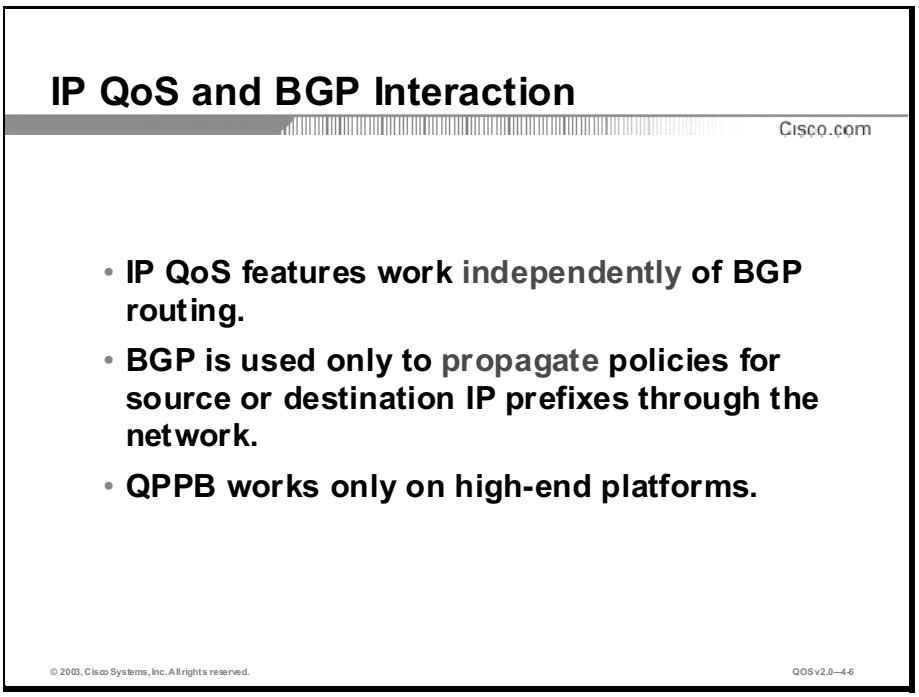

When using QPPB, the IP QoS feature works independently from BGP routing. BGP is only used to propagate the QoS policy.

In QBBP configurations, the network administrator specifies whether to use IP precedence or the QoS group ID obtained from the source (input) address or destination (output) address entry in the route table.

You can specify either the input or output address.

QPPB works only on high-end routers:

- Cisco 7200 series
- Cisco 7500 series
- Cisco 7000 series with the RSP 7000 and RSP 7000CI
- Cisco 10000 series
- Cisco 12000 series

Cisco Express Forwarding

This topic presents a review of CEF switching on Cisco IOS platforms.

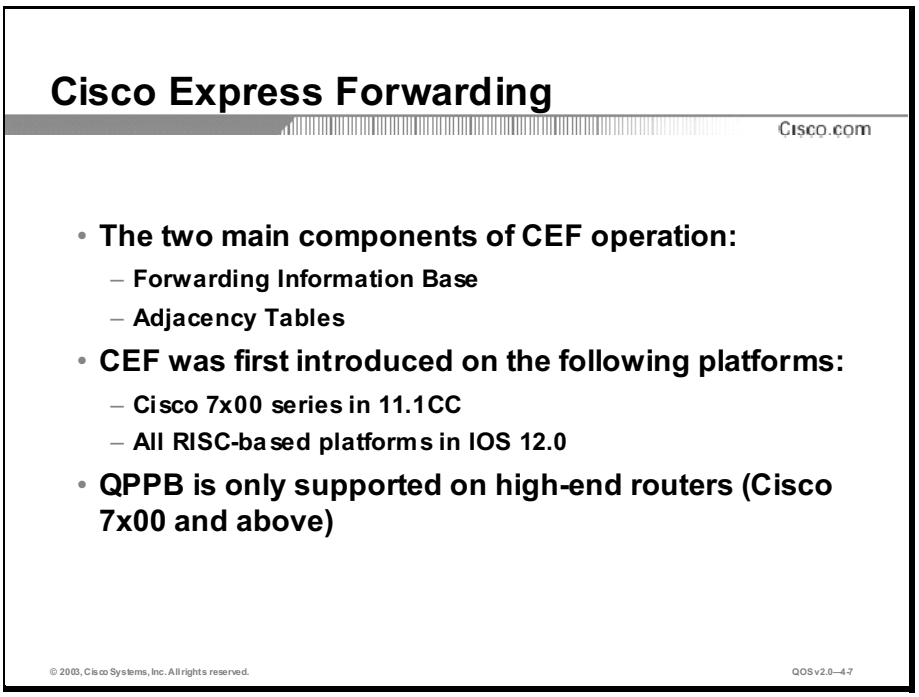

When the router is initialized for CEF, two main tables are built inside the router:

- The Forwarding Information Base (FIB), which lists all paths to all reachable networks, together with the output interface information
- The adjacency table, which lists all required next-hops on output interfaces

To enable scalable forwarding, CEF builds a forwarding table called the FIB. Contrary to demand-switching methods, the FIB is not a small subset of the routing table. The FIB is a full extract of the routing table, with all the forwarding parameters precalculated at the time of FIB creation, and updated with any topology (routing table) changes.

The second table is the adjacency table. This table contains all the Layer 2 next-hops, which are currently being used by the router to forward traffic.

The two tables are interconnected, so that every destination network is linked to its appropriate local next-hop adjacency. Many destinations can be linked to the same next-hop adjacency, removing redundancy and increasing manageability of CEF tables. Moreover, a single destination can point to multiple next-hop adjacencies, enabling flexible traffic load balancing.

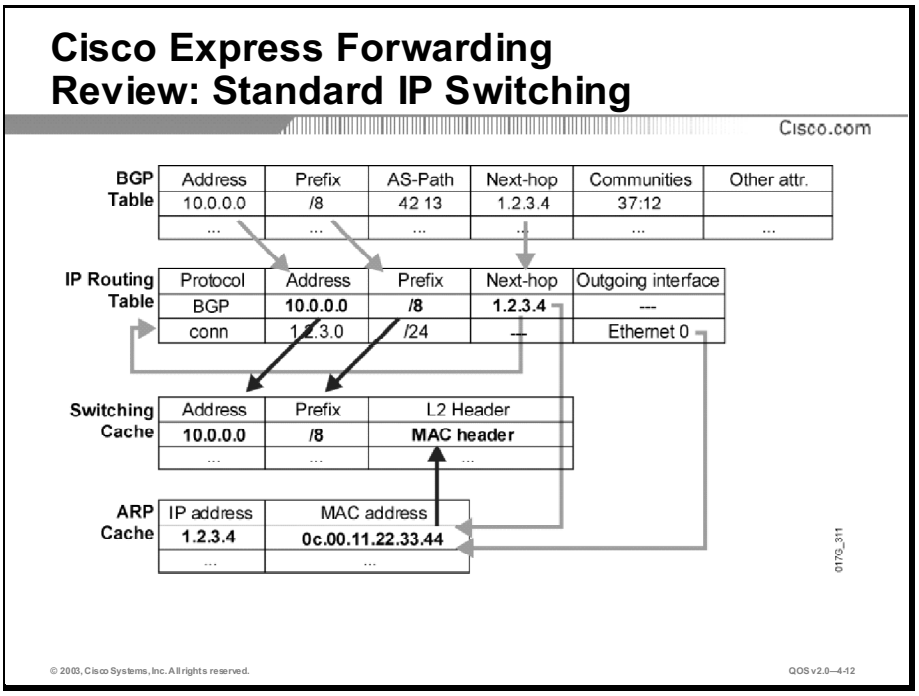

The figure shows a sequence of events when process switching and fast switching are used for destinations learned through BGP.

- When a BGP update is received and processed, an entry is created in the routing table.
- When the first packet arrives for this destination, the router tries to find the destination in the fast-switching cache. Because it is not there, process switching has to switch the packet when the process is run. The process performs a recursive lookup to find the outgoing interface. It may possibly trigger an ARP request or find the Layer 2 address in the Address Resolution Protocol (ARP) cache. Finally, it creates an entry in the fast-switching cache.
- All subsequent packets for the same destination are fast-switched:

The switching occurs in the interrupt code (the packet is processed immediately).

Fast destination lookup is performed (no recursion).

The encapsulation uses a pregenerated Layer 2 header that contains the destination as well as Layer 2 source (MAC) address (no ARP request or ARP cache lookup is necessary).

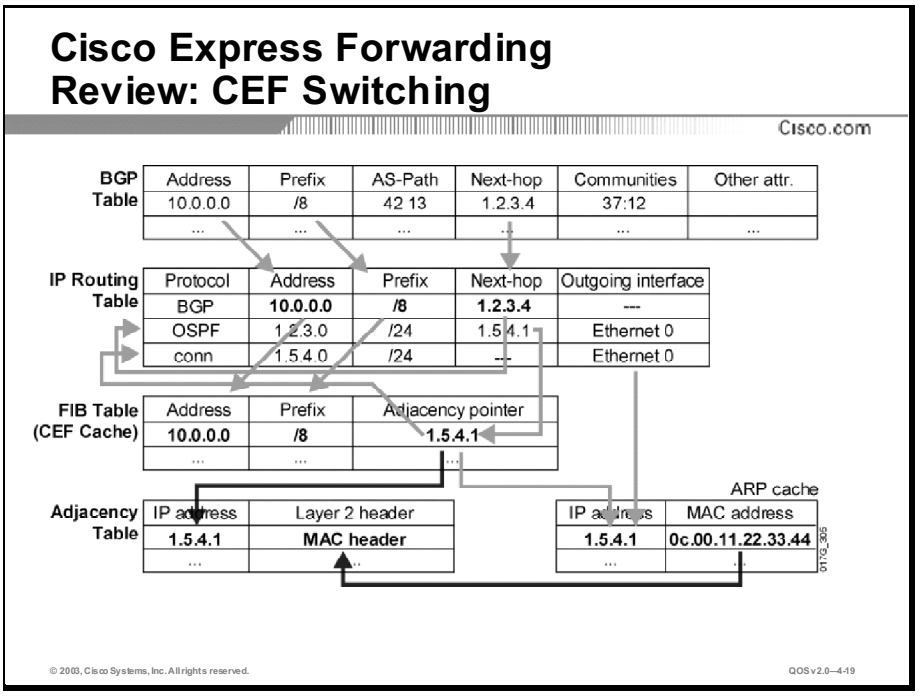

The generation of entries in the FIB table is not packet-triggered but change-triggered. When something changes in the IP routing table, the change is also reflected in the FIB table.

As the FIB contains the complete IP switching table, the router can make definitive decisions based on the FIB. Whenever a router receives a packet that should be CEF-switched, but the destination is not in the FIB, the packet is dropped.

The FIB table is also different from other fast-switching caches in that it does not contain information about the outgoing interface and the corresponding Layer 2 header. That information is stored in a separate table—the adjacency table. This table is more or less a copy of the ARP cache, but instead of holding only the destination MAC address, it holds the Layer 2 header (source and destination MAC address).

The figure illustrates how the CEF switching entries are built. When a route is added or changed in the main routing table (for example, learned via BGP), a new FIB entry is created, and the next hop is calculated via recursive lookups to the routing table (if necessary). The FIB entry is then linked to the next-hop adjacency entry, which provides the necessary Layer 2 information used to forward the packet on the output medium.

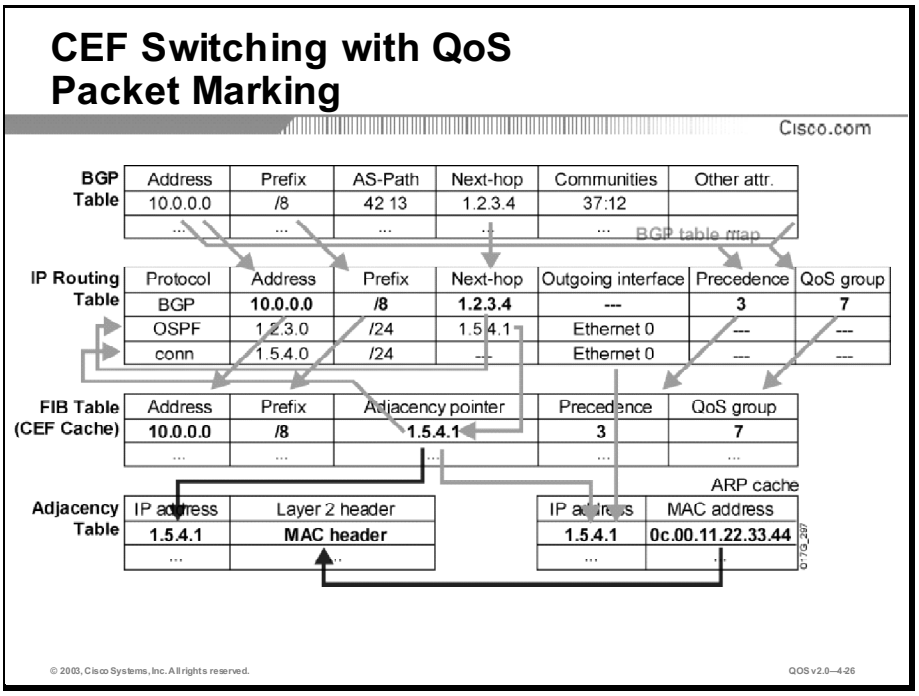

In the figure above, the tables from the previous page are displayed, with the difference that BGP communities being translated to IP precedence and QoS group are inserted into the FIB table also.

QPPB Configuration Tasks

This topic lists the configuration tasks required to enable the QPPB feature.

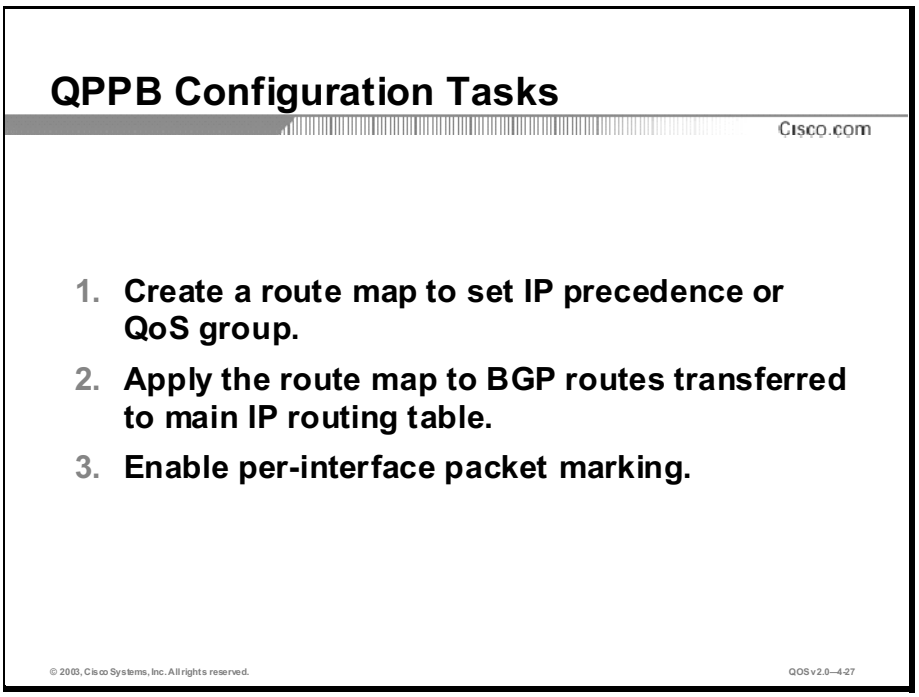

The tasks required to enable QPPB are as follows:

- Create a route map(s) to set IP precedence or QoS group. The **route-map** command is used to accomplish this task as follows:

```
route-map <route-map name> permit 10
 match community <community-list>
 set ip precedence <ip precedence value>
 set ip qos-group <qos-group #>
```

- Apply the route map to BGP routes that are in the BGP table. The **table-map** command is used to accomplish this task as follows:

```
router bgp <as #>
 table-map <route-map name>
```

- Enable the required interface(s) for packet marking. The **bgp-policy** command is used to accomplish this task as follows:

```
interface X
 bgp-policy <source | destination> ip-prec-map
```


Configuring QPPB

This topic identifies the Cisco IOS commands that are required to configure QPPB.

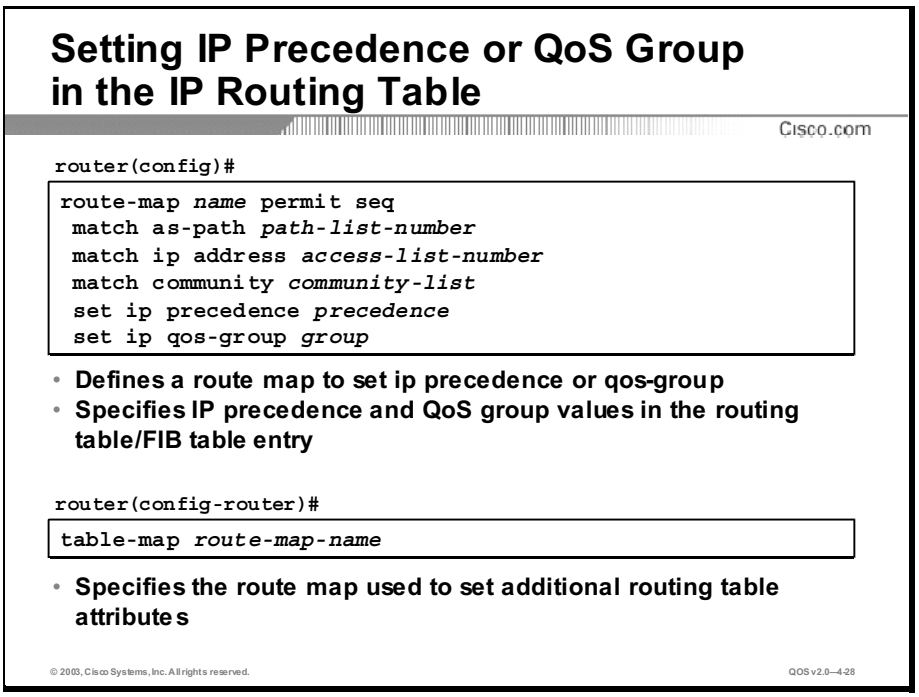

Use the **route-map** command to define a route map to match based on a bgp community list, bgp as-path, or access-list and to set the ip precedence or qos-group. To set the precedence value (and an optional IP number or IP name) in the IP header, use the **set ip precedence** route-map configuration command. To leave the precedence value unchanged, use the **no** form of this command.

set ip precedence [*precedence* | *name*]

Syntax Description

Parameter	Description
precedence | *name*	A number or name that sets the precedence bits in the IP header. The values for the *precedence* argument and the corresponding *name* argument are listed in the following table from least to most important.

The following table lists the values for the *precedence* argument and the corresponding *name* argument for precedence values in the IP header. They are listed from least to most important.

Precedence	Name
0	routine
1	priority
2	immediate
3	flash
4	flash-override
5	critical
6	internet
7	network

To set a group ID that can be used later to classify packets, use the **set qos-group** QoS policy map configuration command. To remove the group ID, use the **no** form of this command.

set qos-group group-id

Syntax Description

Parameter	Description
group-id	Group ID number in the range from 0 to 99.

Note: To display QoS group information, use the **show ip cef** command.

Use the bgp **table-map** command to apply the route map to the BGP routing process. This will populate the corresponding BGP routes in the IP routing table and Forwarding Information Base (FIB) with the CoS (IP precedence and/or qos-group) information. To modify metric and tag values when the IP routing table is updated with BGP learned routes, use the **table-map** command in address family or router configuration mode. To disable this function, use the **no** form of the command.

table-map map-name

Syntax Description

Parameter	Description
map-name	Route map name, from the **route-map** command.

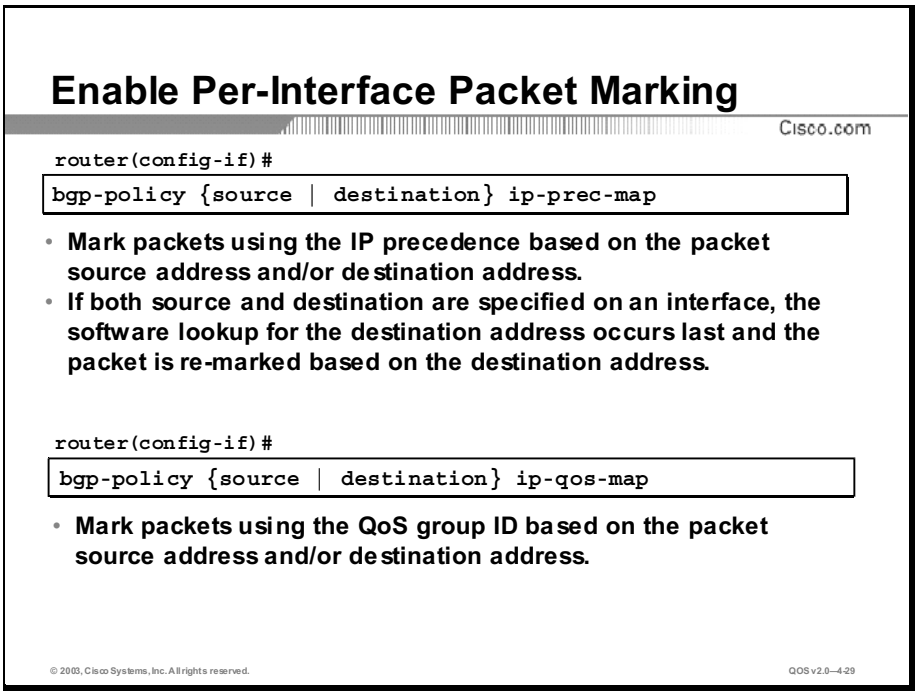

After the IP routing table and the FIB table contain the CoS information (IP precedence or QoS group), CEF-based markings can be configured on the input interfaces by using the **bgp-policy** interface configuration command.

Using the **bgp-policy** interface configuration command, CEF-based markings can be performed based on the source or destination address of an incoming packet. Use the source option to mark packets sourced from a customer. Use the destination option to mark packets destined to a customer.

The packets can be marked with the ip precedence or qos-group value from the FIB table. Use the **ip-prec-map** option to mark the packets with ip precedence and use the **ip-qos-map** option to mark the packets with qos-group.

bgp-policy {source | destination} {ip-prec-map | ip-qos-map}

Syntax Description

Parameter	Description
source	The IP precedence bit or QoS group ID from the source address entry in the route table
destination	The IP precedence bit or QoS group ID from the destination address entry in the route table
ip-prec-map	QoS policy based on the IP precedence
ip-qos-map	The QoS policy based on the QoS group ID

Note: If you specify both **source** and **destination** on the interface, the software looks up the source address in the routing table and classifies the packet based on the source address first; then the software looks up the destination address in the routing table and reclassifies it based on the destination address.

Example: Configuration

```
router bgp 30
table-map precedence-map
neighbor 20.20.20.1 remote-as 10
neighbor 20.20.20.1 send-community
!
ip bgp-community new-format
!
! Match community 1 and set the IP precedence to priority and
set the QoS group to 1
route-map precedence-map permit 10
match community 1
set ip precedence priority
set ip qos-group 1
!
! Match community 2 and set the IP precedence to immediate
route-map precedence-map permit 20
match community 2
set ip precedence immediate
!
ip community-list 1 permit 60:1
ip community-list 2 permit 60:2
!

interface HSSI 5/0
 no ip address
 encapsulation frame-relay
!
interface HSSI 5/0/0.1 point-to-point
 ip address 20.20.20.1 255.255.255.0
 bgp-policy source ip-prec-map
 no ip mroute-cache
 frame-relay interface-dlci 20 IETF
```


In this example the community attribute is being matched and then the action is taken on those attributes. If the community is 60:1 (ip community list 1) its IP precedence will be set to priority and the qos group will be set to 1as specified in the route map precedence map.

If the community attribute is 60:2 (ip community list 2) its IP precedence will be set to immediate.

The policy is then applied to the interface HSSI 5/0/0/0.1, using the **bgp-policy source** command. The **ip-prec-map** keyword indicates that the QoS policy is based on IP precedence.

Example: Configuring QPPB

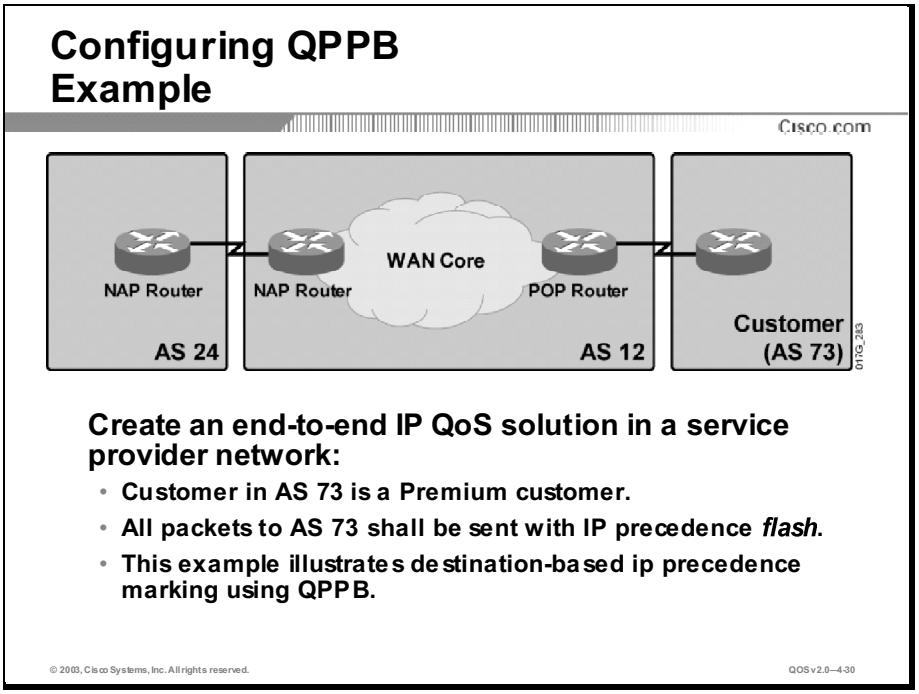

The figure shows an example of configuring QPPB.

In a service provider network an end-to-end IP QoS solution must be created.

The requirements are:

- Customer in AS 73 is a *Premium* customer
- All packets to AS 73 shall be sent with IP precedence flash

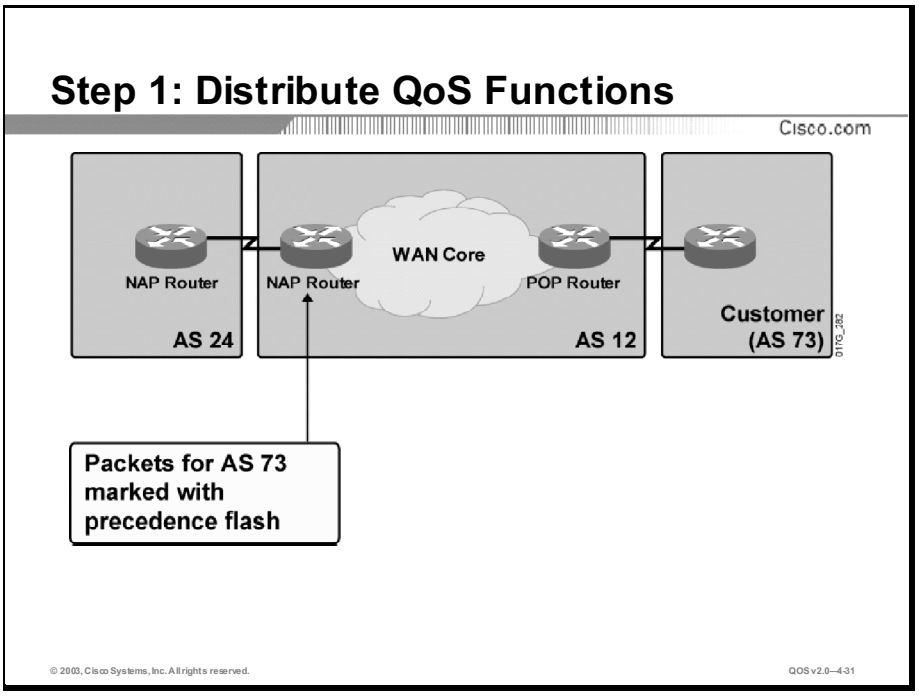

Because we are going to create an end-to-end QoS solution, the figure shows the first step requirements:

- Routes that are received from AS 24 and destined for AS 73 will have IP precedence set to *flash* on the **NAP router** (in AS 12).

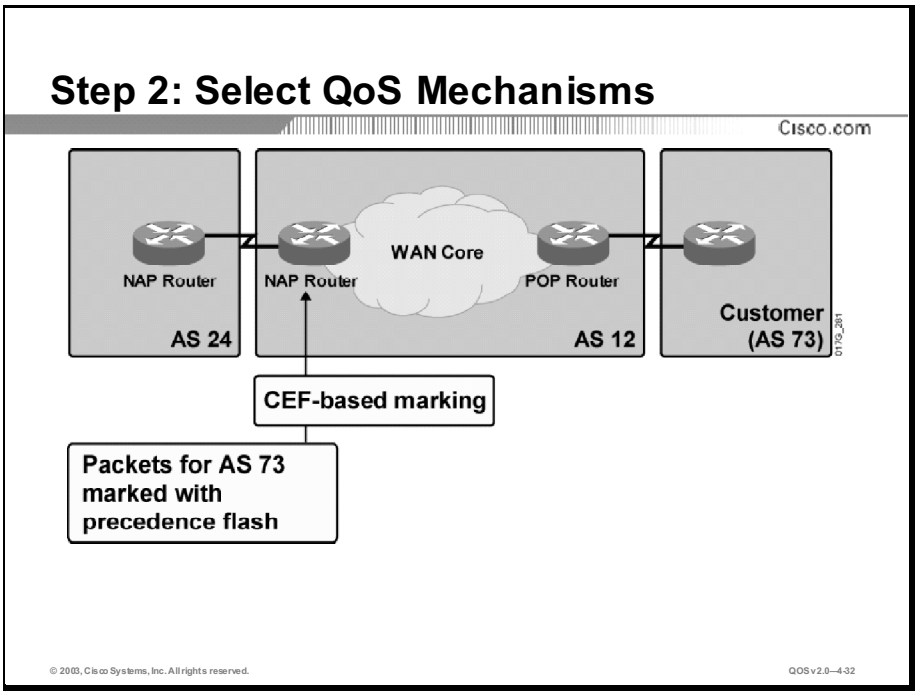

This figure shows the second requirement:

- Enable CEF-based marking on the NAP router serial interface connecting to AS 24.

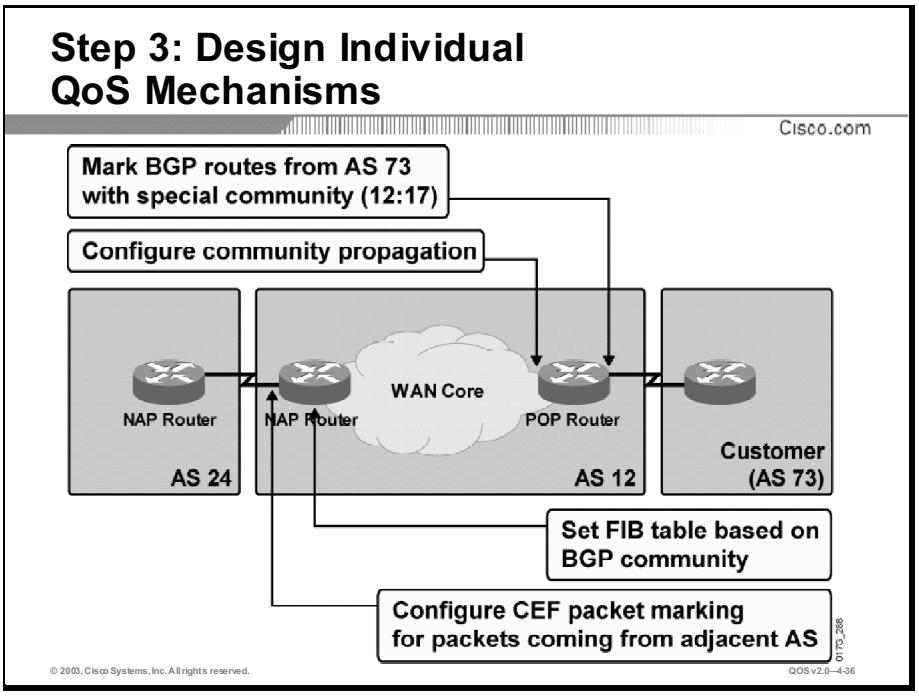

This figure shows the third step requirement:

- BGP routes that are received into AS 12 from AS 73 will be marked with a community value of 12:17 on the points of presence (POP) router.
- Community propagation will have to be configured on the POP router so that the community value of 12:17, set on the POP router, will be propagated to the NAP router.
- All the BGP routes with a community of 12:17 in the IP routing table and the FIB table on the AS 12 NAP routers will contain the IP precedence flash.

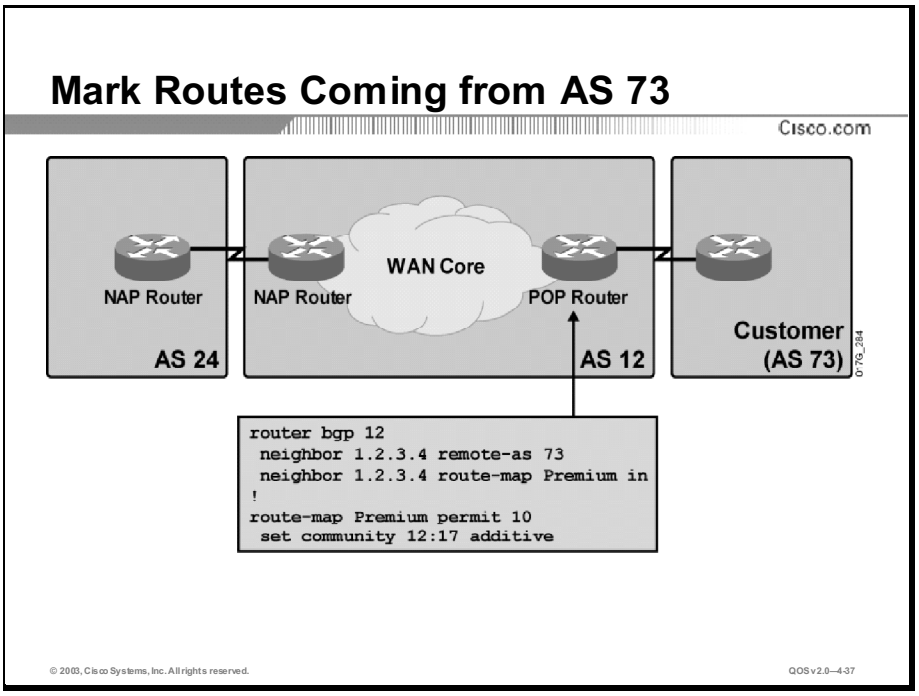

The figure shows the configuration that is necessary to meet the requirement that routes coming from AS 73 will be marked with the special community value of 12:17.

Configuration on the POP router to mark BGP routes from AS 73 with the community value 12:17:

```
router bgp 12
 neighbor 1.2.3.4 remote-as 73
 neighbor 1.2.3.4 route-map Premium in
!
route-map Premium permit 10
 set community 12:17 additive
```


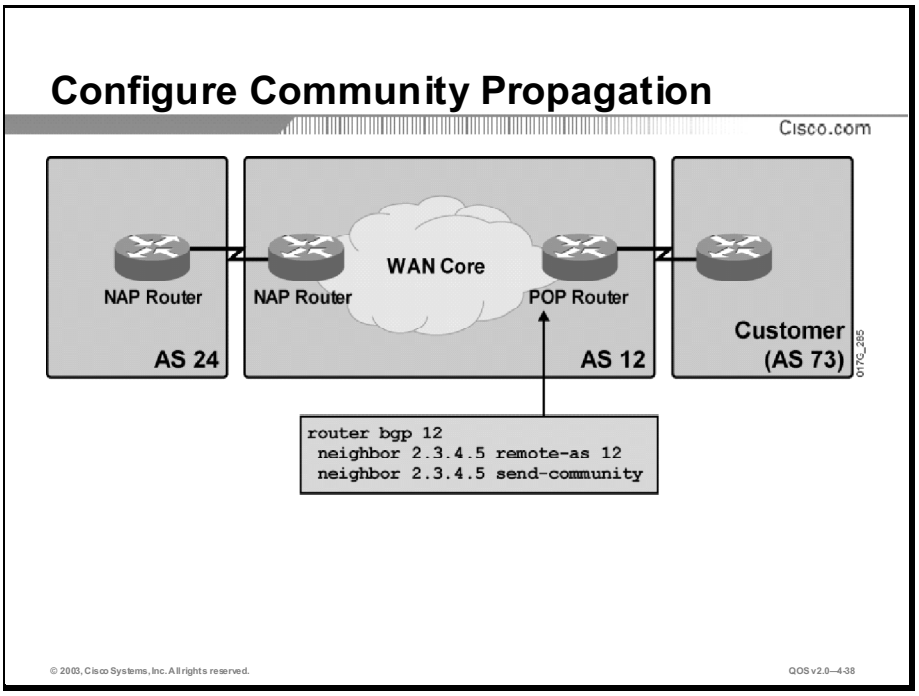

The figure shows the configuration that is necessary to propagate the special community value, 12:17, which has been added on the POP router and will be seen on the NAP router.

Configuration on the POP router to mark to propagate the community value (12:17) to the NAP router:

```
router bgp 12
neighbor 2.3.4.5 remote-as 12
neighbor 2.3.4.5 send-community
```

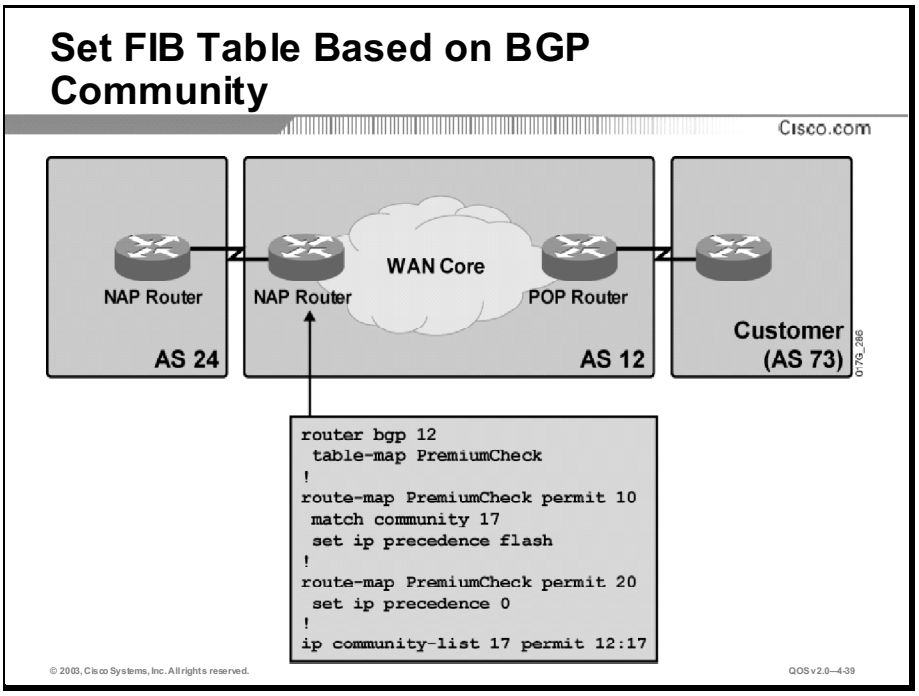

The NAP router in AS 12 uses a route map to translate BGP community values into appropriate IP precedence values. The figure illustrates how all BGP routes carrying BGP community 12:17 are tagged with IP precedence flash in the routing table and the FIB table. All other BGP routes are tagged with IP precedence 0.

Configuration on the **NAP** router to set/change the IP precedence of those BGP routes that match the community value (12:17):

```
router bgp 12
 table-map PremiumCheck
!
route-map PremiumCheck permit 10
 match community 17
 set ip precedence flash
!
route-map PremiumCheck permit 20
 set ip precedence 0
!
ip community-list 17 permit 12:17
```

The configuration shows that if the route map PremiumCheck matches the community attribute of 12:17, the corresponding packet will have its IP precedence changed to flash, as is required in the example.

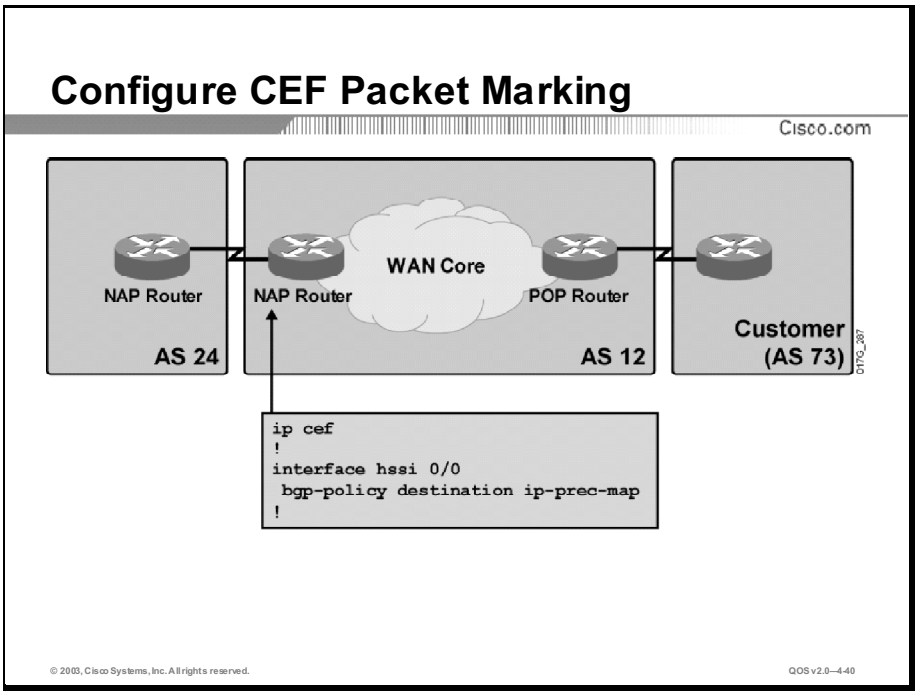

The last configuration step is to enable CEF-based marking on NAP router in AS 12. This example requires that all packets going to (destination-based marking) the customer (AS 73) network be marked with IP precedence flash.

In this case, the AS 12 NAP router HSSI 0/0 interface connects to AS 24. Therefore, the **bgp-policy destination ip-prec-map** command is configured under the HSSI 0/0 interface to enable destination-based CEF-based marking. All packets from AS 24 destined to the customer AS 73 will be marked with IP precedence flash.

QPPB marking is only available in combination with CEF switching. The global **ip cef** command enables CEF switching on all interfaces that support CEF.

Configuration on the **NAP** router to configure CEF packet marking:

```
ip cef
!
interface hssi 0/0
 bgp-policy destination ip-prec-map
```

Summary

This topic summarizes the key points discussed in this lesson.

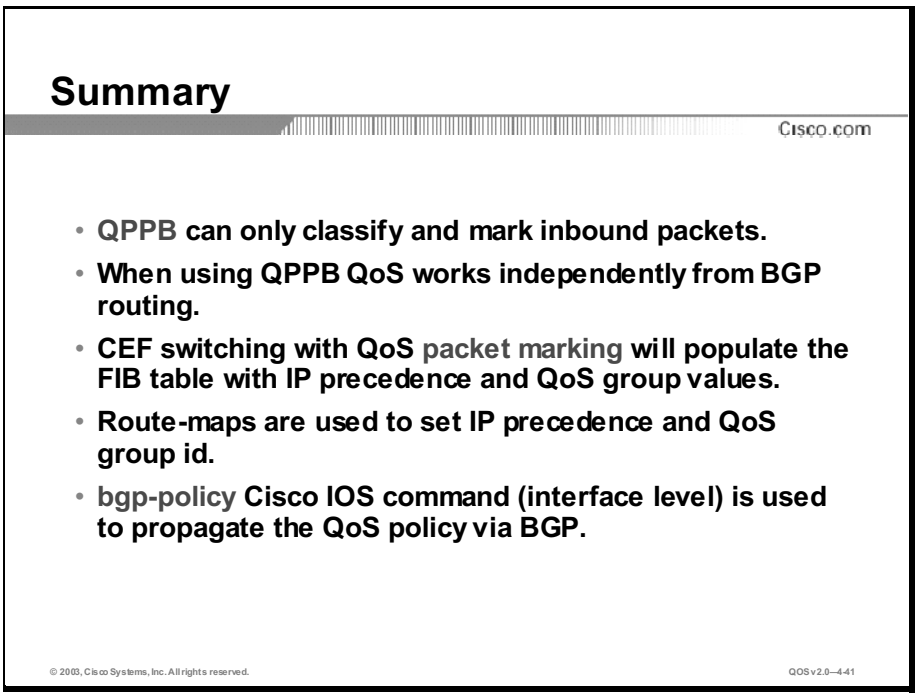

References

For additional information, refer to these resources:

- For more information on QPPB, refer to “Classification Overview” at the following URL: http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122cgcr/fqos_c/fqcprt1/qcfclass.pdf
- For more information on QPPB, refer to “Quality of Service Policy Propagation via Border Gateway Protocol” at the following URL: http://www.cisco.com/univercd/cc/td/doc/product/software/ios111/cc111/bgpprop.pdf
- For more information on configuring QPPB, refer to “Configuring QoS Policy Propagation via Border Gateway Protocol” at the following URL: http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122cgcr/fqos_c/fqcprt1/qcfprop.pdf

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) BGP translates the selected attribute into which two of the following: (Choose two.)

A) BGP communities

B) DSCP value

C) IP precedence

D) QoS group

Q2) When using QPPB, BGP is used only to propagate QoS policies for which two of the following? (Choose two.)

A) COS value

B) DSCP value

C) Source IP Prefix

D) Destination IP Prefix

Q3) The FIB table (CEF cache) is different from other fast-switching caches in that it does not contain information about the outgoing interface and _____.

A) IP address

B) adjacency pointer

C) corresponding L2 header

D) corresponding L3 header

Q4) Which of the following is not a QPPB configuration task?

A) disable CEF

B) enable per-interface packet marking

C) create a route map to set IP precedence or QoS group

D) apply route map to BGP routes transferred to IP routing table

Q5) Which of the following commands enable per-interface packet marking?

A) **bgp-policy**

B) **set qos group**

C) **set ip precedence**

D) **set packet-marking**

Quiz Answer Key

Q1) C, D

Relates to: QoS Policy Propagation Through BGP

Q2) C, D

Relates to: IP QoS and BGP Interaction

Q3) C

Relates to: Cisco Express Forwarding

Q4) A

Relates to: QPPB Configuration Tasks

Q5) A

Relates to: Configuring QPPB

Configuring LAN Classification and Marking

Overview

A switch may be the fastest switch in the world, but if you have many inputs to the switch and fewer outputs or have larger input pipes than output pipes, the switch will experience congestion. At times of congestion, if the congestion management features are not in place, packets will be dropped. When packets are dropped, retransmissions occur. When retransmissions occur, the network load can increase. In networks that are already congested, this can add to existing performance issues and potentially further degrade performance. With converging networks, congestion management is even more critical. Latency-sensitive traffic such as voice and video can be severely impacted if delays are incurred. Simply adding more buffers to a switch will not necessarily alleviate congestion problems. Latency-sensitive traffic must to be switched as fast as possible. First, you need to identify this important traffic through classification techniques, and then implement buffer management techniques to avoid the higher priority traffic from being dropped during congestion.

This lesson will introduce the learner to classification and marking as it is implemented on Cisco Catalyst switches. Topics covered include LAN classification and marking options and platforms, and configuring and monitoring LAN-based classification and marking.

Relevance

Classification is a fundamental requirement for any network deployment of QoS. As such, it is of major importance in the converged networks of today.

Objectives

Upon completing this lesson, you will be able to describe LAN-based methods for implementing classification and marking. This includes being able to meet these objectives:

- Describe LAN-based classification and marking using a Layer 2 Catalyst workgroup switch
- Describe QoS trust boundaries and their significance in LAN-based classification and marking
- Identify the different classification and marking options available on Cisco L2 and L3 switching platforms
- Identify the Cisco IOS commands required to configure LAN-based classification and marking
- Identify the Cisco IOS commands required to monitor LAN-based classification and marking

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Basic knowledge of internetworking with TCP/IP concepts
- Basic knowledge of the Cisco IOS command-line interface

Outline

The outline lists the topics included in this lesson.

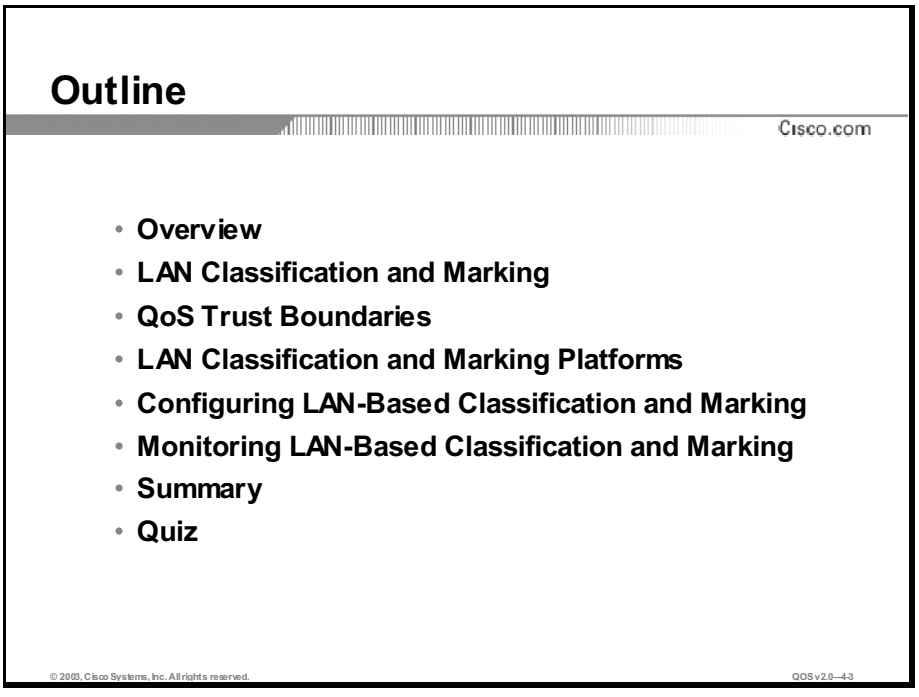

LAN Classification and Marking

This topic provides an introduction to QoS classification and marking in a LAN environment.

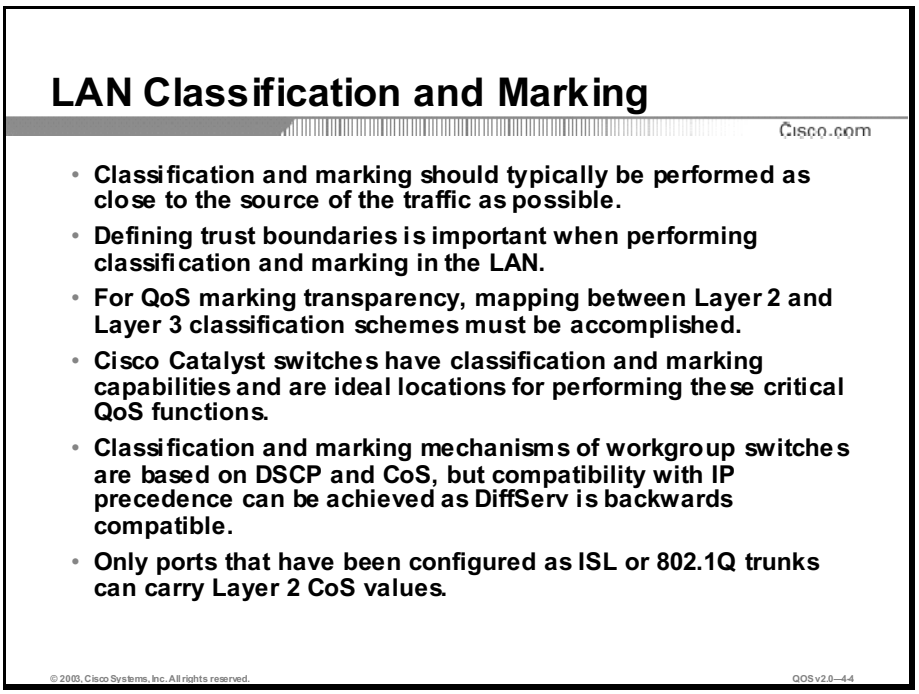

In the Catalyst line of multilayer switches is the capability to provide QoS at Layer 2 or Layer 3, depending on the switch type. At Layer 2, the frame uses CoS in 802.1p and Inter-Switch Link (ISL). CoS uses 3 bits, just like IP precedence, and maps well from Layer 2 to Layer 3, and vice versa.

The switches have the capability to differentiate frames based on CoS settings. If multiple queues are present, frames can be placed in different queues and serviced via weighted round robin (WRR). This allows each queue to have different service levels.

Classification is only performed on a Catalyst switch if QoS has been globally enabled on the switch.

QoS Trust Boundaries

This topic describes QoS trust boundaries and their significance in LAN-based classification and marking.

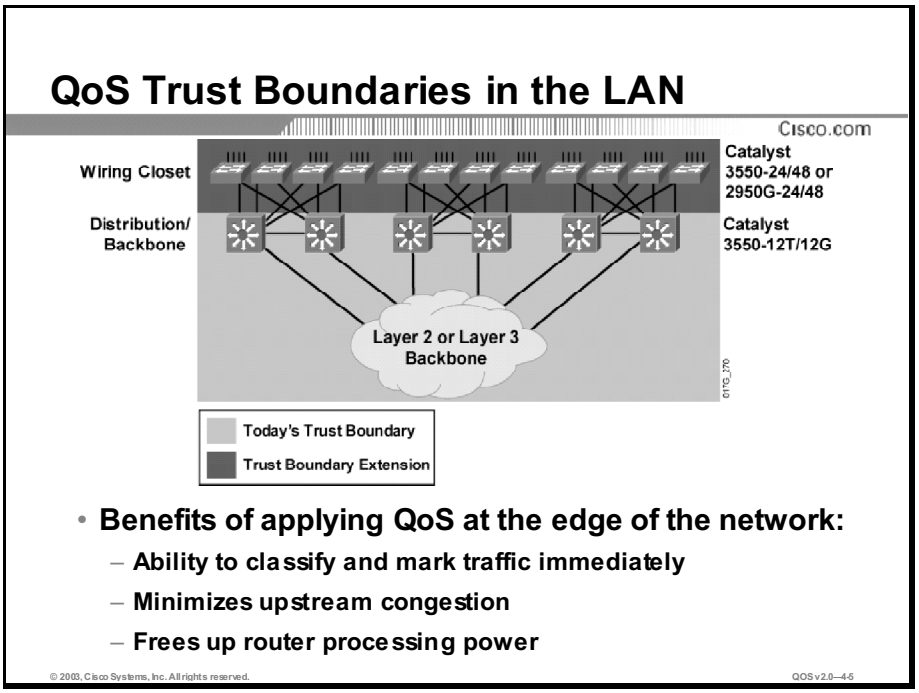

It is recommended that QoS be applied as close to the source of the traffic as possible.

Some of the benefits of applying QoS at the edge (or close to the source of the traffic) are as follows:

- Ability to classify and mark traffic immediately. This will reduce the upstream devices CPU utilization of the upstream device, thus reducing the possibility that priority traffic, such as voice, would be delayed at some point further in the network.
- Frees up router processing power.

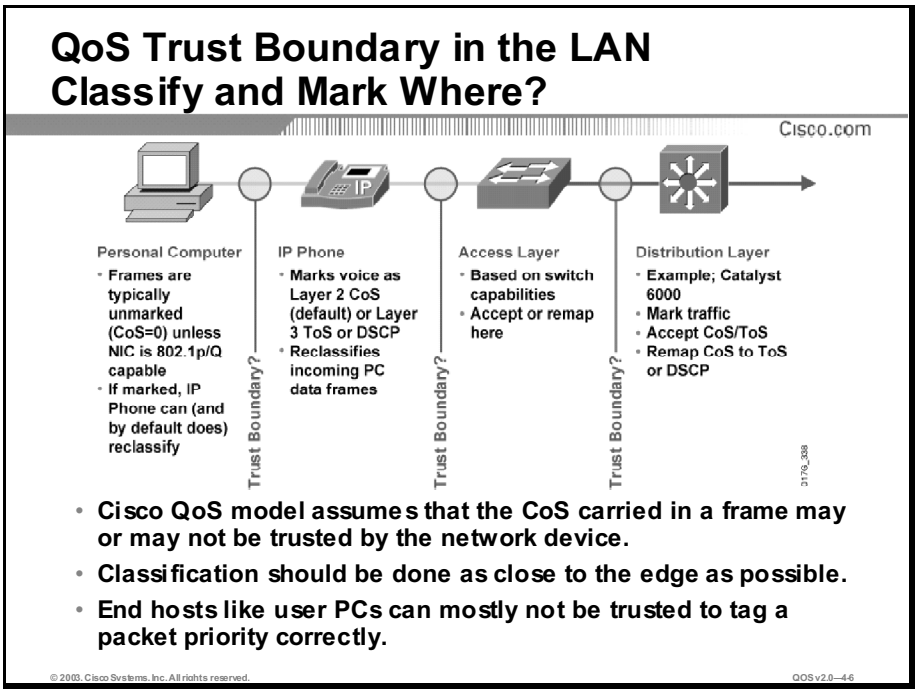

Classification should take place at the network edge, typically in the wiring closet or within video endpoints or IP Phones themselves.

The figure demonstrates this with an IP telephony example. Packets can be marked as important by using Layer 2 CoS settings in the user priority bits of the 802.1p portion of the 802.1p/Q field or the IP precedence/DSCP bits in the ToS/DS field in the IPv4 header. Cisco IP Phones can mark voice packets as high priority using CoS as well as ToS. By default, the IP Phone sends 802.1p tagged packets with the CoS and ToS set to a value of 5.

Because most PCs do not have an 802.1Q-capable network interface card (NIC), they send the packets untagged. This means that the frames do not have an 802.1p field. Also, unless the applications running on the PC send packets with a specific CoS value, this field is zero. A special case is where the TCP/IP stack in the PC has been modified to send all packets with a ToS value other than zero. Typically, this does not happen and the ToS value is zero.

Even if the PC is sending tagged frames with a specific CoS value, Cisco IP Phones can zero out this value before sending the frames to the switch. This is the default behavior. Voice frames coming from the IP Phone have a CoS of 5 and data frames coming from the PC have a CoS of 0. When the switch receives these frames, it can take into account these values for further processing based on its capabilities.

The switch uses its queues (available on a per-port basis) to buffer incoming frames before sending them to the switching engine. (It is important to remember that input queuing comes into play only when there is congestion.) The switch uses the CoS value(s) to put the frames in appropriate queues. The switch can also employ mechanisms, such as WRED, to make intelligent drops within a queue (also known as congestion avoidance) and WRR to provide more bandwidth to some queues than to others (also known as congestion management).

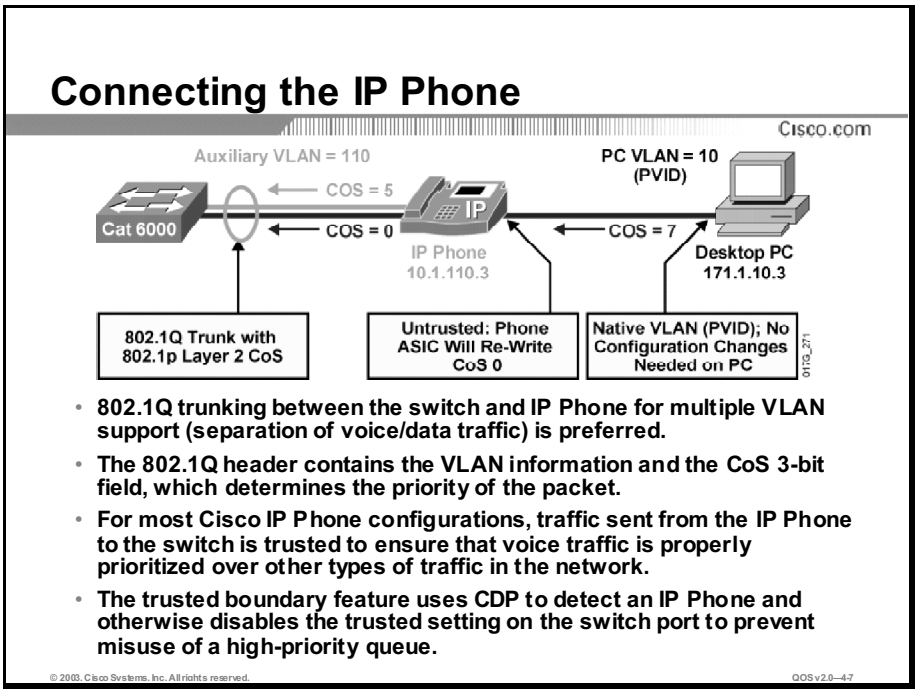

In a typical network, you connect a Cisco IP Phone to a switch port as shown in the figure. Traffic sent from the telephone to the switch is typically marked with a tag that uses the 802.1Q header. The header contains the VLAN information and the CoS 3-bit field, which determines the priority of the packet. For most Cisco IP Phone configurations, the traffic sent from the telephone to the switch is trusted to ensure that voice traffic is properly prioritized over other types of traffic in the network. By using the **mls qos trust device cisco-phone** and the **mls qos trust cos** interface configuration commands, you can configure the switch port to which the telephone is connected to trust the CoS labels of all traffic received on that port.

LAN Classification and Marking Platforms

This topic will discuss several of the Catalyst switches, highlighting their capabilities to perform QoS functions.

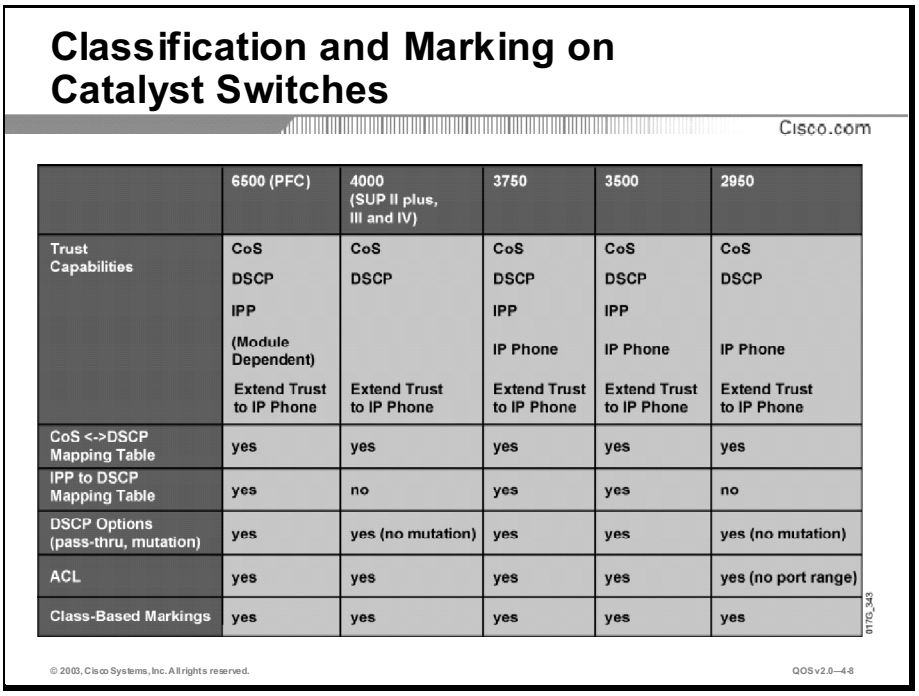

Catalyst 6500

With a Layer 2 switching engine and a policy feature card (PFC), QoS can classify traffic that is addressed to specified MAC address/VLAN pairs to be marked with a configured CoS value. Classification and marking with a Layer 2 switching engine uses Layer 2 CoS values. Classification and marking with a Layer 2 switching engine does not use or set Layer 3 IP precedence or DSCP values. Classification with a Layer 3 switching engine uses Layer 2, 3, and 4 values. Marking with a Layer 3 switching engine uses Layer 2 CoS values and Layer 3 IP precedence or DSCP values.

QoS schedules traffic through the transmit queues based on CoS values and uses CoS-value-based transmit-queue drop thresholds to avoid congestion in traffic that is transmitted from Ethernet ports. The implementation of scheduling and congestion avoidance is hardware-dependent, and with each specific platform different queue capabilities exist.

Queues are defined as a number of queues, the type of queue and the number of drop thresholds per queue. Here are a few examples:

- **2q2t:** Indicates two standard queues, each with two configurable tail-drop thresholds.
- **1p2q2t:** Indicates one strict-priority queue and two standard queues, each with two configurable WRED-drop thresholds.
- **1p3q1t:** Indicates one strict-priority queue and three standard queues, each with one configurable WRED-drop threshold (on 1p3q1t ports, each standard queue also has one nonconfigurable tail-drop threshold).
- **and so on.**

With 1p3q1t, the three standard transmit queues each have one WRED-drop threshold and one nonconfigurable tail-drop threshold.

- Frames with CoS 5 go to the strict-priority transmit queue (queue 4), where the switch drops frames only when the buffer is 100 percent full.
- Frames with CoS 0 and 1 go to the low-priority standard transmit queue (queue 1).
- Frames with CoS 2, 3, or 4 go to the medium-priority standard transmit queue (queue 2).
- Frames with CoS 6 or 7 go to the high-priority standard transmit queue (queue 3).

Catalyst 4000

Classification on the Catalyst 4000 is enabled only if QoS is globally enabled on the switch. By default, QoS is globally disabled and classification does not occur.

Specify which fields in the frame or packet that you want to use to classify incoming traffic.

For IP traffic, you have the following classification options:

- Trust the IP DSCP in the incoming packet (configure the port to trust DSCP), and assign the same DSCP to the packet for internal use.
- Trust the CoS value (if present) in the incoming packet, and generate the DSCP by using the CoS-to-DSCP map.
- Perform the classification based on a configured IP standard or extended ACL, which examines various fields in the IP header. If no ACL is configured, the packet is assigned the default DSCP based on the trust state of the ingress port; otherwise, the policy map specifies the DSCP to assign to the incoming frame.

A packet can be classified for QoS using multiple match criteria, and the classification can specify whether the packet should match all of the specified match criteria or at least one of the match criteria. To define a QoS classifier, provide the match criteria using the “match” statements in a class map. In the “match” statements, specify the fields in the packet to match on, or use IP standard or IP extended ACLs.

During QoS processing, the switch represents the priority of all traffic with an internal DSCP value:

- During classification, QoS uses configurable mapping tables to derive the internal DSCP from received CoS. These maps include the CoS-to-DSCP map.
- During policing, QoS can assign another DSCP value (if the packet is out of profile and the policer specifies a marked down DSCP value).
- Before the traffic reaches the scheduling stage, QoS uses the internal DSCP to select one of the four egress queues for output processing. The DSCP-to-egress queue mapping can be configured using the **qos map dscp to tx-queue** command.

The CoS-to-DSCP and DSCP-to-CoS map have default values that might or might not be appropriate for your network.

Each physical port has four transmit queues (egress queues). Each packet that needs to be transmitted is enqueued to one of the transmit queues. The transmit queues are then serviced based on the transmit queue scheduling algorithm.

After the final transmit DSCP is computed (including any markdown of DSCP), the transmit DSCP-to transmit-queue mapping configuration determines the transmit queue. The packet is placed in the transmit queue of the transmit port, determined from the transmit DSCP. Use the **qos map dscp to tx-queue** command to configure the transmit DSCP to transmit queue mapping.

The transmit queue 3 on each port can be configured as the priority queue using the **priority high tx-queue** configuration command in the interface configuration mode. When transmit queue 3 is configured with higher priority, packets in transmit queue 3 are scheduled ahead of packets in other queues.

Catalyst 3550

Each Gigabit-capable Ethernet port has four egress queues, one of which can be the egress expedite or priority queue. If the expedite (priority) queue is enabled, WRR services it until it is empty before servicing the other three queues. Ingress frame or packet classification options include:

- Non-IP traffic
  - Use the port default. If the frame does not contain a CoS value, the switch assigns the default port CoS value to the incoming frame. Then, the switch uses the configurable CoS-to-DSCP map to generate the internal DSCP value.
  - Trust the CoS value in the incoming frame (configure the port to trust CoS). Then, the switch uses the configurable CoS-to-DSCP map to generate the internal DSCP value. Layer 2 ISL frame headers carry the CoS value in the three least-significant bits of the 1-byte user field. Layer 2 802.1Q frame headers carry the CoS value in the three most significant bits of the TCI field. CoS values range from 0 for low priority to 7 for high priority.
  - The trust DSCP and trust IP precedence configurations are meaningless for non-IP traffic. If you configure a port with either of these options and non-IP traffic is received, the switch assigns the default port CoS value and generates the internal DSCP from the CoS-to-DSCP map.
  - Perform the classification based on the configured Layer 2 MAC ACL, which can examine the MAC source address, the MAC destination address, and the Ethertype field. If no ACL is configured, the packet is assigned the default DSCP of 0, which means best-effort traffic; otherwise, the policy map specifies the DSCP to assign to the incoming frame.
- IP traffic
  - Trust the IP DSCP in the incoming packet (configure the port to trust DSCP), and assign the same DSCP to the packet for internal use. The IETF defines the 6 most significant bits of the 1-byte ToS field as the DSCP. The priority represented by a particular DSCP value is configurable. DSCP values range from 0 to 63.
  - Trust the IP precedence in the incoming packet (configure the port to trust IP precedence), and generate a DSCP by using the configurable IP-precedence-to-DSCP map. The IP version 4 specification defines the three most significant bits of the 1-byte ToS field as the IP precedence. IP precedence values range from 0 for low priority to 7 for high priority.
  - Trust the CoS value (if present) in the incoming packet, and generate the DSCP by using the CoS-to-DSCP map.

- Perform the classification based on a configured IP standard or an extended ACL, which examines various fields in the IP header. If no ACL is configured, the packet is assigned the default DSCP of 0, which means best-effort traffic; otherwise, the policy map specifies the DSCP to assign to the incoming frame.

- Class maps and policy maps
  - A class map is a mechanism that you use to name and to isolate a specific traffic flow (or class) from all other traffic. The class map defines the criteria that is used to match against a specific traffic flow to further classify it; the criteria can include matching the access group defined by the ACL, matching a specific list of DSCP or IP precedence values, or matching a specific list of VLAN IDs associated with another class map that defines the actual criteria (for example, to match a standard or extended ACL). If you have more than one type of traffic that you want to classify, you can create another class map and use a different name. After a packet is matched against the class-map criteria, you further classify it through the use of a policy map.
  - A policy map specifies which traffic class to act on. Actions can include trusting the CoS, DSCP, or IP precedence values in the traffic class; setting a specific DSCP or IP precedence value in the traffic class; or specifying the traffic bandwidth limitations and the action to take when the traffic is out of profile. Before a policy map can be effective, you must attach it to an interface.

Classification and Marking on Catalyst 2950 Switches

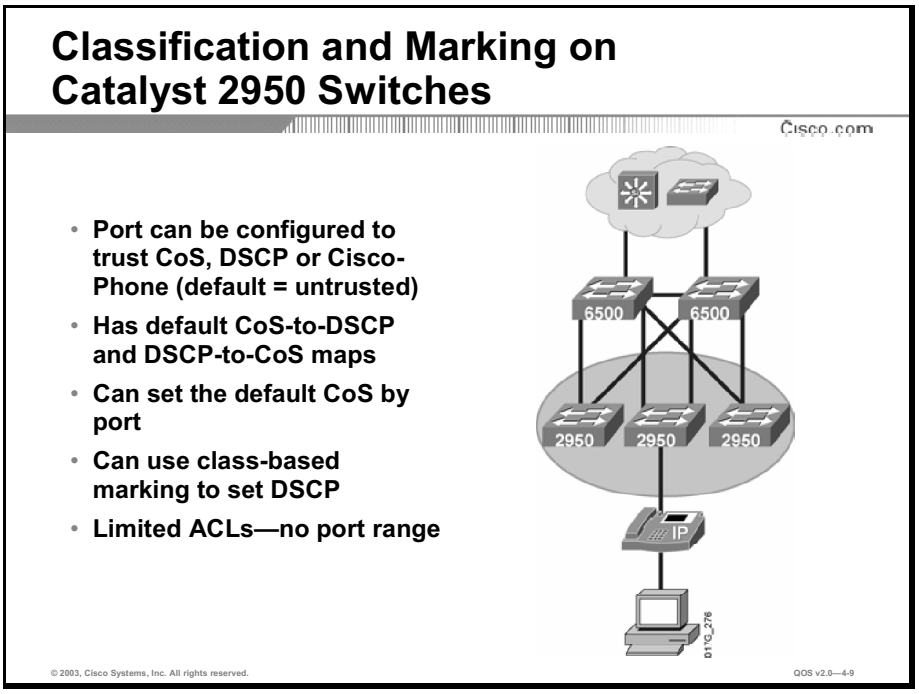

Cisco Catalyst 2950 series switches offer superior and highly granular QoS based on Layer 2 through Layer 4 information to ensure that network traffic is classified and prioritized, and that congestion is avoided in the best possible manner.

Cisco Catalyst 2950 series switches can classify, reclassify, police (determine if the packet is in or out of predetermined profiles and affect actions on the packet), and mark or drop the incoming packets before the packet is placed in the shared buffer. Packet classification allows the network elements to discriminate between various traffic flows and enforce policies based on Layer 2 and Layer 3 QoS fields.

The QoS implementation is based on the DiffServ architecture, an emerging standard from the IETF. This architecture specifies that each packet is classified upon entry into the network. The classification is carried in the IP packet header, using 6 bits from the deprecated IP ToS field to carry the classification (*class*) information.

Classification can also be carried in the Layer 2 frame:

- Prioritization values in Layer 2 frames
  - Layer 2 802.1Q frame headers used in trunks except for native VLAN frames.
  - Other frame types cannot carry Layer 2 CoS values.
- Prioritization bits in Layer 3 packets
  - Layer 3 IP packets with DSCP values 0, 8, 10, 16, 18, 24, 26, 32, 34, 40, 46, 48, and 56 only.

The Catalyst 2950 switch supports four egress queues, which allow the network administrator to be more discriminating in assigning priorities for the various applications on the LAN. Strict-priority scheduling configuration helps ensure that time-sensitive applications (such as voice) always follow an expedited path through the switch fabric. WRR scheduling, another significant enhancement, ensures that lower-priority traffic receives attention without comprising the priority settings administered by a network manager. These features allow

network administrators to prioritize mission-critical, time-sensitive traffic, such as voice (IP telephony traffic), ERP (Oracle, SAP, and so on), and CAD/CAM over less time-sensitive applications such as FTP or e-mail (Simple Mail Transfer Protocol [SMTP]).

Actions at the egress interface include queuing and scheduling:

- Queuing evaluates the CoS value and determines which of the four egress queues in which to place the packet.
- Scheduling services the four egress queues based on their configured WRR.

The Catalyst 2950 supports packet classification based on QoS ACLs as follows:

- You can use IP standard, IP extended, and Layer 2 MAC ACLs to define a group of packets with the same characteristics (*class*). In the QoS context, the permit and deny actions in the access control entries (ACEs) have different meanings than with security ACLs.
- If a match with a permit action is encountered (first-match principle), the specified QoS-related action is taken.
- If no match with a permit action is encountered and all the ACEs have been examined, no QoS processing occurs on the packet.
- If multiple ACLs are configured on an interface, the packet matches the first ACL with a permit action, and QoS processing begins.
- Configuration of a deny action is not supported in QoS ACLs on the switch.

After a traffic class has been defined with the ACL, you can attach a policy to it. A policy might contain multiple classes with actions specified for each one of them. A policy might include commands to classify the class as a particular aggregate (for example, assign a DSCP) or rate-limit the class. This policy is then attached to a particular port on which it becomes effective.

You implement IP ACLs to classify IP traffic by using the **access-list** global configuration command; you implement Layer 2 MAC ACLs to classify Layer 2 traffic by using the **mac access-list extended** global configuration command.

In the case of frames that arrive without a CoS value (such as untagged frames), these switches support classification based on a default CoS value per port assigned by the network administrator. After the frames have been classified or reclassified using one of the above modes, they are assigned to the appropriate queue at the egress port.

Note: To use the features described in this chapter, you must have the enhanced software image (EI) installed on your switch.

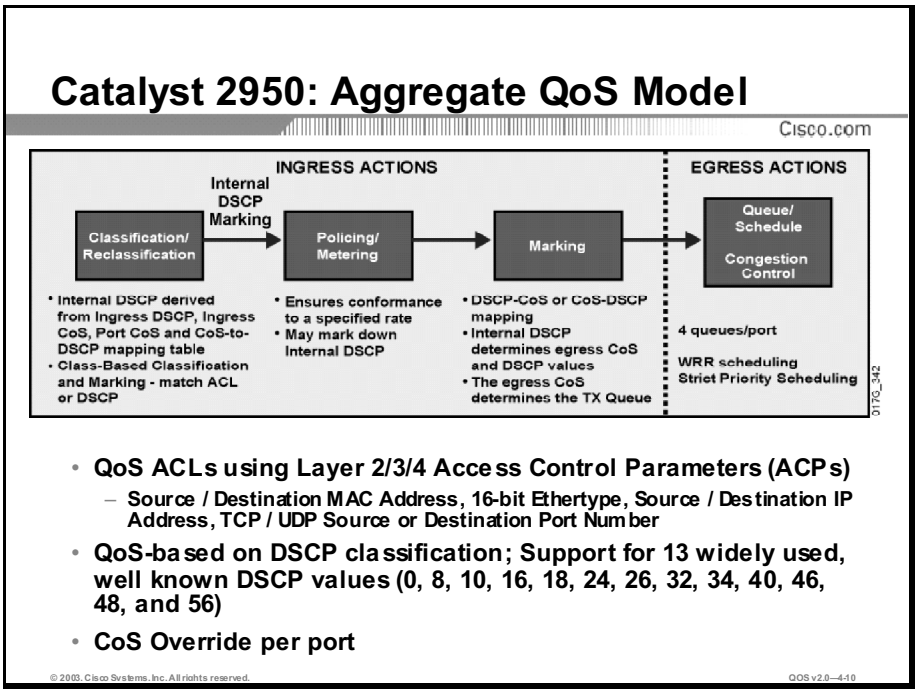

The example in the illustration provides a quick synopsis of what happens on the Catalyst 2950 regarding QoS.

On incoming packets, classification and reclassification are performed by identifying packet groups using either DSCP or CoS. Policing and metering, if configured, are then performed on the packets to ensure compliance to configure rates. Marking is the last action performed on incoming packets based on the CoS-to-DSCP or DSCP-to-CoS mappings.

Outgoing packets are scheduled and queued for congestion control. There are 4 queues per port and are scheduled based on WRR and strict priority scheduling.

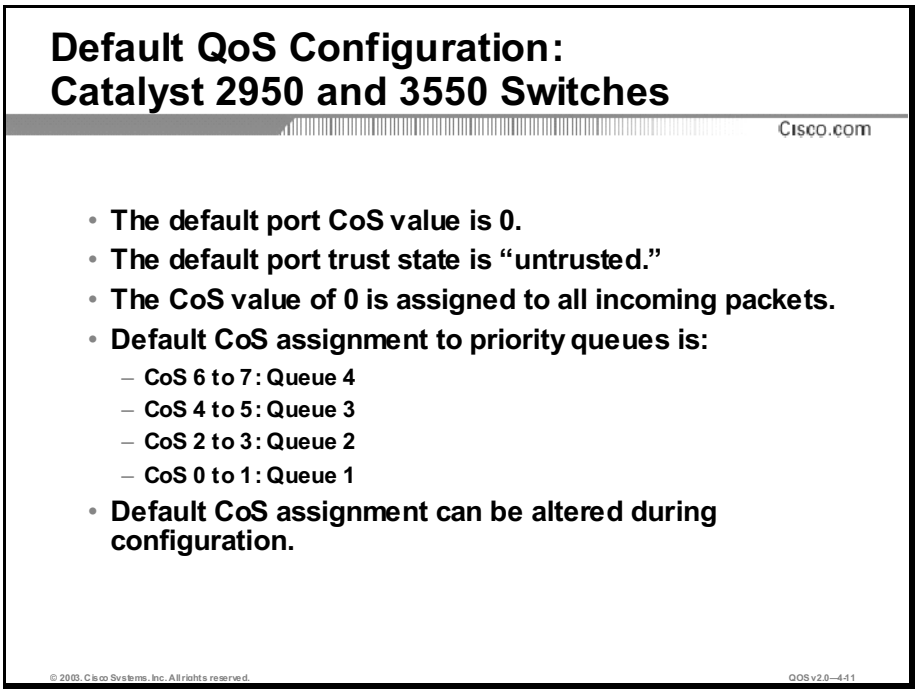

The default QoS settings for the Catalyst 2950 and 3550 switches is as follows:

- The default port CoS value is 0.
- The CoS value of 0 is assigned to all incoming packets.
- The default port trust state is untrusted. If a port is connected to an IP Phone, should change the default port config to trust the CoS setting from the IP Phone using the **mls qos trust** command.
- No policy maps are configured.
- No policers are configured.
- Default CoS assignment to priority queues is:
  - CoS 6 to 7: Queue 4
  - CoS 4 to 5: Queue 3
  - CoS 2 to 3: Queue 2
  - CoS 0 to 1: Queue 1

Mapping Tables: Catalyst 2950 and 3550 Switches

Cisco.com

- During QoS processing, the switch represents the priority of all traffic (including non-IP traffic) with an internal DSCP value.
- During classification, QoS uses configurable mapping tables to derive the internal DSCP (a 6-bit value) from received CoS value.

Default CoS-to-DSCP Map								
CoS value	0	1	2	3	4	5	6	7
DSCP value	0	8	16	24	32	40	48	56

- Before the traffic reaches the scheduling stage, QoS uses the configurable DSCP-to-CoS map to derive a CoS value from the internal DSCP value.

Default DSCP-to-CoS Map								
DSCP values	0	8,10	16,18	24,26	32,34	40,46	48	56
CoS values	0	1	2	3	4	5	6	7

Actions at the egress interface include queuing and scheduling:

- Queuing evaluates the internal DSCP and determines which of the four egress queues in which to place the packet. The DSCP value is mapped to a CoS value, which selects one of the queues.
- Scheduling services the four egress queues based on their configured WRR weights and thresholds. One of the queues can be the expedite queue, which is serviced until empty before the other queues are serviced. Congestion avoidance techniques include tail drop and WRED on Gigabit-capable Ethernet ports and tail drop (with only one threshold) on 10/100 Ethernet ports.

During QoS processing, the switch represents the priority of all traffic (including non-IP traffic) with an internal DSCP value:

- During classification, QoS uses configurable mapping tables to derive the internal DSCP (a 6-bit value) from received CoS or IP precedence (3-bit) values. These maps include the CoS-to-DSCP map and the IP-precedence-to-DSCP map.

On an ingress interface configured in the DSCP-trusted state, if the DSCP values are different between the QoS domains, you can apply the configurable DSCP-to-DSCP-mutation map to the interface that is on the boundary between the two QoS domains.

- During policing, QoS can assign another DSCP value to an IP or non-IP packet (if the packet is out of profile and the policer specifies a marked down DSCP value). This configurable map is called the policed-DSCP map.
- Before the traffic reaches the scheduling stage, QoS uses the configurable DSCP-to-CoS map to derive a CoS value from the internal DSCP value. Through the CoS-to-egress-queue map, the CoS values select one of the four egress queues for output processing.

The CoS-to-DSCP, DSCP-to-CoS, and the IP-precedence-to-DSCP (Catalyst 3550 only) map have default values that might or might not be appropriate for your network.

The default DSCP-to-DSCP-mutation (Catalyst 3550 only) map and the default policed-DSCP map are null maps; they map an incoming DSCP value to the same DSCP value.

The DSCP-to-DSCP-mutation map is the only map you apply to a specific Gigabit-capable Ethernet port or to a group of 10/100 Ethernet ports.

All other maps apply to the entire switch.

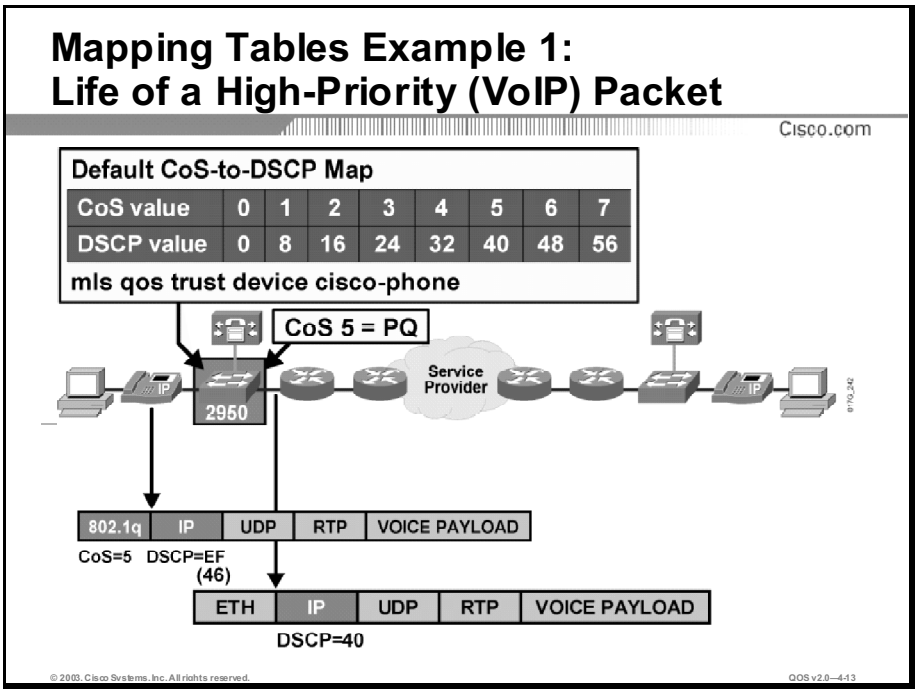

This figure provides an example of a CoS value mapped to the DSCP value in a Catalyst 2950 switch.

The trust boundary has been established on the switch port to trust the CoS setting from the IP Phone. By default, the CoS and DSCP value of a packet coming from a Cisco IP Phone is set to CoS 5 and DSCP EF (numeric 46).

On the output of the switch, in the Layer 3 header, the DSCP will be set to 40 using the default CoS-to-DSCP map.

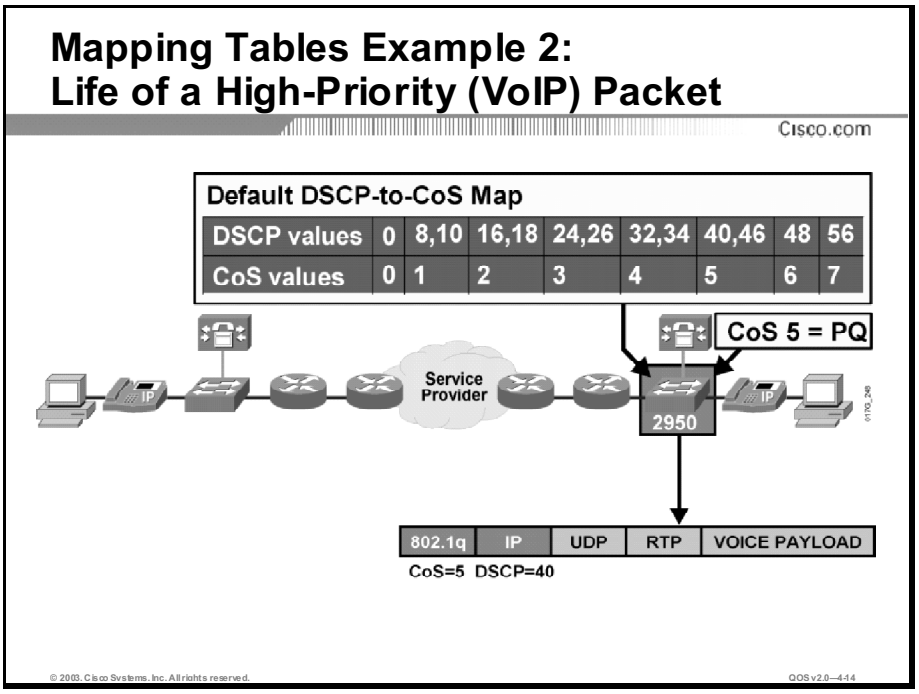

This figure shows the previous packet as it arrives at its destination after traversing the network.

In this example, the switch port connecting to the router is set to trust DSCP. Therefore, the Layer 3 header will have a DSCP value of 40 (from the previous slide) and as it traverses the switch, its CoS value is set to 5 using the default DSCP-to-CoS map.

Configuring LAN-Based Classification and Marking

This topic identifies the Cisco IOS commands that are required to configure LAN-based classification and marking.

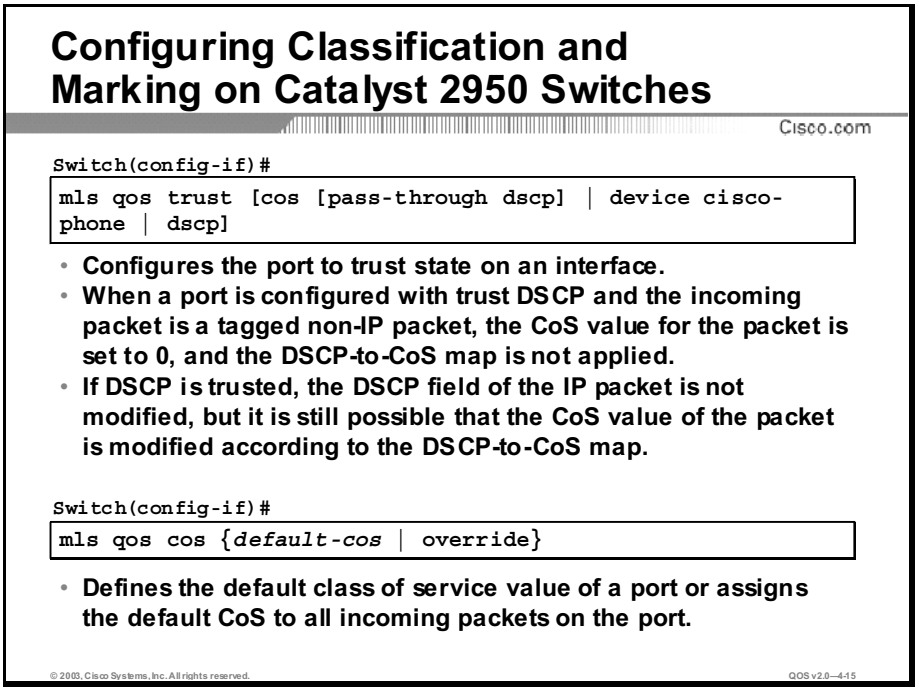

This figure shows some of the QoS configuration commands that are necessary for Catalyst 2950 switches. The defaults for its interfaces are as follows:

- The port is not trusted.
- Pass-through mode is disabled.
- Trusted boundary is disabled.
- If no keyword is specified and the switch is running the EI, the default is dscp.

mls qos trust [cos [pass-through dscp] | device cisco-phone | dscp]

Syntax Description

Parameter	Description
cos	(Optional) Specifies that the CoS bits in incoming frames are trusted and derives the internal DSCP value from the CoS bits.
cos pass-through dscp	(Optional) Configure the interface to classify ingress packets by trusting the CoS value and to send packets without modifying the DSCP value (pass-through mode).
device cisco-phone	(Optional) Classify ingress packets by trusting the value sent from the Cisco IP Phone (trusted boundary).
dscp	(Optional) Classify ingress packets with packet DSCP values (most significant 6 bits of the 8-bit service-type field). For non-IP packets, the packet CoS value is set to 0. This keyword is available only if your switch is running the EI software.

To define the default CoS value for an interface, use the **mls qos cos** command. Use the **no** form of this command to remove a prior entry. QoS assigns the CoS value specified with **mls qos cos** interface configuration command to untagged frames received on trusted and untrusted ports. The default cos value is 0.

mls qos cos *cos-value*

Syntax Description

Parameter	Description
cos-value	Default CoS value for the interface; valid values are from 0 to 7.

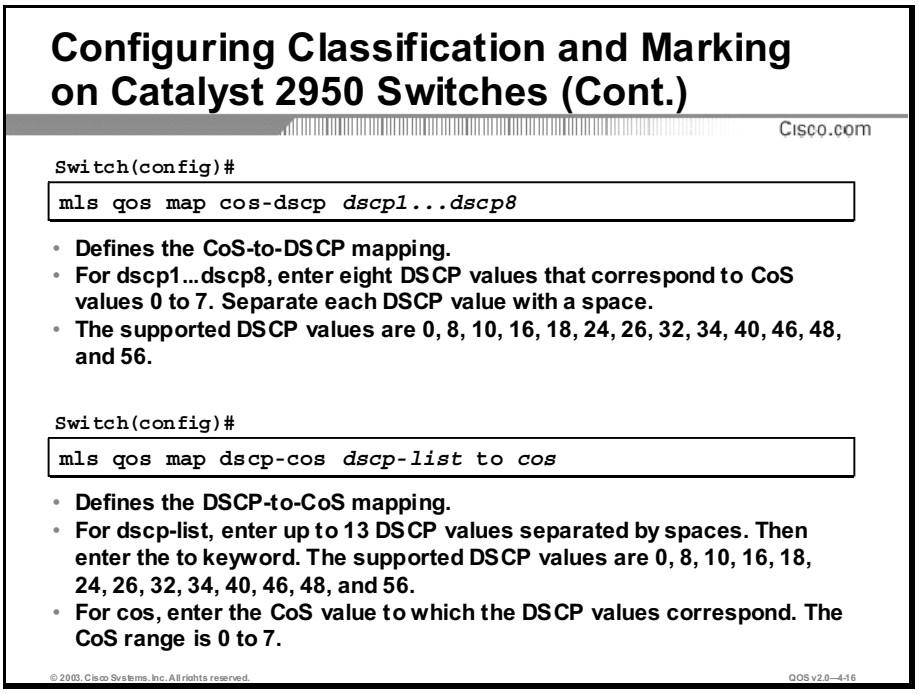

The commands listed in the figure show how to change the default CoS-to-DSCP and DSCP-to-CoS mappings.

CoS-to-DSCP Default Mapping

Marker	Value							
CoS Values	0	1	2	3	4	5	6	7
DSCP Values	0	8	16	24	32	40	48	56

To define the ingress CoS-to-DSCP mapping for trusted interfaces, use the **mls qos map cos-dscp** command. The CoS-to-DSCP map is used to map the CoS of packets arriving on trusted interfaces (or flows) to a DSCP where the trust type is trust-cos. This map is a table of eight CoS values (0 through 7) and their corresponding DSCP values. Use the **no** form of this command to remove a prior entry.

mls qos map cos-dscp *values*

Syntax Description

Parameter	Description
values	Eight DSCP values, separated by spaces, corresponding to the CoS values; valid values are from 0 to 63.

DSCP-to-CoS Default Mapping

Marker	Value							
DSCP Values	0	8, 10	16, 18	24, 26	32, 34	40, 42	48	56
CoS Values	0	1	2	3	4	5	6	7

To define an egress DSCP-to-CoS mapping, use the **mls qos map dscp-cos** command. The DSCP-to-CoS map is used to map the final DSCP classification to a final CoS. You use the DSCP-to-CoS map to map DSCP values in incoming packets to a CoS value, which is used to select one of the four egress queues. The CoS map is written into the ISL header or 802.1Q tag of the transmitted packet on trunk interfaces and contains a table of 64 DSCP values and the corresponding CoS values. You can enter up to eight DSCP values separated by a space. You can enter up to eight CoS values separated by a space. Use the **no** form of this command to remove a prior entry.

mls qos map dscp-cos *dscp-values* **to** *cos-values*

Syntax Description

Parameter	Description
dscp-values	DSCP values; valid values are from 0 to 63.
to	Defines mapping.
cos-values	CoS values; valid values are from 0 to 63.

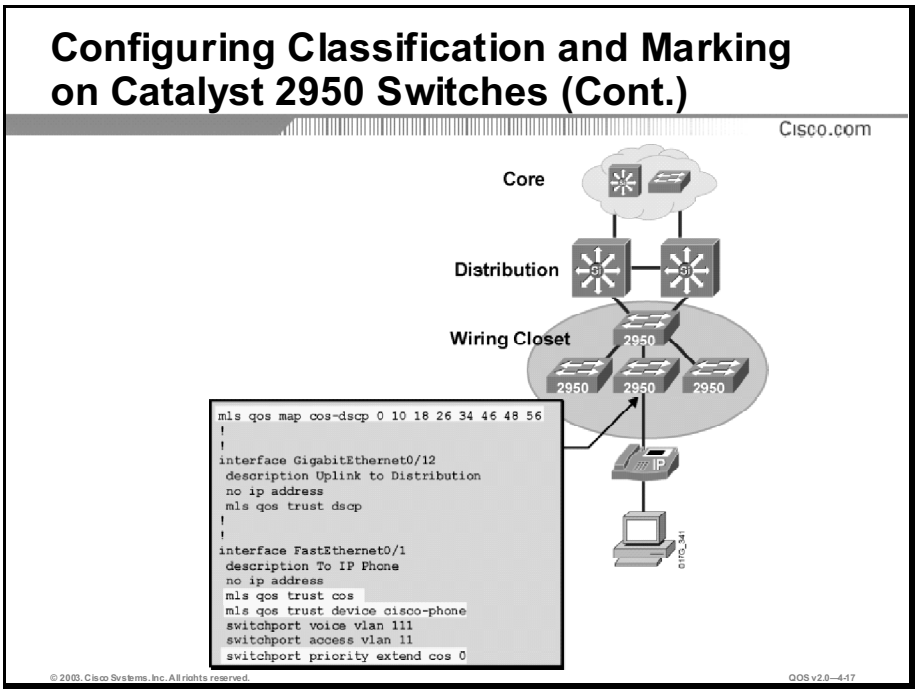

This figure shows a configuration example on a Catalyst 2950 switch where the CoS-to-DSCP map has been changed from the default.

The default map is:

Marker	Value							
CoS Values	0	1	2	3	4	5	6	7
DSCP Values	0	8	16	24	32	40	48	56

And the map after configuration is:

Marker	Value							
CoS Values	0	1	2	3	4	5	6	7
DSCP Values	0	10	18	26	34	46	48	56

Also we see that interface has been set to trust the CoS value using the **mls qos** trust command using both the cos and cisco-phone options. The result of the configuration is that the switch interface to trust CoS only when a Cisco IP Phone is attached. The switch uses CDP to detect if a Cisco IP Phone is attached and also passes the voice VLAN ID information to the Cisco IP Phone using CDP.

The last command in the configuration is the **switchport priority extend cos 0** command. The **switchport priority extend cos 0** interface configuration command is used to enable the IP Phone to override the CoS marking from the PC attached to the IP Phone with a CoS value of 0.

Use the **switchport priority extend** interface configuration command to set a port priority for the incoming untagged frames or the priority of frames received by the IP Phone connected to the specified port. Use the **no** form of this command to return to the default setting.

switchport priority extend {**cos** *value* | **trust**}

Syntax Description

Parameter	Description
cos *value*	Set the IP Phone port to override the priority received from PC or the attached device. The CoS value is a number from 0 to 7. 7 is the highest priority. The default is 0.
trust	Set the IP Phone port to trust the priority received from PC or the attached device.

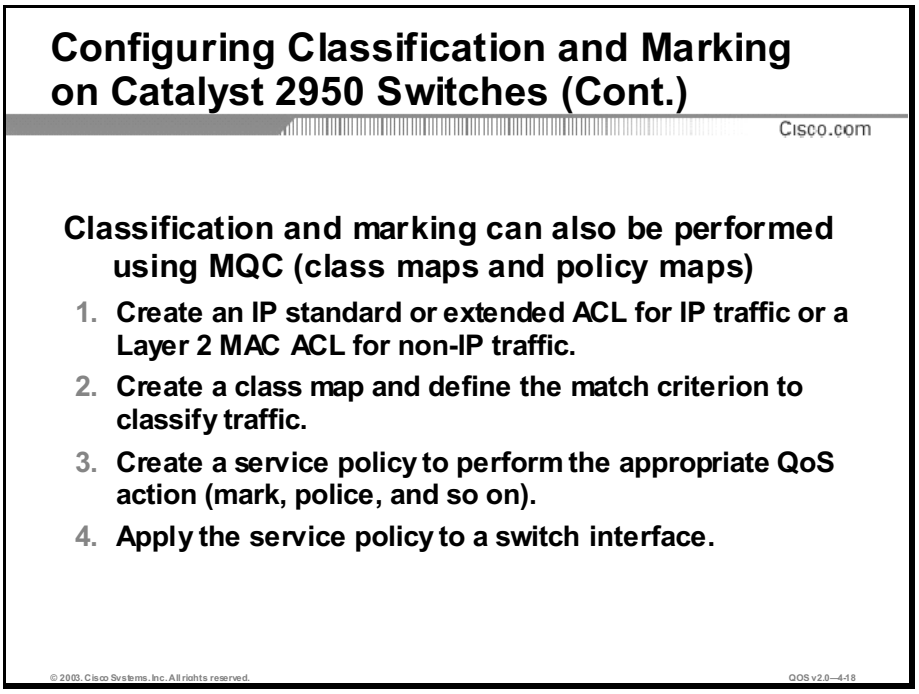

A class map is a mechanism that you use to isolate and name a specific traffic flow (or class) from all other traffic. The class map defines the criteria used to match against a specific traffic flow to further classify it; the criteria can include matching the access group defined by the ACL. If you have more than one type of traffic that you want to classify, you can create another class map and use a different name. After a packet is matched against the class-map criteria, you further classify it through the use of a policy map.

A policy map specifies which traffic class to act on. Actions can include setting a specific DSCP value in the traffic class or specifying the traffic bandwidth limitations and the action to take when the traffic is out of profile. Before a policy map can be effective, you must attach it to an interface.

You create a class map by using the **class-map** global configuration command or the **class** policy-map configuration command. You should use the **class-map** global configuration command when the map is shared among many ports. When you enter the **class-map** global configuration command, the switch enters the class-map configuration mode. In this mode, you define the match criterion for the traffic by using the **match** class-map configuration command.

You create and name a policy map by using the **policy-map** global configuration command. When you enter this command, the switch enters the policy-map configuration mode. In this mode, you specify the actions to take on a specific traffic class by using the **class** policy-map configuration or **set** policy-map class configuration command. To make the policy map effective, you attach it to an interface by using the **service-policy** interface configuration command.

The policy map can also contain commands that define the policer, the bandwidth limitations of the traffic, and the action to take if the limits are exceeded.

You use the **class-map** global configuration command to isolate a specific traffic flow (or class) from all other traffic and to name it. The class map defines the criteria to use to match against a specific traffic flow to further classify it. Match statements can only include ACLs.

The match criterion is defined with one match statement entered within the class-map configuration mode.

A policy map specifies which traffic class to act on. Actions can include trusting the CoS or DSCP values in the traffic class; setting a specific DSCP value in the traffic class; and specifying the traffic bandwidth limitations for each matched traffic class (policer) and the action to take when the traffic is out of profile (marking).

A policy map also has these characteristics:

- A policy map can contain multiple class statements, each with different match criteria and policers.
- A separate policy-map class can exist for each type of traffic received through an interface.

You can attach only one policy map per interface in the input direction.

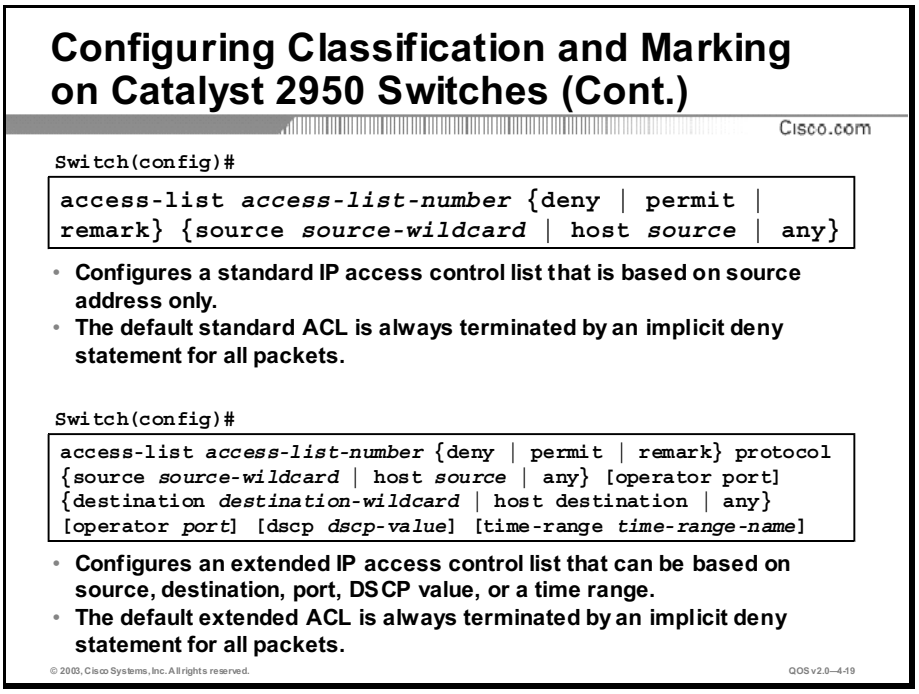

You can use IP standard, IP extended, and Layer 2 MAC ACLs to define a group of packets with the same characteristics (*class*). In the QoS context, the permit and deny actions in the ACEs have different meanings than with security ACLs:

- If a match with a permit action is encountered (first-match principle), the specified QoS-related action is taken.
- If no match with a permit action is encountered and all the ACEs have been examined, no QoS processing occurs on the packet.
- If multiple ACLs are configured on an interface, the packet matches the first ACL with a permit action, and QoS processing begins.
- Configuration of a deny action is not supported in QoS ACLs on the switch.

Use the standard version of the **access-list** global configuration command to configure a standard IP ACL. Use the **no** form of this command to remove a standard IP ACL.

access-list *access-list-number* **{deny | permit | remark}** {***source source-wildcard*** **| host** source **| any}**

Syntax Description

Parameter	Description
access-list-number	Number of an ACL, from 1 to 99 or from 1300 to 1999.
deny	Deny access if conditions are matched.
permit	Permit access if conditions are matched.
remark	ACL entry comment up to 100 characters.
source source-wildcard | **host** source | **any**	Define a source IP address and wildcard. The *source* is the source address of the network or host from which the packet is being sent, specified in one of these ways: ■ The 32-bit quantity in dotted-decimal format. The *source-wildcard* applies wildcard bits to the source. ■ The keyword **host**, followed by the 32-bit quantity in dotted-decimal format, as an abbreviation for *source* and *source-wildcard* of *source* 0.0.0.0. ■ The keyword **any** as an abbreviation for *source* and *source-wildcard* of 0.0.0.0 255.255.255.255. You do not need to enter a source-wildcard.

Use the extended version of the **access-list** global configuration command to configure an extended IP ACL. Use the **no** form of this command to remove an extended IP ACL.

access-list *access-list-number* **{deny | permit | remark}** *protocol* **{***source source-wildcard* **| host** source **| any} [***operator*** port] {***destination destination-wildcard*** **| host** destination **| any} [***operator*** port] [dscp** *dscp-value***] [time-range** *time-range-name***]**

Syntax Description

Parameter	Description
access-list-number	Number of an ACL, from 100 to 199 or from 2000 to 2699.
protocol	Name of an IP protocol. *protocol* can be **ip**, **tcp**, or **udp**.
deny	Deny access if conditions are matched.
permit	Permit access if conditions are matched.
remark	ACL entry comment up to 100 characters.
source source-wildcard | **host** source | **any**	Define a source IP address and wildcard. The *source* is the source address of the network or host from which the packet is being sent, specified in one of these ways: ■ The 32-bit quantity in dotted-decimal format. The *source-wildcard* applies wildcard bits to the source. ■ The keyword **host**, followed by the 32-bit quantity in dotted-decimal format, as an abbreviation for *source* and *source-wildcard* of *source* 0.0.0.0. ■ The keyword **any** as an abbreviation for *source* and *source-wildcard* of 0.0.0.0 255.255.255.255. You do not need to enter a source-wildcard.
destination destination-wildcard | **host** destination | **any**	Define a destination IP address and wildcard. The *destination* is the destination address of the network or host from which the packet is being sent, specified in one of these ways: ■ The 32-bit quantity in dotted-decimal format. The *destination-wildcard* applies wildcard bits to the destination. ■ The keyword **host**, followed by the 32-bit quantity in dotted-decimal format, as an abbreviation for *destination* and *destination-wildcard* of *source* 0.0.0.0. ■ The keyword **any** as an abbreviation for *destination* and *destination-wildcard* of 0.0.0.0 255.255.255.255. You do not need to enter a destination-wildcard.
operator port	(Optional) Define a source or destination port. The *operator* can be only **eq** (equal). If *operator* is after the source IP address and wildcard, conditions match when the source port matches the defined port. If *operator* is after the destination IP address and wildcard, conditions match when the destination port matches the defined port. The *port* is a decimal number or name of a Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) port. The number can be from 0 to 65535. Use TCP port names only for TCP traffic. Use UDP port names only for UDP traffic.

Parameter	Description
dscp *dscp-value*	(Optional) Define a Differentiated Services Code Point (DSCP) value to classify traffic. For the *dscp-value*, enter any of the 13 supported DSCP values (0, 8, 10, 16, 18, 24, 26, 32, 34, 40, 46, 48, and 56), or use the question mark (**?**) to see a list of available values.
time-range *time-range-name*	(Optional) For the **time-range** keyword, enter a meaningful name to identify the time range. For a more detailed explanation of this keyword, refer to the software configuration guide.

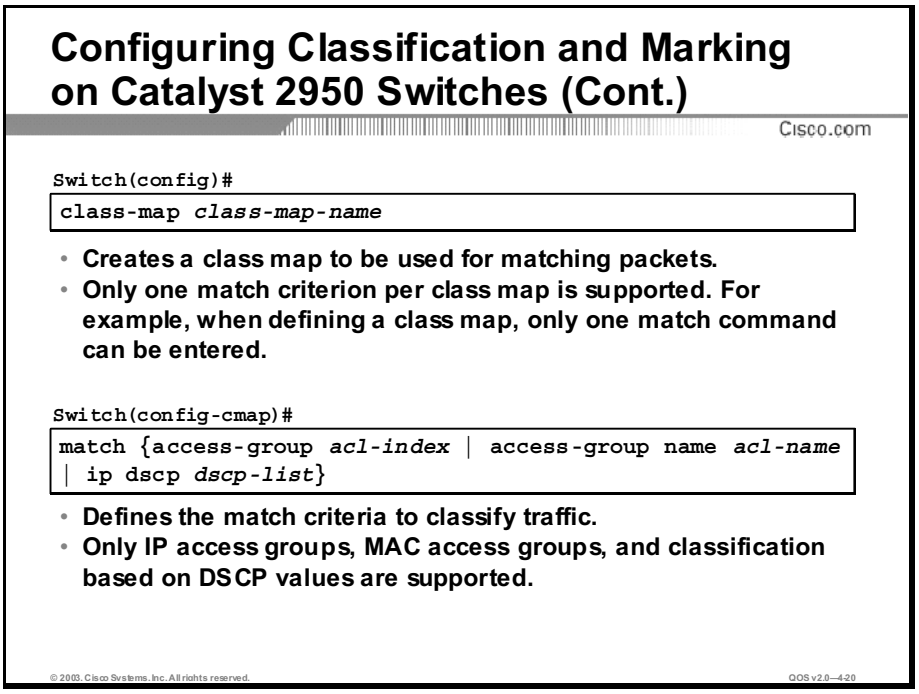

MQC class maps can also be used on Catalyst 2950 switches for packet classification purposes. However the **match** command used in conjunction with the class map has different parameters when executed on a Catalyst switch.

Use the **match** class-map configuration command to define the match criteria to classify traffic. Use the **no** form of this command to remove the match criteria.

match {access-group *acl-index* | **access-group name** *acl-name* | **ip dscp** ***dscp-list***}

Syntax Description

Parameter	Description
`access-group acl-index`	Number of an IP standard or extended ACL. For an IP standard ACL, the ACL index range is 1 to 99 and 1300 to 1999. For an IP extended ACL, the ACL index range is 100 to 199 and 2000 to 2699.
`access-group name` `acl-name`	Name of an IP standard or extended ACL or name of an extended MAC ACL. **Note** The ACL name must begin with an alphabetic character to prevent ambiguity with numbered ACLs. A name also cannot contain a space or quotation mark.
`ip dscp` *`dscp-list`*	List of up to eight DSCP values for each match statement to match against incoming packets. Separate each value with a space. The supported DSCP values are 0, 8, 10, 16, 18, 24, 26, 32, 34, 40, 46, 48, and 56.

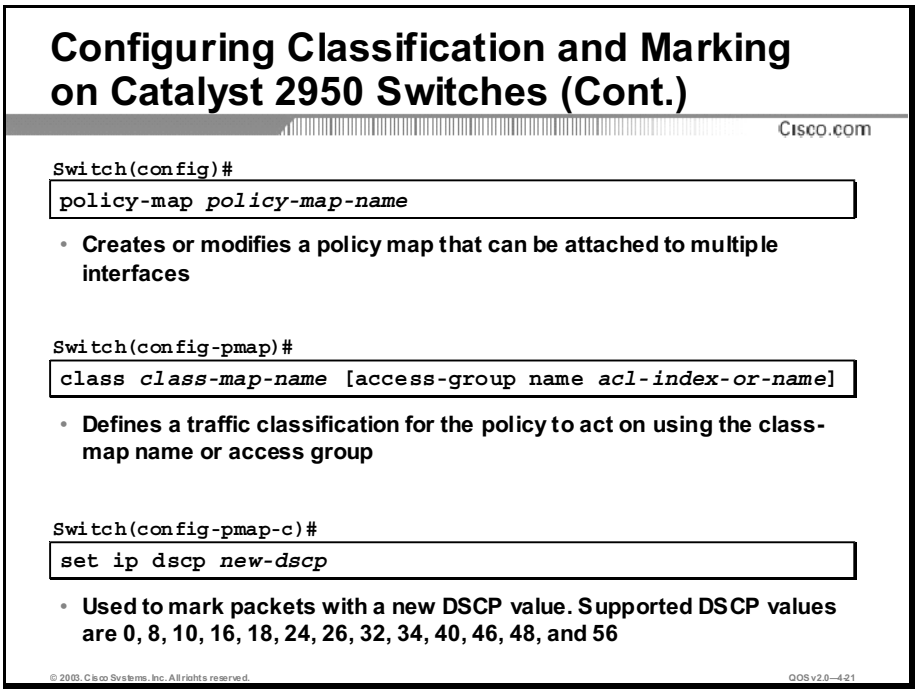

Recall that a policy map creates or modifies a policy that can be attached to multiple interfaces. The class command defines traffic classification for the policy to act on based on the class map or the access group. Use the **class** policy-map configuration command to define a traffic classification for the policy to act on using the class map name or access group. Use the **no** form of this command to delete an existing class map.

class *class-map-name* [**access-group name** *acl-index-or-name*]

Syntax Description

Parameter	Description
access-group name acl-index-or-name	(Optional) Number or name of an IP standard or extended ACL or name of an extended MAC ACL. For an IP standard ACL, the index range is 1 to 99 and 1300 to 1999; for an IP extended ACL, the index range is 100 to 199 and 2000 to 2699.

Use the **set** policy-map class configuration command to classify IP traffic by setting a DSCP value.

set ip dscp *new-dscp*

Syntax Description

Parameter	Description
new-dscp	New DSCP value assigned to the classified traffic. The supported DSCP values are 0, 8, 10, 16, 18, 24, 26, 32, 34, 40, 46, 48, and 56.

Configuring Classification and Marking on Catalyst 2950 Switches (Cont.)

Cisco.com

```
Switch(config-if)#
service-policy input policy-map-name
```

- **Applies a policy map defined by the policy-map command to the input of a particular interface**

```
mac access-list extended maclist1
  permit host 0001.0000.0001 host 0002.0000.0001
!
class-map macclass1
 match access-group name maclist1
!
policy-map macpolicy1
 class macclass1
   set ip dscp 26
!
interface gigabitethernet0/1
 switchport mode trunk
 mls qos trust cos
 service-policy input macpolicy1
```


The last step in configuring a policy is to apply the policy to the interface.

In the above example an extended access-list has been created for a mac address, maclist1. A class-map, macclass1 has been created that will match any MAC address permitted by the access-list maclist1.

If there is a match for the class map macclass1 the DSCP field will be set to 26 as defined in the policy-map macpolicy1.

This policy map has been implemented on the gigabit Ethernet port 0/1 for incoming packets.

Monitoring LAN-Based Classification and Marking

This topic describes some of the Cisco IOS commands that can be used to monitor QoS on Catalyst switches.

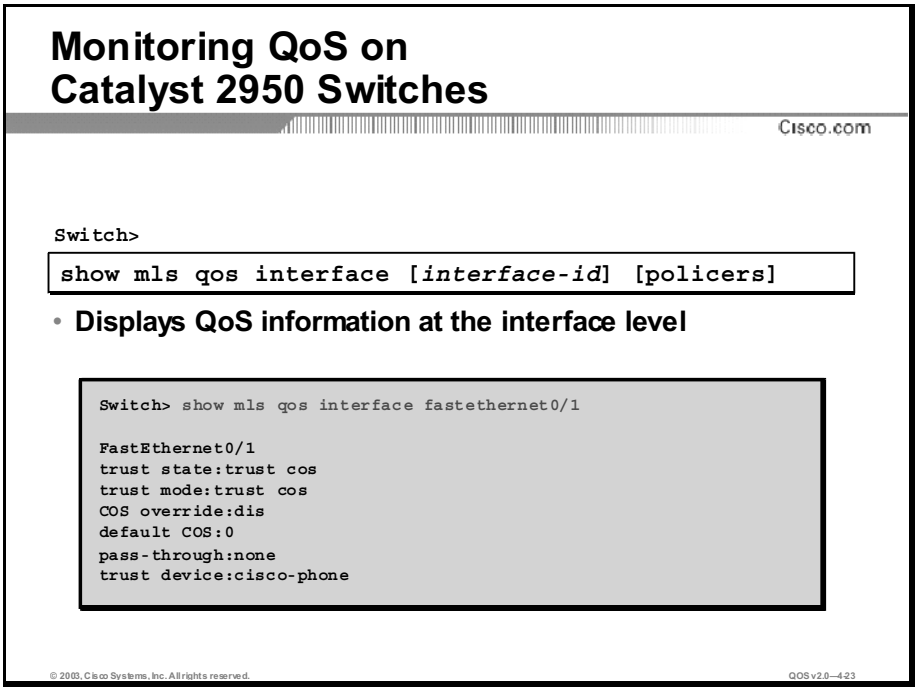

After QoS has been configured on a Catalyst switch, the network administrator will want to verify proper operation of QoS and the policies they may have configured. In the example, we see that the trust state has been set for CoS and that the default value of CoS is 0.

Use the **show mls qos interface** user EXEC command to display QoS information at the interface level.

show mls qos interface [*interface-id*] [policers]

Syntax Description

Parameter	Description
interface-id	(Optional) Display QoS information for the specified interface.
policers	(Optional) Display all the policers configured on the interface, their settings, and the number of policers unassigned (available only when the switch is running the EI software).

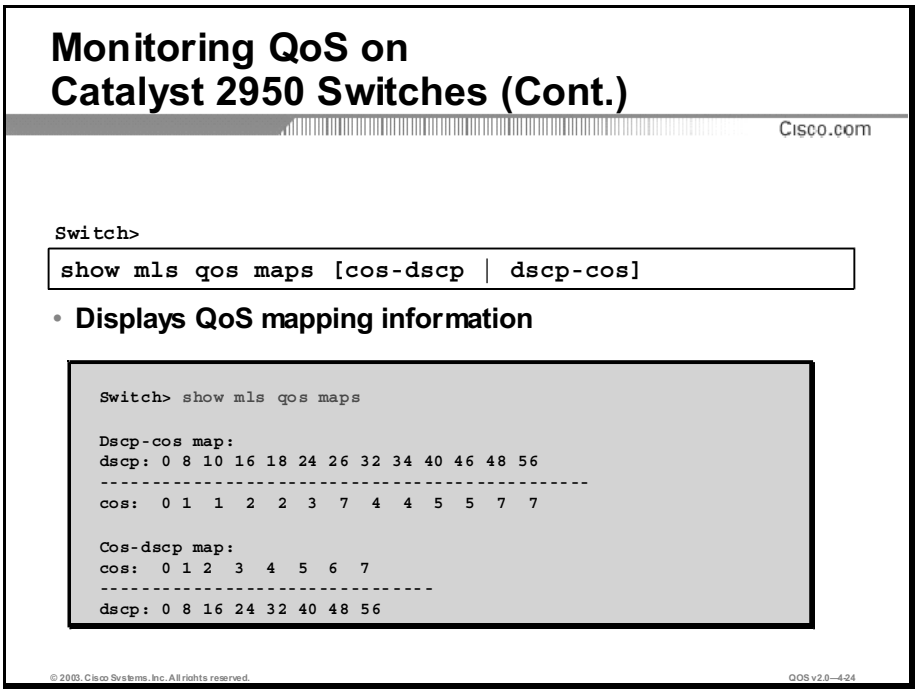

Another important monitoring command is shown above. The **show mls qos maps** command will display the CoS-to-DSCP and DSCP-to-CoS mappings.

Use the **show mls qos maps** user EXEC command to display QoS mapping information. Maps are used to generate an internal DSCP value, which represents the priority of the traffic.

show mls qos maps [**cos-dscp** | **dscp-cos**]

Syntax Description

Parameter	Description
cos-dscp	(Optional) Display CoS-to-DSCP map.
dscp-cos	(Optional) Display DSCP-to-CoS map.

Summary

This topic summarizes the key points discussed in this lesson.

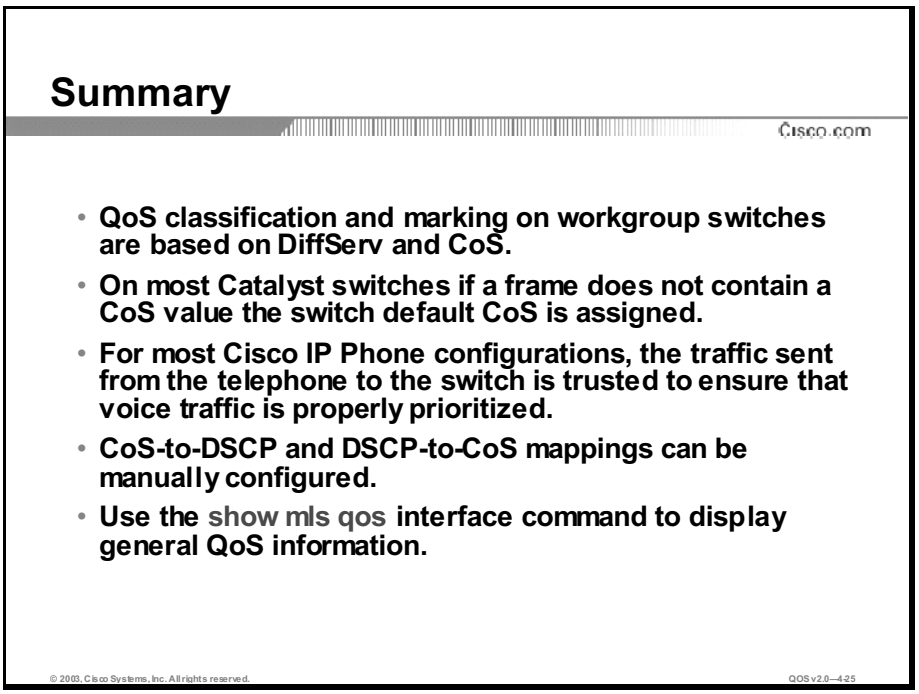

References

For additional information, refer to these resources:

- For more information on classification and marking on the Catalyst 2950, refer to "Configuring QoS" at the following URL: http://www.cisco.com/univercd/cc/td/doc/product/lan/cat2950/12111yj4/lrescg/swqos.htm
- For more information on configuring classification and marking on the Catalyst 2950, refer to "LAN Based Packet Classification" at the following URL: http://www.cisco.com/application/pdf/en/us/guest/products/ps628/c1051/ccmigration_09186a0080150b6c.pdf
- For more information on classification and marking on the Catalyst 4000, refer to "Configuring QoS" at the following URL: http://www.cisco.com/univercd/cc/td/doc/product/lan/cat4000/rel7_1/config/qos.htm
- For more information on classification and marking on the Catalyst 6500, refer to "Configuring QoS" at the following URL: http://www.cisco.com/univercd/cc/td/doc/product/lan/cat6000/sw_7_6/confg_gd/qos.htm

Next Steps

For the associated lab exercise, refer to the following section of the course Lab Guide:

- Lab Exercise 4-4: LAN-Based Packet Classification and Marking

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) If multiple queues are present, frames can be placed in different queues and serviced via _____.

A) Weighted Fair Queue

B) Weighted Round Robin

C) CB-Weighted Fair Queue

D) Weighted Random Early Detection

Q2) What is the default CoS value as used by Cisco IP Phones for voice packets?

A) 0

B) 3

C) 5

D) 7

Q3) The default CoS value on Catalyst 2950 and 3550 switches is _____.

A) 0

B) 3

C) 5

D) 7

Q4) The command to assign a default CoS value on a Catalyst switch is _____.

A) **mls qos**

B) **qos cos**

C) **mls qos cos**

D) **qos mls cos**

Q5) The command to display the CoS-to-DSCP and DSCP-to-CoS maps is _____.

A) **show maps**

B) **show mls maps**

C) **show qos maps**

D) **show mls qos maps**

Quiz Answer Key

Q1) B

Relates to: LAN-Based Classification and Marking

Q2) C

Relates to: QoS Trust Boundaries

Q3) A

Relates to: LAN Classification and Marking Platforms

Q4) C

Relates to: Configuring LAN-Based Classification and Marking

Q5) D

Relates to: Monitoring LAN-Based Classification and Marking

Module Assessment

Overview

Use this assessment to test what you learned in this module. The correct answers and solutions are found in the Module Assessment Answer Key.

Quiz: Classification and Marking

Complete the Quiz to assess what you have learned in the module.

Objectives

This activity tests your knowledge on how to meet these objectives:

- Explain the purpose of classification and marking and how they can be used to define a QoS service class
- Use MQC CLI commands to classify packets
- Use class-based marking to assign packets to a specific service class
- Use NBAR to discover network protocols and applications, and to classify packets
- Use the QoS pre-classify feature to classify GRE, IPSec, and L2F and L2TP encapsulated packets
- Explain how to implement classification and marking in an interdomain network using QPPB
- Describe LAN-based methods for implementing classification and marking

Instructions

Complete these steps:

Step 1 Answer all questions in this quiz by selecting the best answer(s) to each question

Step 2 Verify your results against the answer key located at the end of this section

Step 3 Review the topics in this module that relates to the questions that you answered incorrectly.

Q1) Classification of packets should occur _____.

A) at the distribution layer
B) anywhere in the core of the network
C) as close to the source of the traffic as possible
D) as close to the destination of the traffic as possible

Q2) To utilize a class map, QoS must be referenced through the use of _____.

A) route map
B) access list
C) policy map
D) service map

Q3) What is a requirement for using CB marking?

A) CEF must be enabled
B) CEF must be disabled
C) CEF can only be used on serial interfaces
D) CEF can only be used on Ethernet interfaces

Q4) What is the MQC feature that allows traffic to be classified by a packet sub-port number?

A) LDPM
B) NBAR
C) service maps
D) service classes

Q5) The QoS for VPN feature is designed to operate on _____.

A) logical interfaces
B) loopback interfaces
C) tunnel interfaces
D) physical interfaces

Q6) Which of the following is the proper command that will allow Cisco routers to make a copy of the inner IP header and to run a QoS classification before encryption based on fields in the inner IP header?

A) **qos classify**
B) **qos pre-classify**
C) **qos nbar classify**
D) **qos vpn classify**

Q7) Which of the following commands will modify metric and tag values when the IP routing table is updated with BGP learned routes?

A) table-map
B) bgp-policy
C) map bpg ip
D) bgp table-map

Q8) Which of the following commands will enable the propagation of the QoS policy via BGP on an interface?

A) table-map
B) bgp-policy
C) bgp send-policy
D) bgp policy-propagation

Q9) Which of the following commands will display both the CoS-to-DSCP and DSCP-to-CoS mappings on a Catalyst switch?

A) show mls maps
B) show mls qos maps
C) show mls maps both
D) show qos mls maps both

Scoring

You have successfully completed the quiz for this lesson when you earn a score of 80 percent or better.

Module Assessment Answer Key

Q1) C

Relates to: Classification and Marking Overview

Q2) C

Relates to: Using MQC Classification

Q3) A

Relates to: Using MQC for Class-Based Marking

Q4) B

Relates to: Using NBAR for Classification

Q5) C

Relates to: QoS Pre-Classify

Q6) B

Relates to: QoS Pre-Classify

Q7) A

Relates to: Configuring QoS Policy Propagation Through BGP

Q8) B

Relates to: Configuring QoS Policy Propagation Through BGP

Q9) B

Relates to: Configuring LAN Classification and Marking

Module Summary

This topic summarizes the key points discussed in this module.

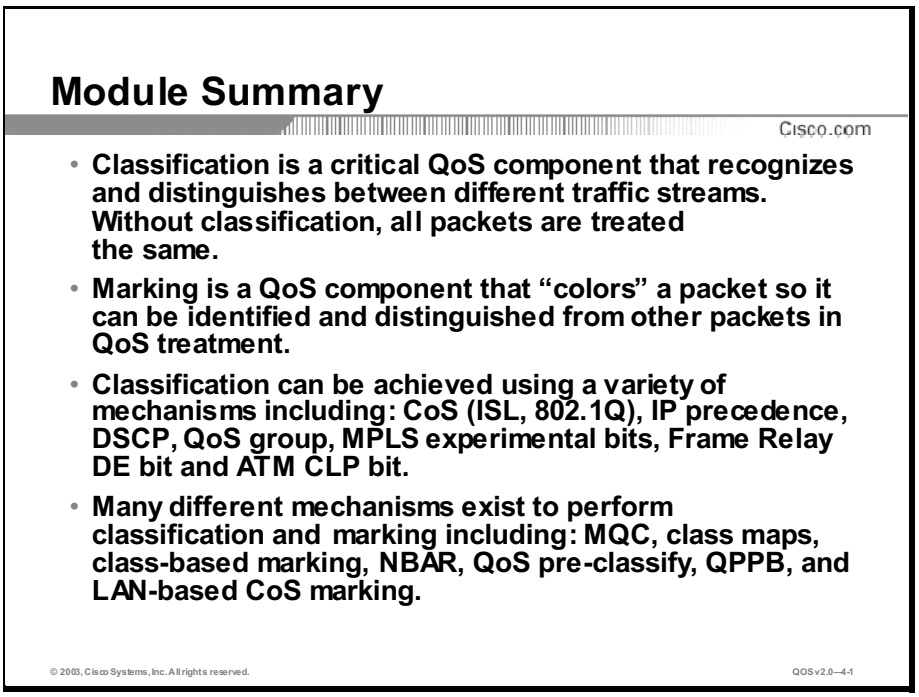

Classification is the process of identifying traffic and categorizing it into different classes. Packet classification allows some packets to be handled more quickly or with a higher priority than other packets. Applications such as voice typically need to be treated faster than a file transfer.

Classification uses a traffic descriptor to categorize a packet within a specific group to define that packet. Typically, used traffic descriptors include: CoS (ISL, 802.1Q) incoming interface, IP precedence, DSCP, QoS group ID, MPLS experimental bits, Frame Relay DE bit, ATM CLP bit, source or destination address, or application.

Marking a packet or frame with its classification allows network devices to easily distinguish the marked packet or frame. Marking is a useful feature in that it allows network devices to easily identify packets or frames as belonging to a specific class. After packets have been identified as belonging to a specific class, QoS mechanisms can be uniformly applied to ensure compliance with administrative QoS policies.

Packet classification can be implemented using such tools MQC class maps and policy maps; NBAR, QoS pre-classify (VPN QoS), QPPB.

Classification and Marking can be done at the network or link layer.

Module 5

Congestion Management

Overview

Congestion can occur in many different locations within a network and is the result of many factors including oversubscription, insufficient packet buffers, traffic aggregation points, network transit points, and WAN links. Increasing link bandwidth is not a simple fix that solves the congestion issue in most cases. Aggressive traffic can fill interface queues and starve more fragile flows such as voice and interactive traffic. The results can be devastating for these delay-sensitive traffic types, making it difficult to meet the service level requirements these applications require. Fortunately, there are many congestion management techniques available on Cisco IOS platforms. These congestion management techniques provide network administrators with an effective means to manage software queues and to allocate the required bandwidth to specific applications when congestion conditions exist.

This module examines the components of queuing systems and the different congestion management mechanisms available on Cisco IOS devices.

Module Objectives

Upon completing this module, you will be able to use Cisco QoS queuing mechanisms to manage network congestion.

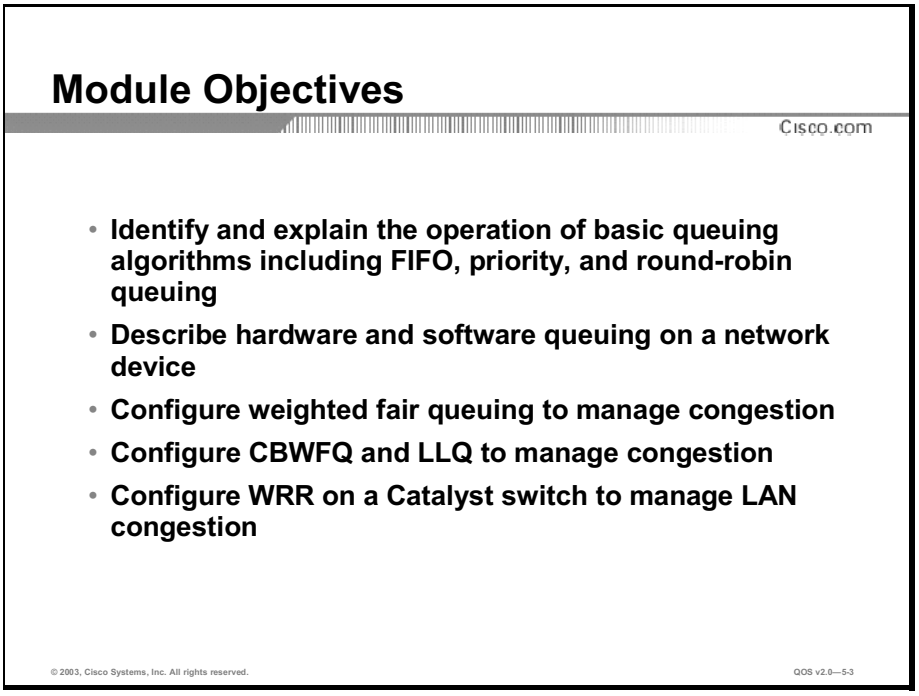

Module Outline

The outline lists the components of this module.

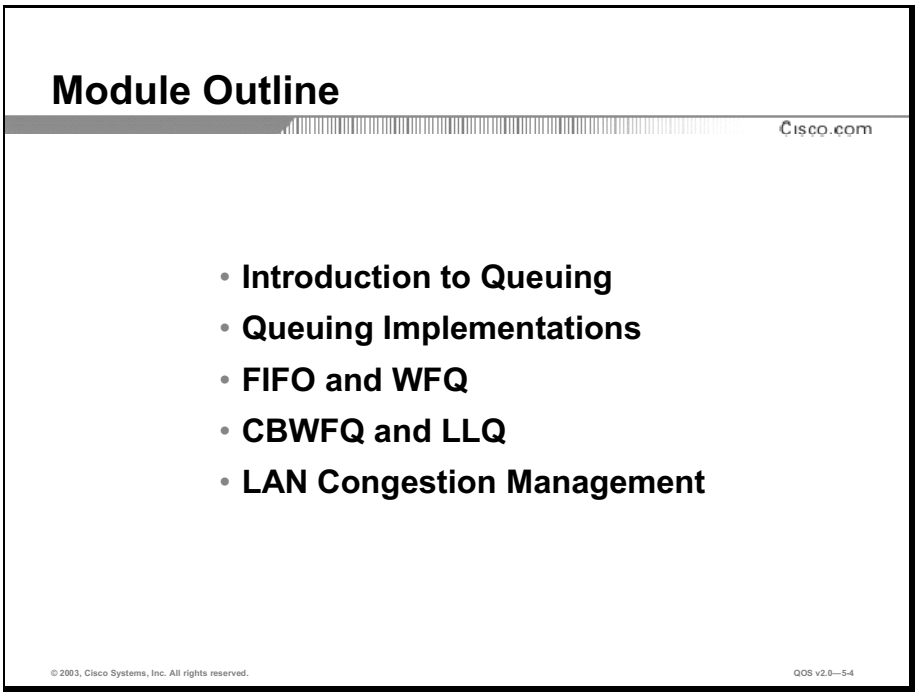

Introduction to Queuing

Overview

Queuing algorithms are one of the primary ways to manage congestion in a network. Network devices handle an overflow of arriving traffic by using a queuing algorithm to sort traffic and determine a method of prioritizing the traffic onto an output link. Each queuing algorithm was designed to solve a specific network traffic problem and has a particular effect on network performance. This lesson describes the basic queuing algorithms.

Relevance

In order to understand how an advanced queuing mechanism such as class-based weighted fair queuing (CBWFQ) will operate on a Cisco router, it is important to first understand the basic queuing mechanisms upon which CBWFQ is built.

Objectives

Upon completing this lesson, you will be able to identify and explain the operation of basic queuing algorithms including FIFO, priority, and round-robin queuing. This includes being able to meet these objectives:

- Explain the need for congestion management mechanisms
- List the different queuing algorithms
- Describe FIFO queuing
- Describe priority queuing
- Describe round-robin queuing
- Describe weighted round-robin queuing
- Describe deficit round-robin queuing

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- No special skills or knowledge are required.

Outline

The outline lists the topics included in this lesson.

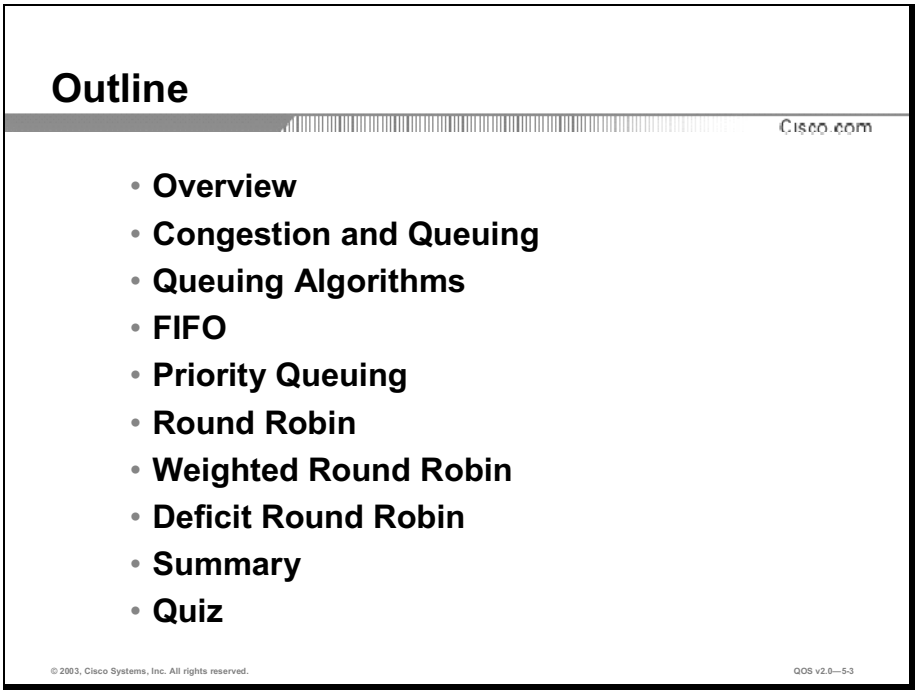

Congestion and Queuing

This topic explains the relationship between congestion and queuing.

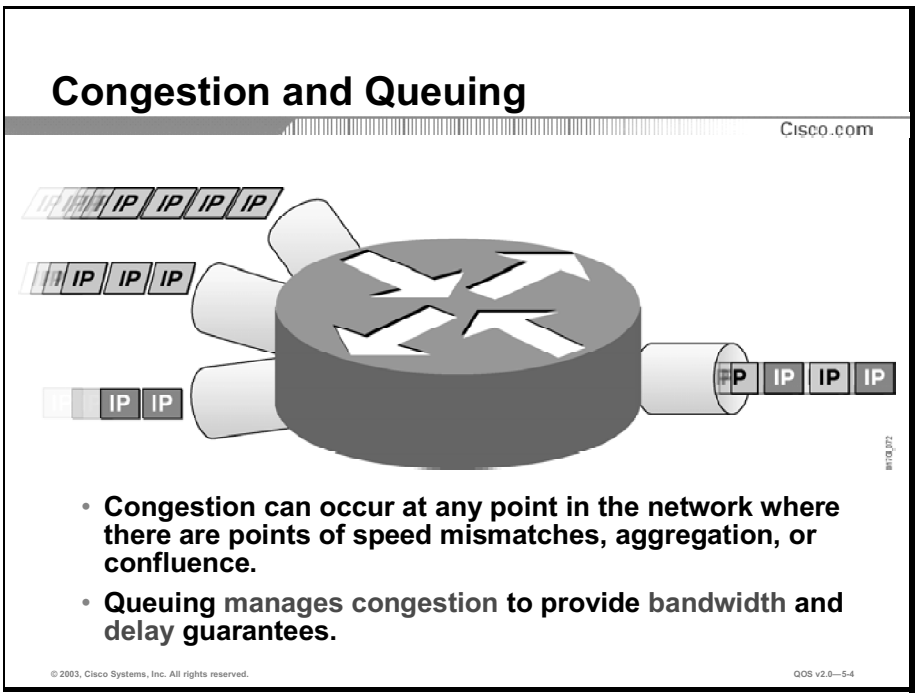

Congestion can occur anywhere within a network where speed mismatches (that is, a 1000-Mbps link feeding a 100-Mbps link), aggregation (that is, multiple 100-Mbps links feeding an upstream 100-Mbps link), or confluence (the flowing together of two or more traffic streams).

Queuing algorithms are used to manage congestion. Many algorithms have been designed to serve different needs. A well-designed queuing algorithm will provide some bandwidth and delay guarantees to priority traffic.

Congestion and Queuing Speed Mismatch

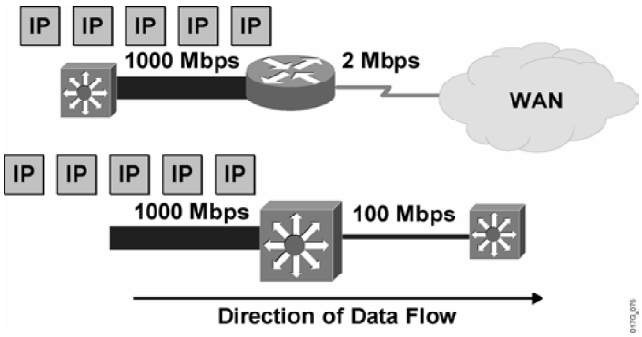

- **Speed mismatches are the most typical cause of congestion.**
- **Possibly persistent when going from LAN to WAN.**
- **Usually transient when going from LAN to LAN.**

Speed mismatches are the most typical cause of congestion in a network.

Speed mismatches are most common when traffic moves from a high-speed LAN environment (100 or 1000 Mbps) to lower-speed WAN links (1 or 2 Mbps). Speed mismatches are also common in LAN-to-LAN environments when, for example, a 1000-Mbps link feeds into a 100-Mbps link. In these situations, congestion tends to be persistent and must continually be managed.

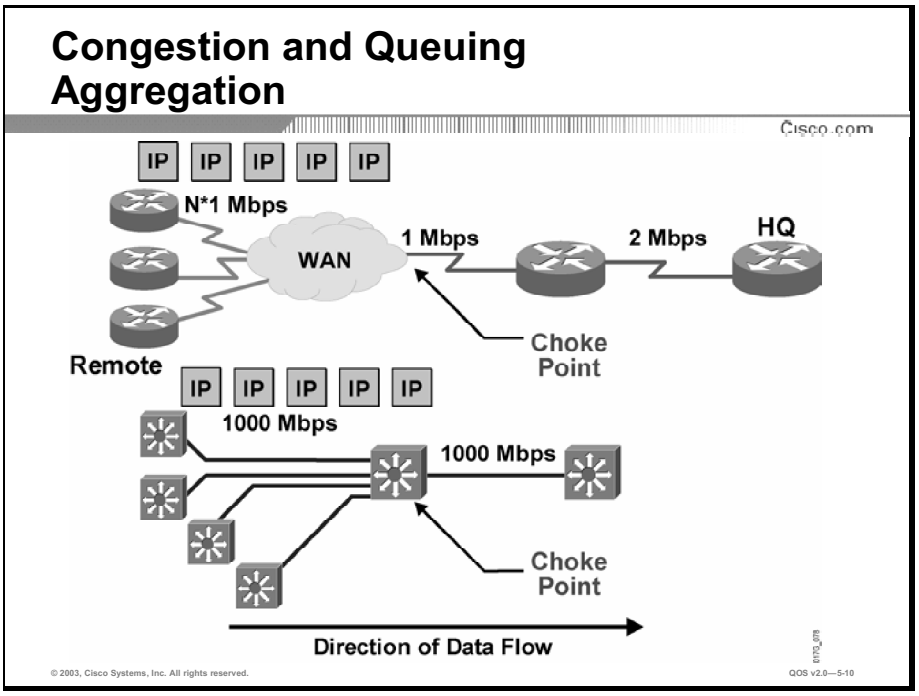

The second most common site of congestion is at points of aggregation in a network.

Typical points of aggregation occur in WANs when multiple remote sites feed back into a central services site.

In a LAN environment, congestion resulting from aggregation often occurs at the distribution layer of networks where the different access layer devices feed traffic to the distribution-level switches.

Queuing Algorithms

This topic lists the different queuing algorithms.

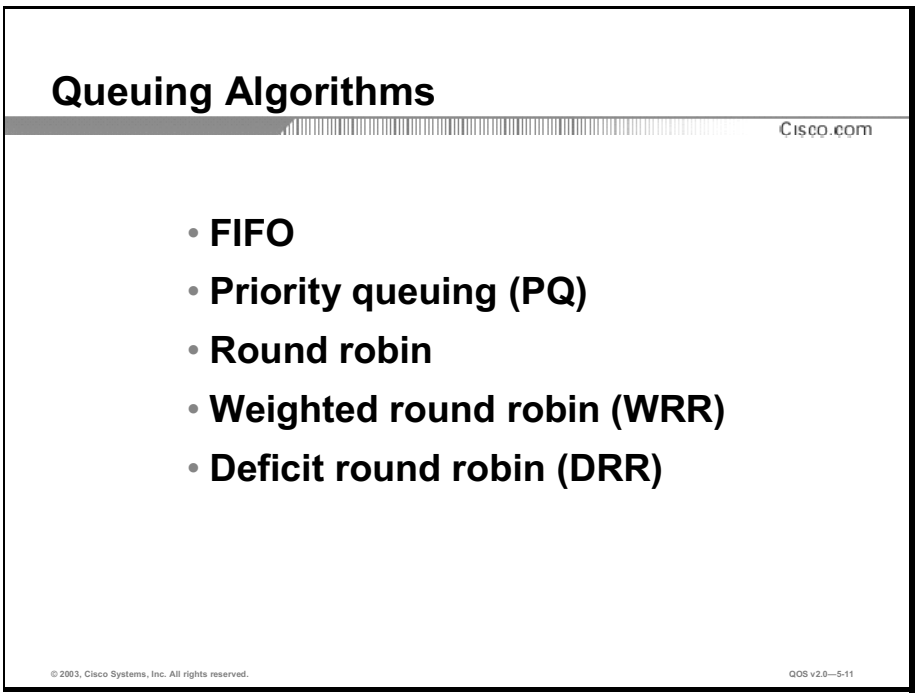

Key queuing algorithms include:

- **FIFO:** The simplest algorithm.
- **Priority queuing (PQ):** Allows traffic to be prioritized.
- **Round robin:** Allows several queues to share bandwidth.
- **Weighted round robin (WRR):** Allows sharing of bandwidth with prioritization.
- **Deficit round robin (DRR):** Resolves problem with some WRR implementations.

FIFO

This topic describes the FIFO queuing algorithm.

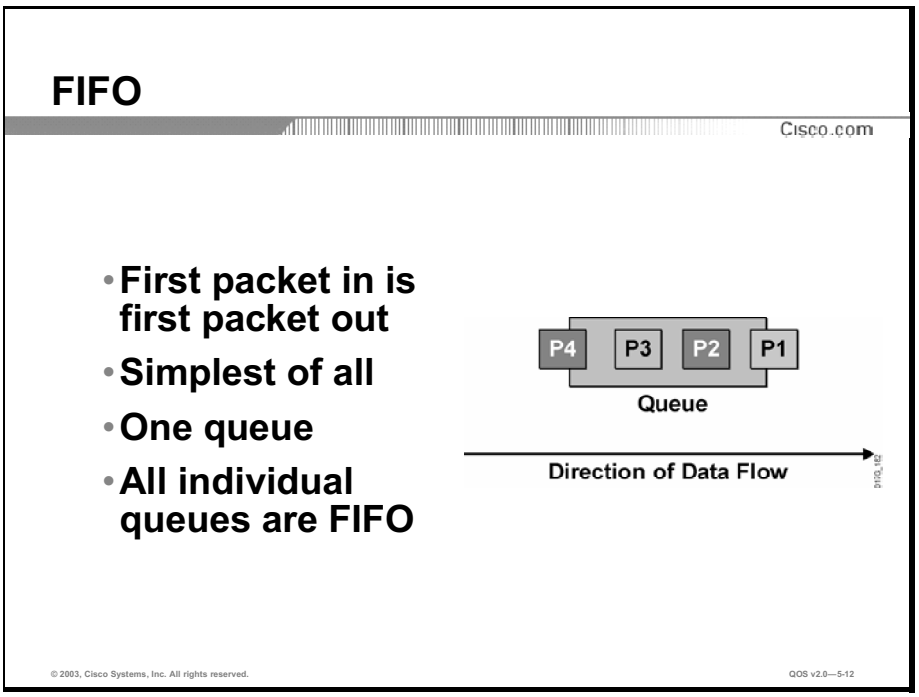

FIFO is the simplest queuing algorithm.

Packets are placed into a single queue and serviced in the order they were received.

All individual queues are, in fact, FIFO queues. Other queuing methods rely upon FIFO as the congestion management mechanism for single queues while utilizing multiple queues to perform more advanced functions such as prioritization.

Priority Queuing

This topic describes the priority queuing algorithm.

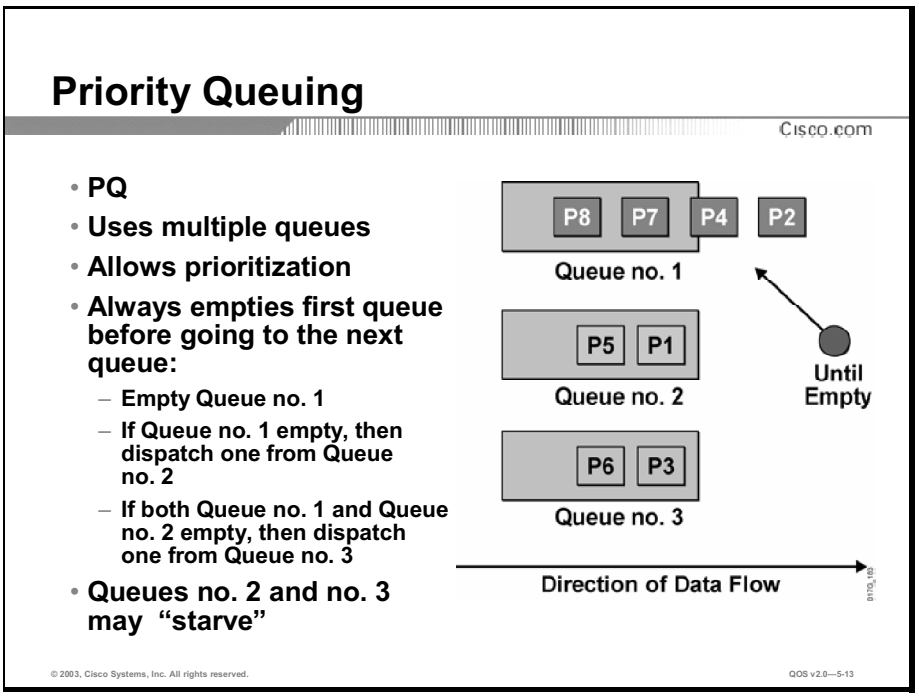

The PQ algorithm is also quite simple.

Each packet is assigned a priority and placed into a hierarchy of queues based on priority. When there are no more packets in the highest queue, the next-lower queue is serviced.

Then, packets are dispatched from the next-highest queue until either the queue is empty or another packet arrives for a higher PQ.

Only when all higher-priority queues are empty will packets be dispatched from a lower queue.

If a packet arrives for a higher queue, the packet from the higher queue is dispatched before any packets in lower-level queues.

The problem with PQ is that queues with lower priority can "starve" if a steady stream of packets continues to arrive for a queue with a higher priority. Packets waiting in the lower-priority queues may never be dispatched.

Round Robin

This topic describes the round-robin queuing algorithm.

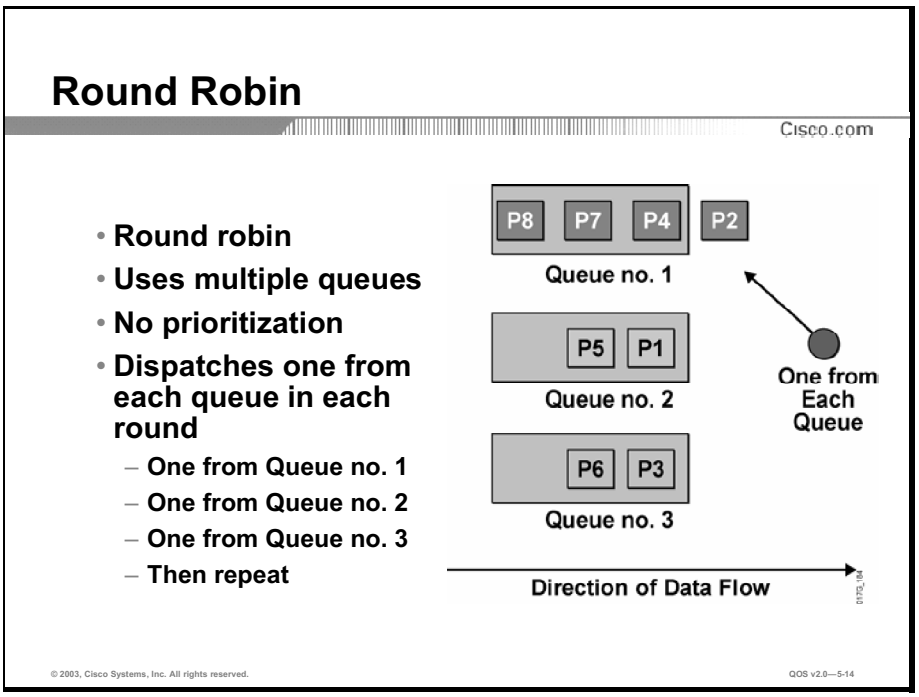

With round-robin queuing, one packet is taken from each queue and then the process repeats.

If all packets are the same size, all queues share the bandwidth equally. If packets being put into one queue are larger, that queue will receive a larger share of bandwidth.

No queue will “starve” with round robin as they all receive an opportunity to dispatch a packet every round.

A limitation of round robin is the inability to prioritize traffic.

Weighted Round Robin

This topic describes the WRR queuing algorithm.

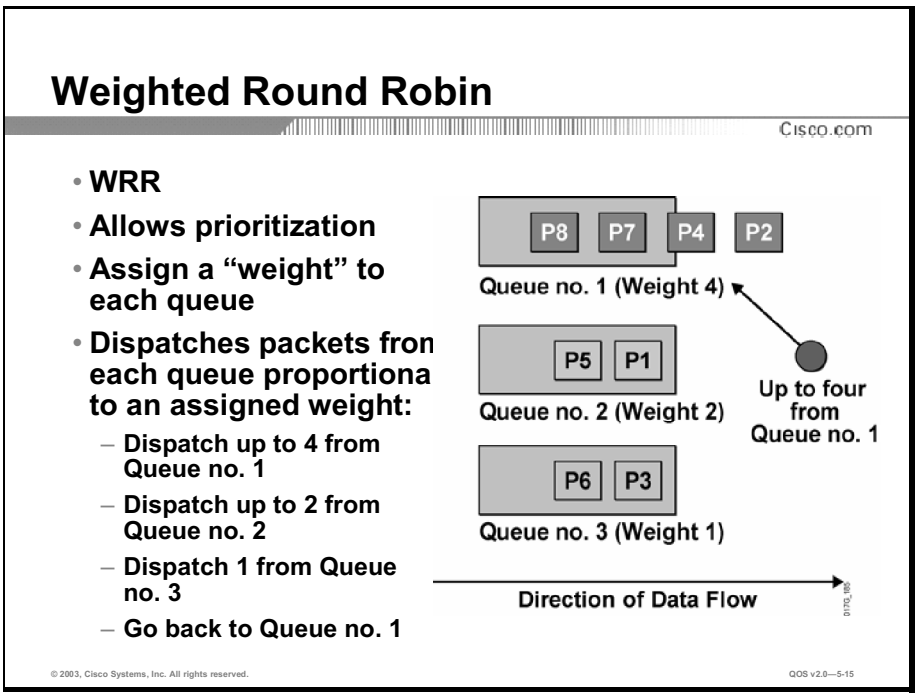

The WRR algorithm was developed to provide prioritization capabilities for round robin.

In WRR, packets are assigned a class (voice, file transfer, and so on) and placed into the queue for that class of service. Packets are accessed round-robin style, but queues can be given priorities called "weights." For example, in a single round, four packets from a high-priority class might be dispatched, followed by two from a middle-priority class, and then one from a low-priority class.

Some implementations of the WRR algorithm will dispatch a configurable number of bytes during each round.

Weighted Round Robin (Cont.)

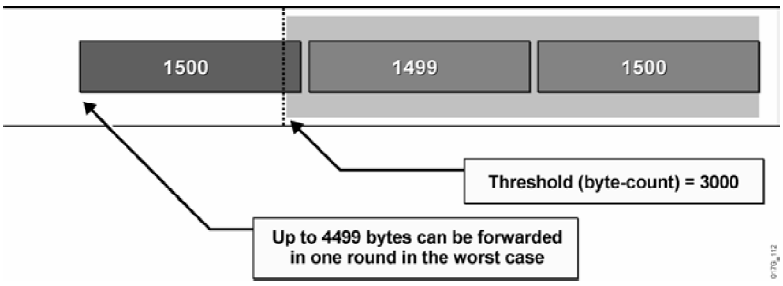

- **Problem with WRR**
  - **Some implementations of WRR dispatch a configurable number of bytes (threshold) from each queue for each round—several packets can be sent in each turn.**
  - **The router is allowed to send the entire packet even if the sum of all bytes is more than the threshold.**

Some implementations of the weighted round robin algorithm provide prioritization by dispatching a configurable number of bytes each round rather than a number of packets (Cisco custom queuing [CQ] mechanism is an example of this implementation).

The figure illustrates the worst-case scenario of the WRR algorithm where the following parameters were used to implement WRR queuing on an interface:

- Maximum Transmission Unit (MTU) of the interface is 1500 bytes.
- The byte-count to be sent each round for the queue is 3000 (twice the MTU).

The example shows how the router first sent two packets with a total size of 2999 bytes. Because this is still within the limit (3000), the router can send the next packet (MTU-sized). The result was that the queue received almost 50 percent more bandwidth in this round than it should.

This is one of the drawbacks of WRR queuing—it does not allocate bandwidth accurately.

The limit or weight of the queue is configured in bytes. The accuracy of WRR queuing depends on the weight (byte-count) and the MTU.

If the ratio between the byte-count and the MTU is too small, WRR queuing will not allocate bandwidth accurately.

If the ratio between the byte-count and the MTU is too large, WRR queuing will cause long delays.

Deficit Round Robin

This topic describes the deficit round-robin queuing algorithm.

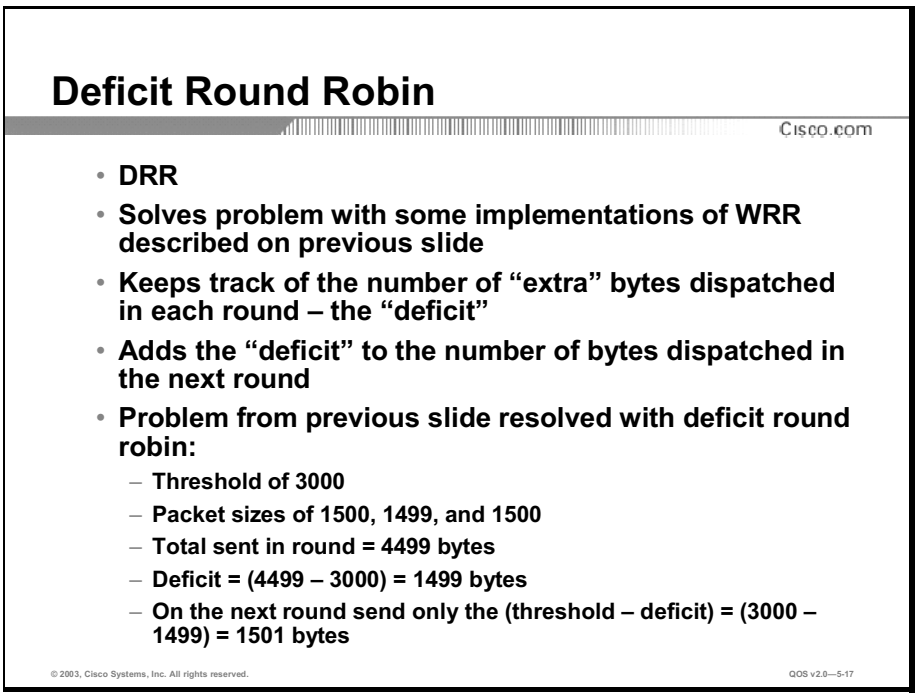

Deficit round robin is an implementation of the WRR algorithm developed to resolve the WRR problem described on the previous page. The Cisco modified deficit round robin (MDRR) method used on the Cisco 12000 series is an implementation of deficit round robin.

Deficit round robin uses a deficit counter to track the number of "extra" bytes dispatched over the number of bytes that was to be configured to be dispatched each round. During the next round, the number of "extra" bytes—the deficit—is effectively subtracted from the configurable number of bytes that are dispatched.

Summary

This topic summarizes the key points discussed in this lesson.

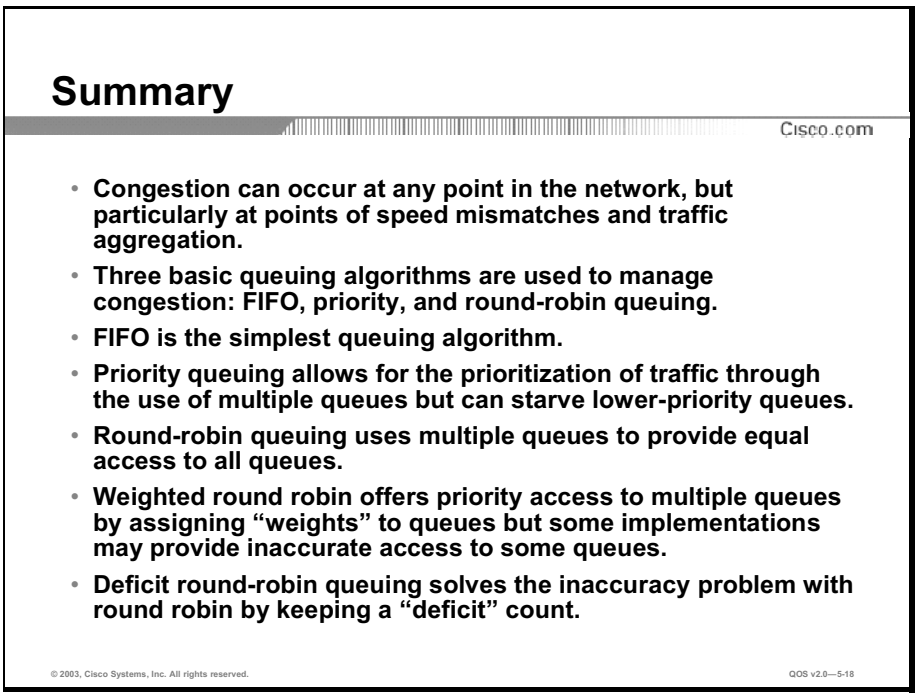

References

For additional information, refer to these resources:

- To learn more about congestion and queuing, refer to "Understanding Delay in Packet Voice Networks" at the following URL: http://www.cisco.com/en/US/tech/tk652/tk698/technologies_white_paper09186a00800a8993.shtml
- To learn more about congestion and queuing, refer to "Understanding Jitter in Packet Voice Networks (Cisco IOS Platforms)" at the following URL: http://www.cisco.com/en/US/tech/tk652/tk698/technologies_tech_note09186a00800945df.shtml

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) Which three of the following represent three likely points in a network where congestion would occur? (Choose three.)

A) points of aggregation

B) points of confluence

C) points of convergence

D) points of speed mismatches

Q2) Which of the following is the simplest queuing algorithm?

A) priority queuing

B) first-in, first-out

C) weighted round robin

D) round robin

Q3) Which of the following queuing algorithms would be most likely to “starve” lower-priority queues?

A) priority queuing

B) weighted round robin

C) round robin

D) deficit round robin

Q4) Which of the following queuing algorithms would be most likely to dispatch an equal number of packets from each queue?

A) round robin

B) priority queuing

C) weighted round robin

D) deficit round robin

Q5) When WRR is configured on a switch with four transmit queues, given that weights have been assigned to each of the queues as follows and that all queues are full, how many packets from queue 4 would be dispatched every time a packet from queue 2 is dispatched?

Queue	Weight
4	8
3	4
2	2
1	1

A) 2

B) 4

C) 16

D) 24

Q6) Given that a deficit round-robin queue is configured to dispatch 4000 bytes each round and it has just dispatched 4500 bytes, what will be the maximum number of bytes the queue will try to dispatch during the next round?

E) 4000

F) 4500

G) 3500

H) 7500

Quiz Answer Key

Q1) A, B, D

Relates to: Congestion and Queuing

Q2) B

Relates to: Queuing Algorithms

Q3) A

Relates to: Priority Queuing

Q4) A

Relates to: Round Robin

Q5) B

Relates to: Weighted Round Robin

Q6) C

Relates to: Deficit Round Robin

Queuing Implementations

Overview

Queuing technologies are one of the primary ways to manage congestion in a network. Network devices handle an overflow of arriving traffic by using a queuing algorithm to sort traffic and determine a method of prioritizing the traffic onto an output link. Each queuing algorithm was designed to solve a specific network traffic problem and has a particular effect on network performance. This lesson explains the underlying principles behind queuing on Cisco networking devices.

Relevance

In order to understand how an advanced queuing mechanism such as CBWFQ will operate on a Cisco router, it is important to first understand how queuing is implemented on Cisco network devices.

Objectives

Upon completing this lesson you will be able to describe hardware and software queuing on a network device. This includes being able to meet these objectives:

- Explain the components of hardware and software queuing systems on Cisco routers
- Explain the effects of tuning the size of the hardware queue on router and network performance
- Describe how congestion affects software interfaces on Cisco routers
- List and describe the basic queuing mechanisms available in Cisco IOS

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Basic knowledge of internetworking with TCP/IP concepts

Outline

The outline lists the topics included in this lesson.

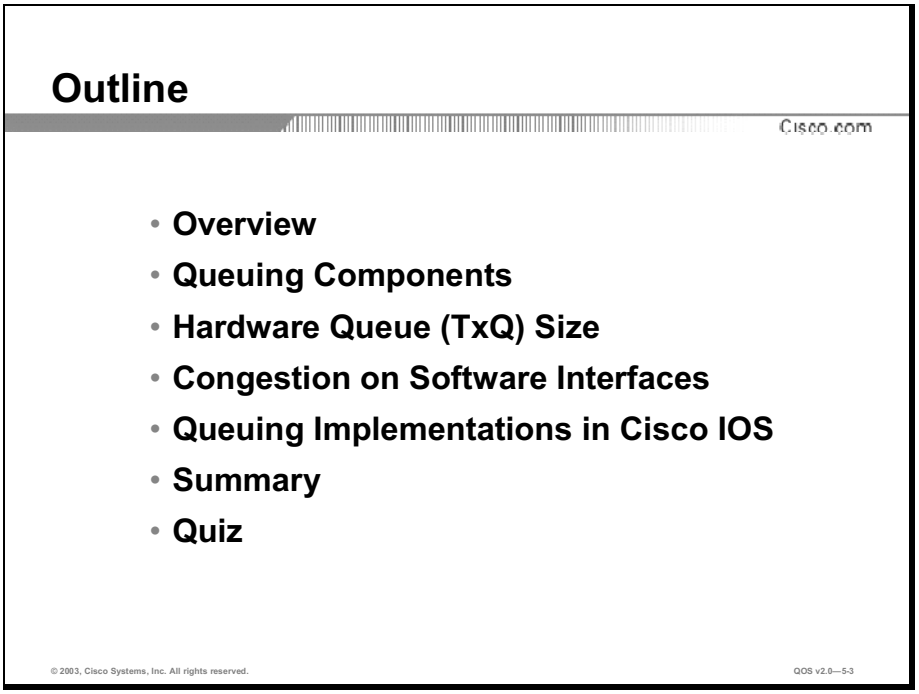

Queuing Components

This topic describes the primary components of a queuing mechanism.

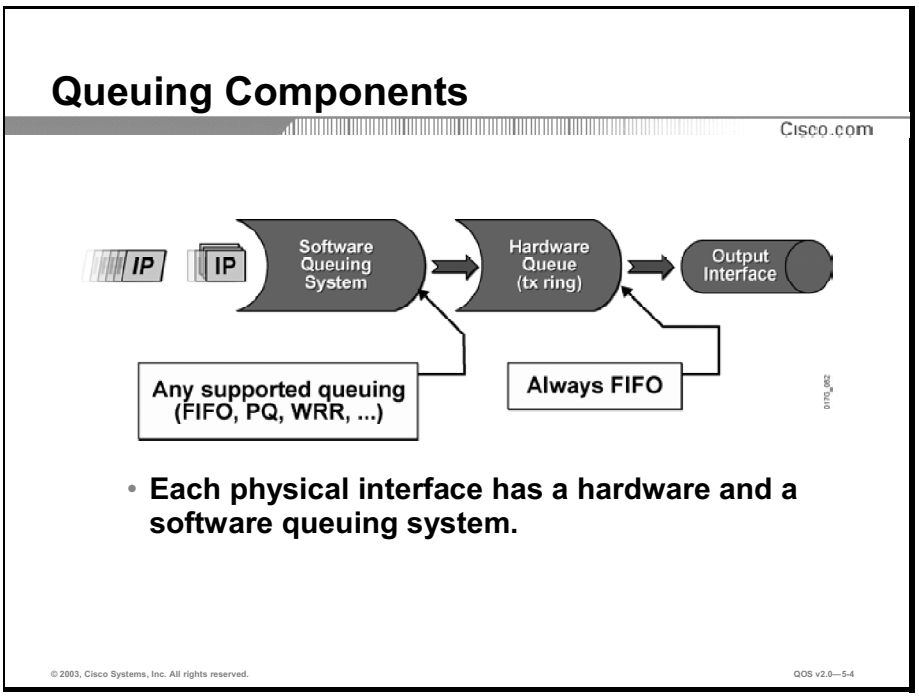

Queuing on routers is necessary to accommodate bursts when the arrival rate of packets is greater than the departure rate, usually because of one the following two reasons:

- Input interface is faster than the output interface
- Output interface is receiving packets coming in from multiple other interfaces

Initial implementations of queuing used a single FIFO strategy. More complex queuing mechanisms were introduced when special requirements need routers to differentiate between packets of different importance.

Queuing was split into two parts:

- **Hardware queue:** Uses FIFO strategy, which is necessary for the interface drivers to transmit packets one by one. The hardware queue is sometimes referred to as the transmit queue (TxQ).
- **Software queue:** Schedules packets into the hardware queue based on the QoS requirements.

Queuing Components (Cont.)

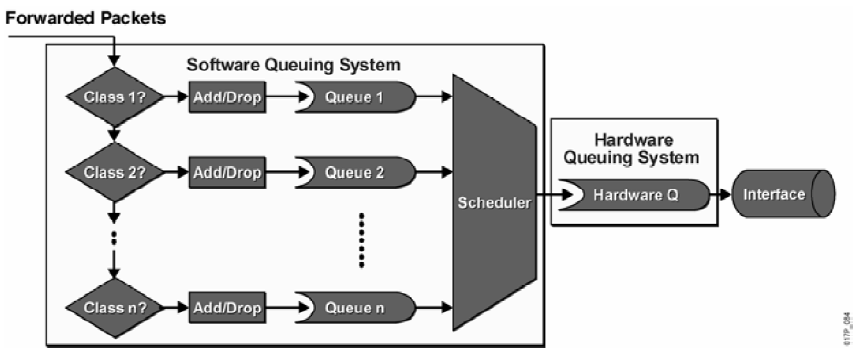

- **The hardware queuing system always uses FIFO queuing.**
- **The software queuing system can be selected and configured depending on the platform and Cisco IOS version.**

The figure illustrates the actions that have to be taken before a packet can be transmitted:

- Most queuing mechanisms include classification of packets.
- After a packet is classified, a router has to determine whether it can put the packet into the queue or it has to drop the packet. Most queuing mechanisms will drop a packet only if the corresponding queue is full (tail drop). Some mechanisms use a more intelligent dropping scheme (WFQ) or a random dropping scheme (weighted random early detection [WRED]).
- If the packet is allowed to be enqueued it will be put into the FIFO queue for that particular class.
- Packets are then taken from the individual per-class queues and put into the hardware queue.

The Software Queue

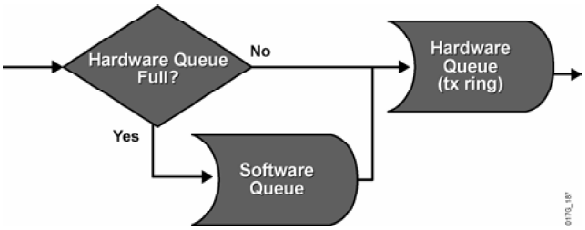

- **Generally, a full hardware queue indicates interface congestion and software queuing is utilized to manage it.**
- **When a packet is being forwarded, the router will bypass the software queue if the hardware queue has space in it (no congestion).**

The implementation of software queuing was optimized for periods when the interface is not congested. The software queuing system is bypassed whenever there is no packet in the software queue and there is room in the hardware queue.

The software queue is, therefore, only used when data must wait to be placed into the hardware queue.

Hardware Queue (TxQ) Size

This topic explains the significance of the size of the hardware queue.

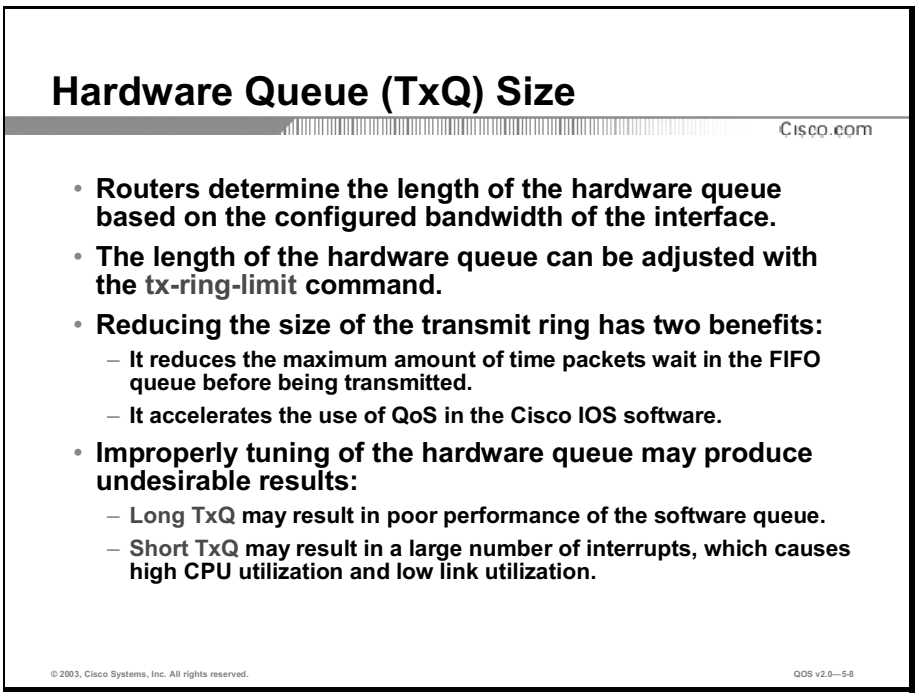

The double queuing strategy (software and hardware queue) has its impacts on the results of overall queuing. Software queues serve a valuable purpose. If the hardware queue is too long, it will contain a large number of packets scheduled in the FIFO fashion. A long FIFO hardware queue most likely defeats the quality of service (QoS) design that required a certain complex software queuing system (for example, CQ).

So why use the hardware queue at all? Or why not just set its length to one? That would force all packets to go through the software queue and be scheduled one by one to the interface for transmission. This approach has the following drawbacks:

- Each time a packet is transmitted, the interface driver interrupts the CPU and requests more packets to be delivered into its hardware queue. Some queuing mechanisms have complex scheduling that takes time to deliver more packets. The interface does not send anything during that time (link utilization is decreased) if the hardware queue is empty because its maximum size is one.
- The CPU schedules packets one by one instead of many at the same time (in the same interrupt interval). This increases the CPU utilization.

Choosing the appropriate length of the hardware queue is very important. The default TxQ size is determined by the Cisco IOS software, based on the bandwidth of the media, and should be fine for most queuing implementations. Some platforms and QoS mechanisms will automatically adjust the TxQ size to an appropriate value. Faster interfaces have longer hardware queues because they produce less delay. Slower interfaces have shorter hardware queues to prevent too much delay in the worst-case scenario where the entire hardware queue is full of MTU-sized packets.

Note: Refer to Cisco IOS software configuration documentation for more information.

The transmit ring serves as a staging area for packets in line to be transmitted. The router needs to enqueue a sufficient number of packets on the transmit ring and ensure that the interface driver has packets with which to fill available cell timeslots.

The primary reason to tune the transmit ring is to reduce latency caused by queueing. On any network interface, queueing forces a choice between latency and the amount of burst that the interface can sustain. Larger queue sizes sustain longer bursts while increasing delay. Tune the size of a queue when you think traffic is experiencing unnecessary delay.

The size of the transmit ring must be small enough to avoid introducing latency because of queuing and it must be large enough to avoid drops and a resulting impact to TCP-based flows. Queuing on the transmit ring introduces a serialization delay that is directly proportional to the depth of the ring. An excessive serialization delay can impact latency budgets for delay-sensitive applications such as voice. Thus, Cisco recommends reducing the size of the transmit ring for VCs carrying voice. Select a value based on the amount of serialization delay, expressed in seconds, introduced by the transmit ring. Use the following formula:

```
((P*8)*D)/S

P = Packet size in bytes. Multiply by eight to convert to
bits.

D = Transmit-ring depth.

S = Speed of the VC in bps.
```

Note: IP packets on the Internet are typically one of three sizes: 64 bytes (for example, control messages), 1500 bytes (for example, file transfers), or 256 bytes (all other traffic). These values produce a typical overall Internet packet size of 250 bytes.

Congestion on Software Interfaces

This topic explains how congestion occurs on software interfaces.

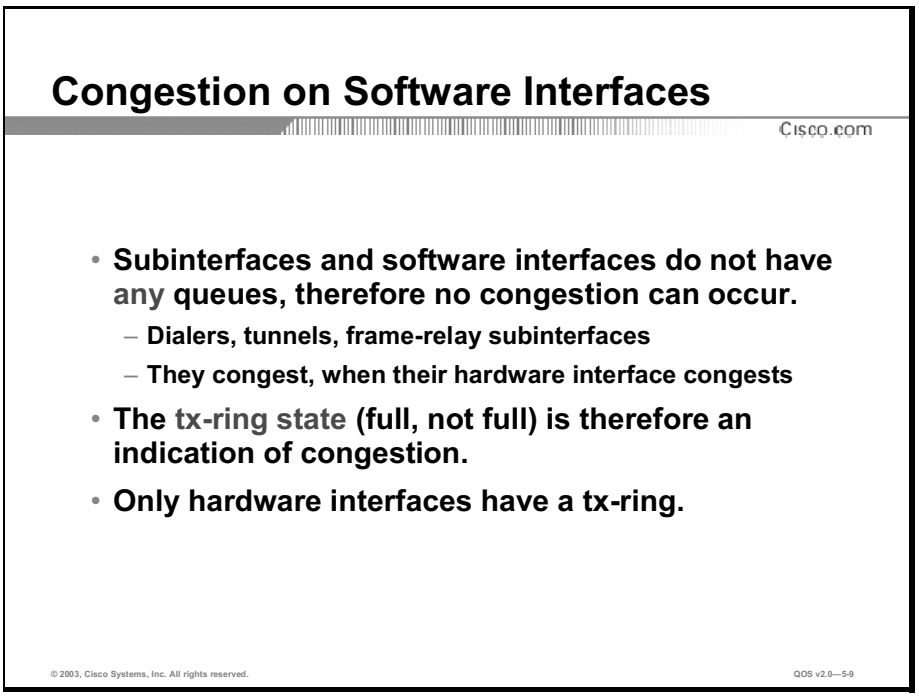

Subinterfaces and software interfaces do not have queues. Therefore, no congestion can occur. These interface types include dialers, tunnels, and frame-relay subinterfaces and will only congest when their hardware interface congests. The transmit (tx) ring state is an indication of congestion for software interfaces.

Queuing Implementations in Cisco IOS

This topic describes the basic software queuing technologies used on Cisco network devices.

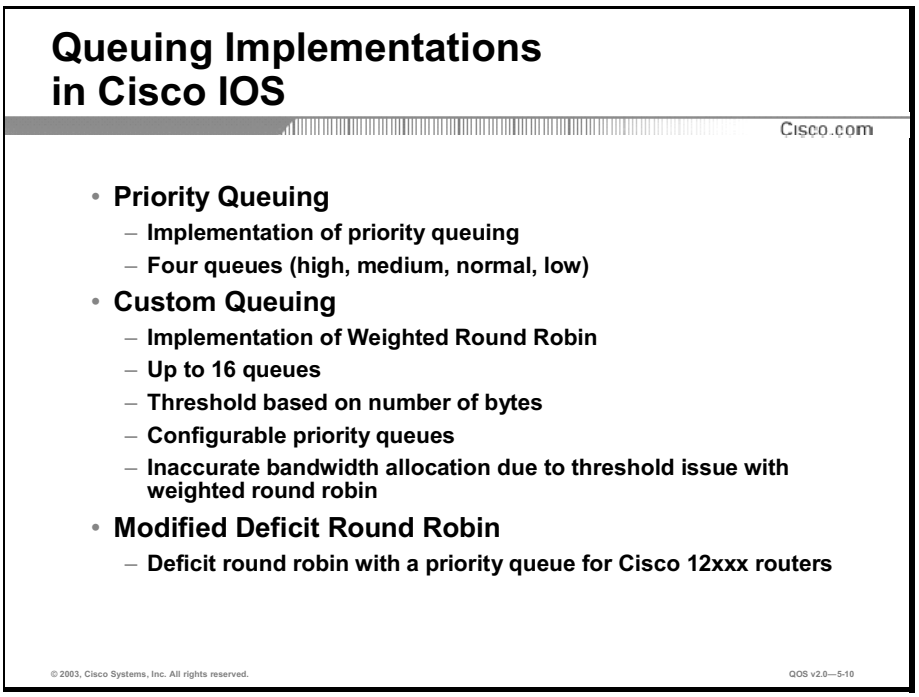

The figure lists some of the available software queuing technologies.

- PQ
  - A Cisco implementation of the priority queuing algorithm
  - Allows four queues to be used for prioritizatoin (high, medium, normal, low)
  - Allows for a variety of classification including source IP address, destination IP address, IP precedence, and DSCP
- Custom queuing
  - A Cisco implementation of WRR
  - Allows up to 16 queues to be used for traffic classification
  - Allows for a variety of classification including: source IP address, destination IP address, IP precedence, and DSCP
  - Tail drop is used within each individual queue
- MDRR
  - A Cisco implementation of deficit round robin
  - Available only on the Cisco 12000 series routers

Summary

This topic summarizes the key points discussed in this lesson.

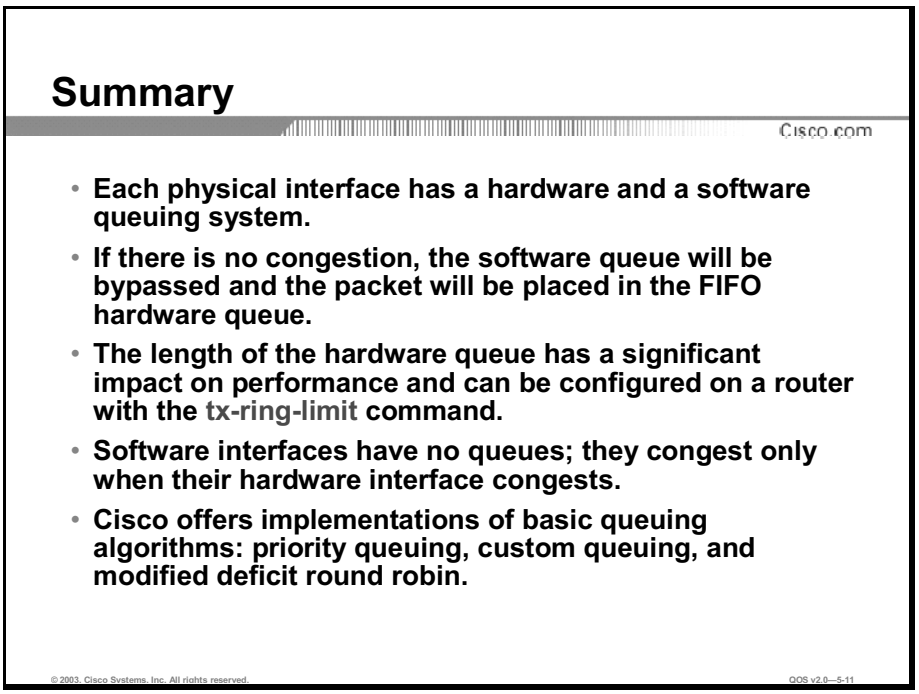

References

For additional information, refer to these resources:

- To learn more about congestion and queuing, refer to “Understanding Delay in Packet Voice Networks” at the following URL: http://www.cisco.com/en/US/tech/tk652/tk698/technologies_white_paper09186a00800a8993.shtml
- To learn more about congestion and queuing, refer to “Understanding Jitter in Packet Voice Networks (Cisco IOS Platforms)” at the following URL: http://www.cisco.com/en/US/tech/tk652/tk698/technologies_tech_note09186a00800945df.shtml

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) Every physical interface must have which two of the following? (Choose two.)

A) priority queuing

B) custom queuing

C) software queuing system

D) hardware queuing system

Q2) Which two would be the likely results of a hardware queue with too short of a TxQ? (Choose two.)

A) jitter

B) low link utilization

C) high CPU utilization

D) poor software queue performance

Q3) How many queues does a software interface have by default?

A) 0

B) 1

C) 2

D) 4

Q4) What would be the likely result of a hardware queue with too long of a TxQ?

A) jitter

B) low link utilization

C) high CPU utilization

D) poor software queue performance

Q5) Which of the following is a Cisco implementation of the weighted round robin algorithm?

A) custom queuing (CQ)

B) low latency queuing (LLQ)

C) weighted fair queuing (WFQ)

D) class-based weighted fair queuing (CBWFQ)

Quiz Answer Key

Q1) C, D

Relates to: Queuing Components

Q2) B, C

Relates to: Hardware Queue (TxQ) Size

Q3) A

Relates to: Congestion on Software Interfaces

Q4) D

Relates to: Hardware Queue (TxQ) Size

Q5) A

Relates to: Queuing Implementations in Cisco IOS

FIFO and WFQ

Overview

FIFO and WFQ are the two primary default queuing mechanisms that are implemented on Cisco routers. WFQ was developed to resolve some of the problems resulting from the use of basic queuing methods such as queue starvation, delay, and jitter. WFQ dynamically divides available bandwidth by a calculation based on the total number of flows and the weight of each given flow. Bandwidth cannot be guaranteed, as the number of flows are constantly changing and thus so is the allocated bandwidth to each flow.

Relevance

WFQ is a key technology for ensuring QoS in a converged network and is used as a key element of the more advanced queuing methods.

Objectives

Upon completing this lesson, given a network with suboptimal QoS performance, you will be able to configure WFQ to manage congestion. This includes being able to meet these objectives:

- Describe the FIFO queuing mechanism
- Give a detailed explanation of WFQ using a block diagram
- Identify the parameters on which WFQ can classify traffic
- Explain the insertion and drop policy used by WFQ using a block diagram
- Explain how finish time is calculated based on weight and used in the operation of WFQ
- Describe the benefits and drawbacks of using WFQ to implement QoS
- Identify the Cisco IOS commands required to configure WFQ on a Cisco router
- Identify the Cisco IOS commands required to monitor WFQ on a Cisco router

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Basic knowledge of internetworking with TCP/IP concepts and knowledge of basic Cisco IOS commands

Outline

The outline lists the topics included in this lesson.

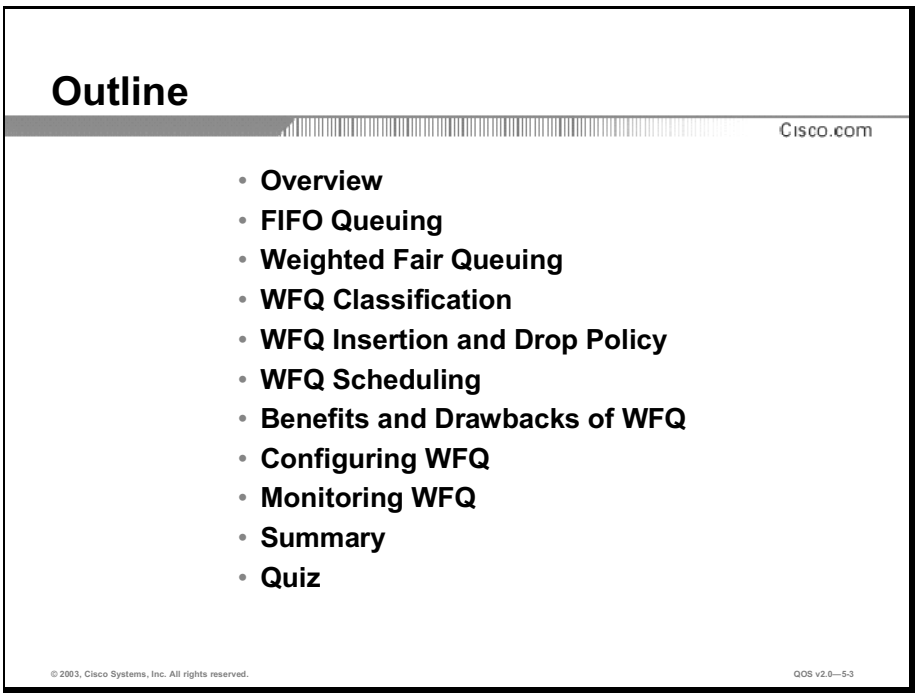

FIFO Queuing

This topic describes FIFO queuing.

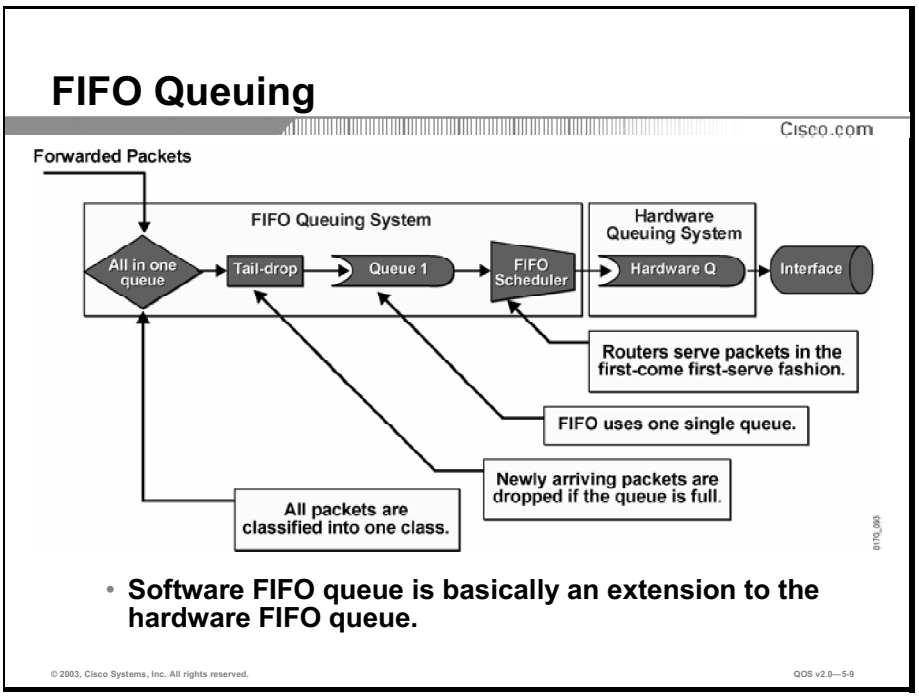

FIFO queuing has no classification because all packets belong to the same class. Packets are dropped when the output queue is full (tail drop). The scheduler services packets in the order they arrived.

Software FIFO queue is basically an extension of the hardware FIFO queue.

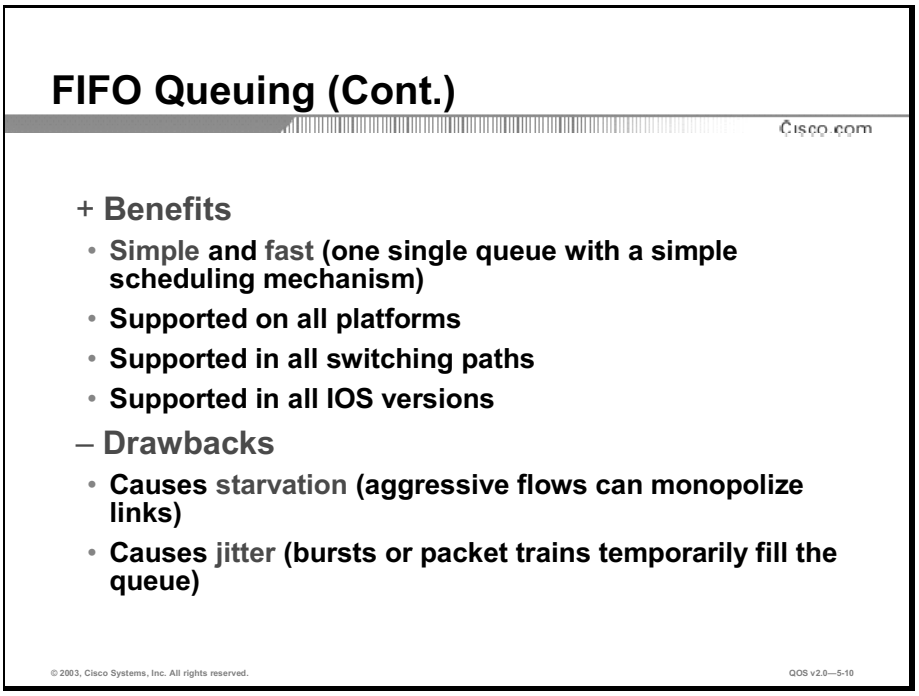

Although FIFO queuing might be regarded as the fairest queuing mechanism, it has a long list of drawbacks:

- FIFO is extremely unfair when an aggressive flow is contesting with a fragile flow. Aggressive flows send a large number of packets, many of which are dropped. Fragile flows send a modest amount of packets and most of them are dropped because the queue is always full because of the aggressive flow. This type of behavior is called starvation.
- Short or long bursts cause a FIFO queue to fill. Packets entering an almost full queue have to wait a long time before they can be transmitted. Another time, the queue might be empty causing packets of the same flow to experience almost no delay. Variation in delay is called jitter.

In spite of all the drawbacks, FIFO is still the most used queuing mechanism because of the following benefits:

- It is simple and fast. Most high-end routers with fast interfaces are not really challenged by the drawbacks mentioned earlier. Furthermore, routers are not capable of complex classification and scheduling when they have to process a large number of packets-per-second. FIFO is, therefore, the most suitable queuing mechanism on these platforms.
- It is supported on all platforms.
- FIFO queuing is supported in all versions of Cisco IOS.

Weighted Fair Queuing

This topic explains the purpose and function of WFQ.

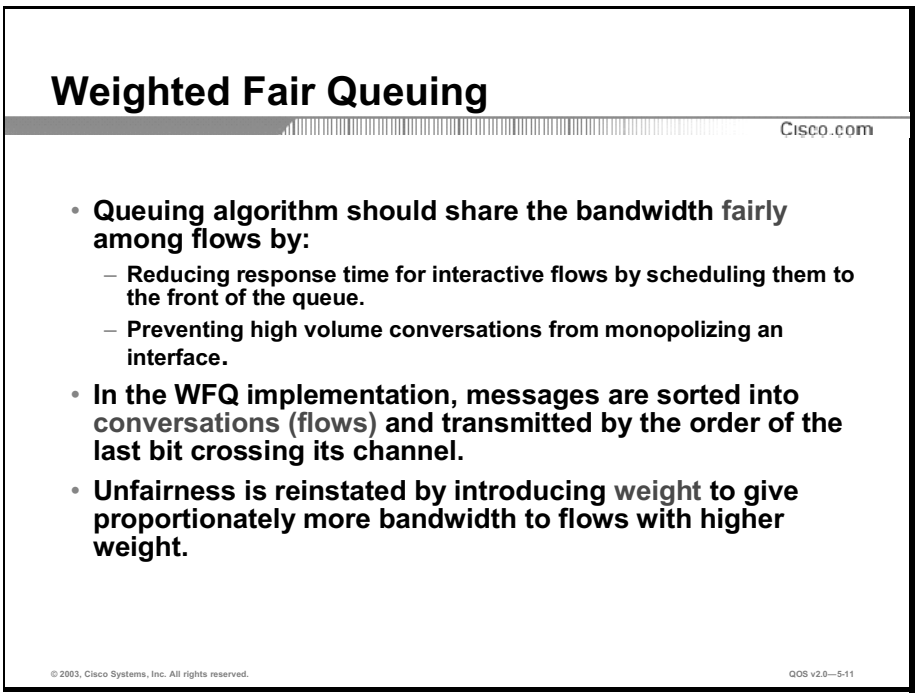

WFQ was introduced as a solution to the problems of the following queuing mechanisms:

- FIFO queuing causes starvation, delay, and jitter.
- PQ causes starvation of other lower-priority classes and suffers from all FIFO problems within each of the four queues.
- CQ causes long delays and also suffers from all FIFO problems within each of the 16 queues.

The idea of WFQ is to:

- Have a dedicated queue for each flow (no starvation, delay, or jitter within the queue).
- Fairly and accurately allocate bandwidth among all flows (minimum scheduling delay, guaranteed service).
- Use IP precedence as weight when allocating bandwidth.

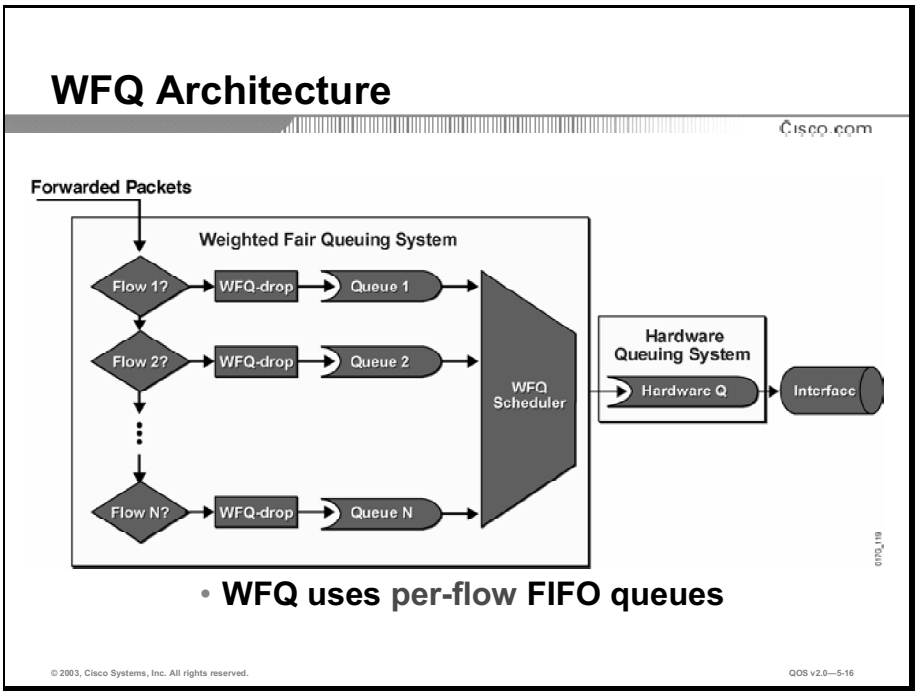

WFQ uses automatic classification. Manually defined classes are not supported.

WFQ dropping is not a simple tail drop. WFQ drops packets of the most aggressive flows.

WFQ scheduler is a simulation of a time-division multiplexing (TDM) system. The bandwidth is fairly distributed to all active flows.

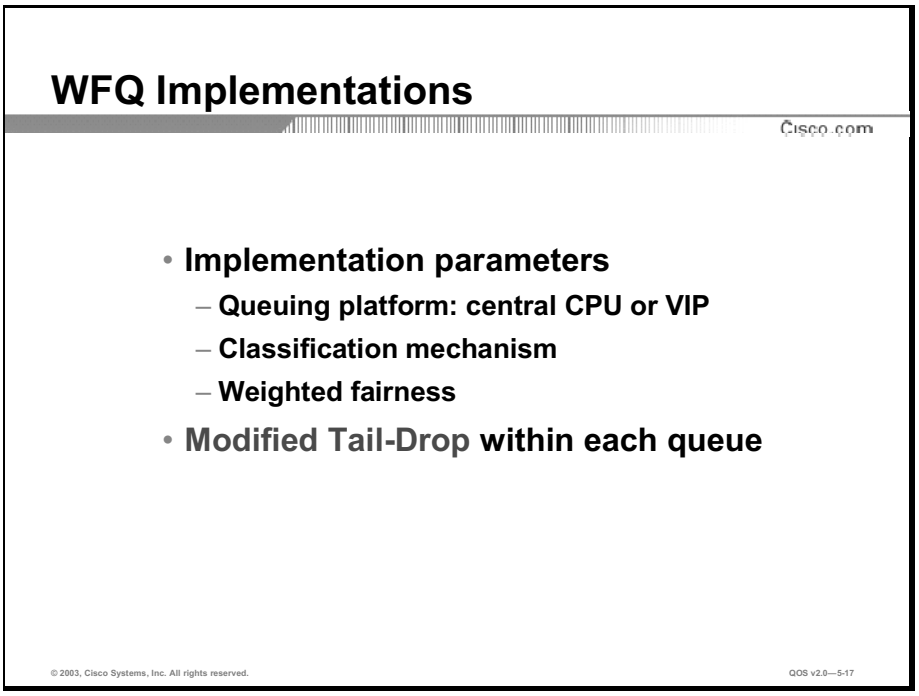

WFQ is supported on most Cisco routers as well as on Versatile Interface Processors (VIP). The implementation of WFQ on the VIP differs slightly from the one discussed in this lesson in the following ways:

- Classification identifies a flow and assigns a queue to the flow.
- Weight is used for scheduling to give proportionately more bandwidth to flows with a higher IP precedence.
- Tail-dropping scheme is improved to drop packets of the most aggressive flows.

WFQ Classification

This topic explains how classification is accomplished with WFQ.

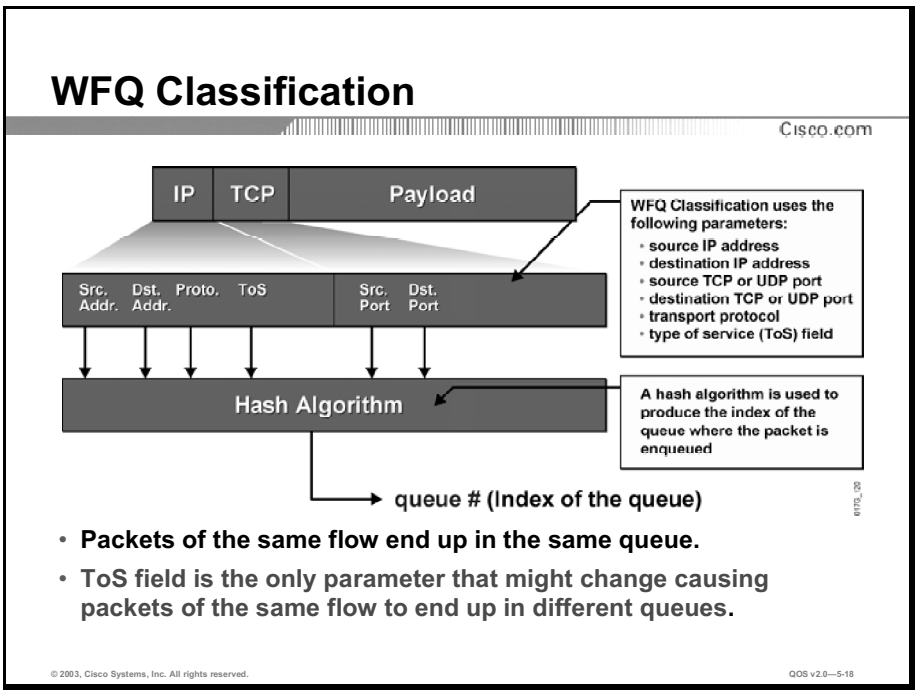

WFQ classification has to identify individual flows. (The term *conversation* is also used to signify flows.) A flow is identified based on the following information taken from the IP header and the TCP or UDP headers:

- Source IP address
- Destination IP address
- Protocol number (identifying TCP or User Datagram Protocol (UDP)
- Type of service field
- Source TCP/UDP port number
- Destination TCP/UDP port number

The following parameters are usually fixed for a single flow, although there are some exceptions:

- A QoS design could mark packets with different IP precedence values even if they belong to the same flow. This kind of behavior should be avoided when using WFQ.
- Some applications change port numbers (for example, TFTP).

If packets of the same flow do not have the same parameters (for example, a different type of service [ToS] field), the packets can end up in different queues and reordering can occur.

The parameters are used as input for a hash algorithm that produces a fixed-length number that is used as the index of the queue.

WFQ Classification (Cont.)

Cisco.com

- **Fixed number of per-flow queues is configured.**
- **A hash function is used to translate flow parameters into queue number.**
- **System packets (8 queues) and RSVP flows (if configured) are mapped into separate queues.**
- **Two or more flows could map into the same queue, resulting in lower per-flow bandwidth.**
- **Important: the number of queues configured has to be larger than the expected number of flows.**

WFQ uses a fixed number of queues. The hash function is used to assign a queue to a flow. There are eight additional queues for system packets and optionally up to 1000 queues for Resource Reservation Protocol (RSVP) flows. The number of dynamic queues WFQ uses by default is based on the interface bandwidth. Using the default interface bandwidth, WFQ uses 256 dynamic queues by default. The number of queues can be configured in the range between 16 and 4096 (the number must be a power of 2). The default number of dynamic queues for different interface bandwidths is shown in the table.

Bandwidth Range	Number of Dynamic Queues
Less than or equal to 64 kbps	16
More than 64 kbps and less than or equal to 128 kbps	32
More than 128 kbps and less than or equal to 256 kbps	64
More than 256 kbps and less than or equal to 512 kbps	128
More than 512 kbps	256

If there are a large number of concurrent flows it is very likely that two flows could end up in the same queue. It is recommended to have several times as many queues as there are flows (on average). This may not be possible in larger environments where the number of concurrent flows is in thousands.

The probability of two flows ending up in the same flow could be calculated using the following formula:

$$P = 1 - \frac{Queues!}{Queues^{Flows} \cdot (Queues - Flows)!}$$

The following table lists the probability values for 3 sizes of the WFQ system (64, 128, and 256 queues), with the number of concurrent flows from 5 to 40. The table shows the probability values for three sizes of the WFQ system.

Flows	64 Queues	128 Queues	256 Queues
5	15%	8%	4%
10	52%	30%	16%
15	83%	57%	34%
20	96%	79%	53%
25	100%	92%	70%
30	100%	98%	83%
35	100%	99%	91%
40	100%	100%	96%

Below is the sample calculation of the probability value for 5 flows and 64 queues:

```
Flows: 5
Queues: 64

Probability = 1 – ((64!) / ((64^5) * (59!)))
            = 1 – ((64 * 63 * 62 * 61 * 60) / (64 * 64 * 64 *
64 * 64)
            = 1 – 0.852105618
            = 0.147894382 or 14.7% (15% rounded off)
```


WFQ Insertion and Drop Policy

This topic explains WFQ insertion and drop policy.

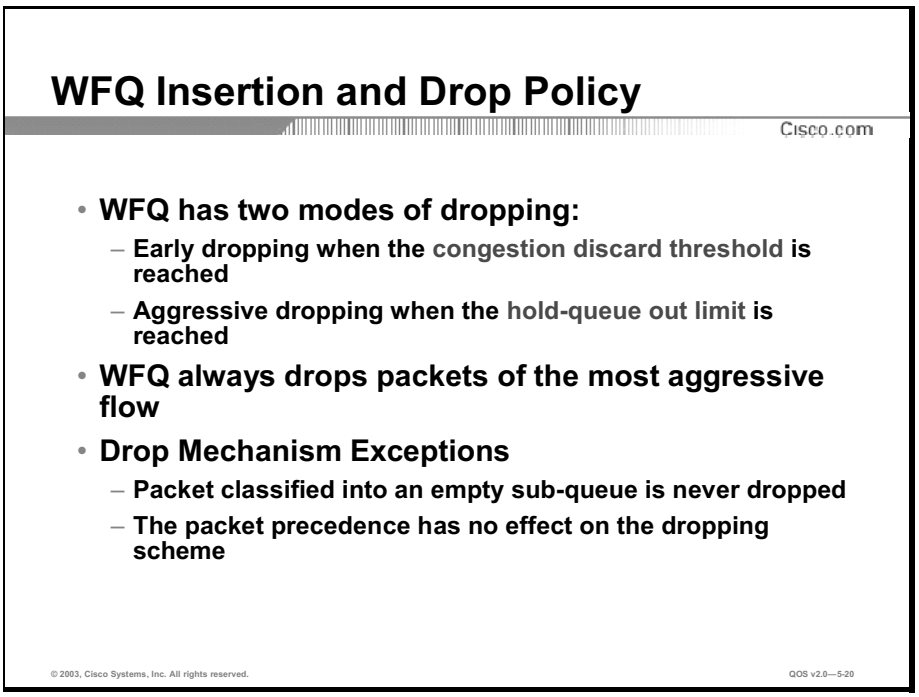

WFQ uses two parameters that affect the dropping of packets.

- The congestive discard threshold (CDT) is used to start dropping packets of the most aggressive flow, even before the hold-queue limit is reached.
- The hold-queue out limit defines the total maximum number of packets that can be in the WFQ system at any time.

There are two exceptions to the WFQ insertion and drop policy as follows:

- If the WFQ system is above the CDT limit the packet is still enqueued if the per-flow queue is empty.
- The dropping strategy is not directly influenced by IP precedence.

WFQ Insertion and Drop Policy (Cont.)

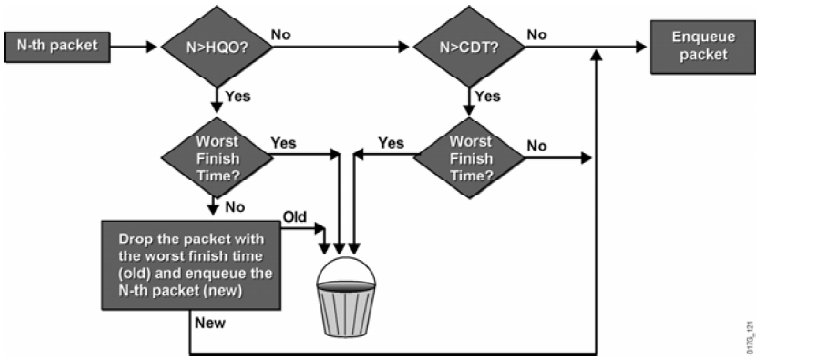

- **HQO (hold-queue out limit) is the max. number of packets that the WFQ system can hold.**
- **CDT is the threshold when WFQ starts dropping packets of the most aggressive flow.**
- **N is the number of packets in the WFQ system when the N-th packet arrives.**

The figure illustrates the dropping scheme of WFQ. The process can be organized into the following steps:

Step 1 Drop the new packet if the WFQ system is full (hold-queue limit reached) and the new packet has the worst finish time (the last in the entire system).

Step 2 Drop the packet with the worst finish time in the WFQ system if the system is full. Enqueue the new packet.

Step 3 When the WFQ system is above the CDT limit, a new packet is dropped if it is the last in the WFQ system, even though the WFQ system is still within the hold-queue limit.

Step 4 When the WFQ system is above the CDT limit, and if a new packet would not be the last in the WFQ system, the new packet can be enqueued and no other packet is dropped.

WFQ Scheduling

This topic describes scheduling in WFQ.

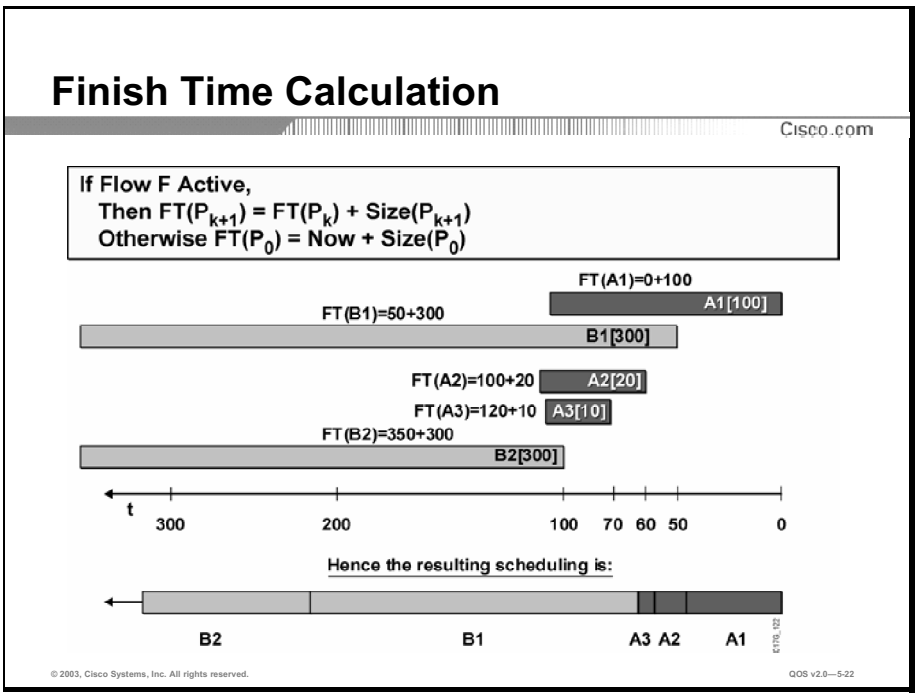

The length of queues (for scheduling purposes) is not in packets but in the time it would take to transmit all the packets in the queue. The end result is that WFQ adapts to the number of active flows (queues) and allocates equal amounts of bandwidth to each flow (queue).

The side effect is that flows with small packets (usually interactive flows) get a much better service because they do not need a lot of bandwidth. They, however, need low-delay, which they get because small packets have a low finish time.

The figure illustrates how two queues (queue A and queue B) are contesting for link bandwidth. For this example, assume the time units are in ms and time T (value 0 is used in the figure) is the starting point.

Queue A is receiving packets in the following order and the following times:

- Packet A1 arrives at time T + 0 ms and would require 100 ms to be transmitted.
- Packet A2 arrives at time T + 60 ms (the input interface is obviously faster than the output interface because the arrival time of packet A2 is before the finish time of packet A1) and would require 20 ms to be transmitted.
- Packet A3 arrives at time T + 70 ms (the input interface is obviously much faster than the output interface) and would require 10 ms to be transmitted.

Queue B is receiving packets in the following order and the following times:

- Packet B1 arrives at time T + 50 ms and would require 300 ms to be transmitted.
- Packet B2 arrives at time T + 100 ms and would also require 300 ms to be transmitted.

The finish time of packets in Queue A are:

- Packet A1 has a finish time which is the sum of the current time (because the queue was empty at the time of arrival) and the time it takes to transmit this packet (100 ms): FT_{A1} = 0 ms + 100 ms = 100 ms.
- Packet A2 has a finish time which is the sum of the finish time of the last packet in queue A (Packet A1) and the time it would take to transmit this packet (20 ms): FT_{A2} = 100 ms + 20 ms = 120 ms.
- Packet A3 has a finish time which is the sum of the finish time of the last packet in Queue A (Packet A2) and the time it would take to transmit this packet (20 ms): FT_{A3} = 120 ms + 10 ms = 130 ms.

The finish time for the packets in queue B are:

- Packet B1 has a finish time which is the sum of the current time (because the queue was empty at the time of arrival) and the time it takes to transmit this packet (300 ms): FT_{B1} = 50 ms + 300 ms = 350 ms.
- Packet B2 has a finish time which is the sum of the finish time of the last packet in queue B (Packet B1) and the time it would take to transmit this packet (300ms): FT_{B2} = 350 ms + 300 ms = 650 ms.

The packets are scheduled into the hardware queue (or the TxQ) in the ascending order of finish times:

1. A1 (100 ms)
2. A2 (120 ms)
3. A3 (130 ms)
4. B1 (350 ms)
5. B2 (650 ms)

Note: WFQ prevents reordering of packets within a single flow (conversation). Small packets are automatically preferred over large packets.

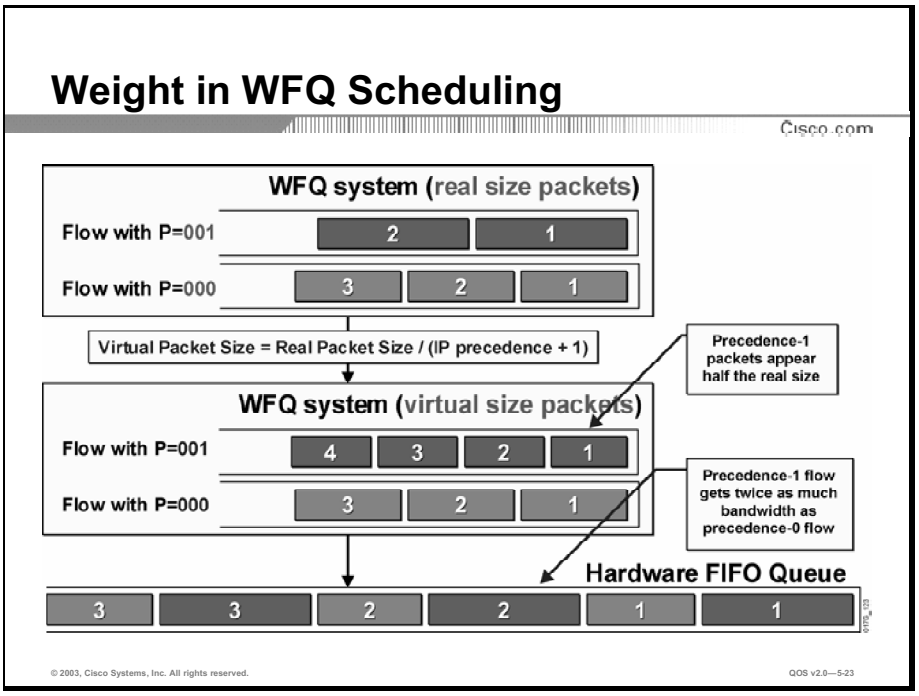

This figure introduces the weight into the finish time calculation. The time it takes to transmit the packet is divided by IP precedence increased by one (to prevent division by zero).

The WFQ implementation in Cisco routers was optimized in the following way:

- The real time it takes to transmit the packet is not relevant. The packet size can be used instead because it is proportional to the transmit time.
- The packet size is not divided by IP precedence (division is a CPU-intensive operation). Instead, the size is multiplied by a fixed value (one multiplication value for each IP precedence value).

Packets with IP precedence one appear half the size they really are. The result is that these packets receive twice as much bandwidth as packets with IP precedence zero.

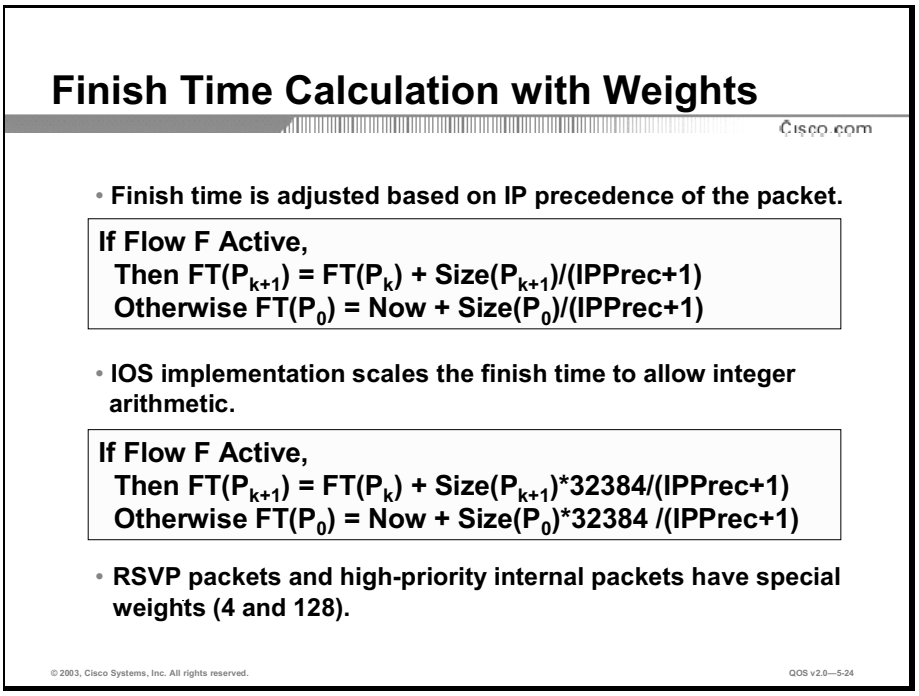

The first formula in the figure is the first optimisation where the finish time is really the sum of packet sizes divided by an increased IP precedence value.

The second formula shows further optimization where, instead of dividing, the packet size is multiplied by 32384/(IP precedence + 1). A number for each different IP precedence value is stored in a table and therefore, does not have to be calculated for each packet.

Packets belonging to RSVP flows and system packets have special low weights that guarantee them more bandwidth.

Note: Cisco IOS versions before 12.0(5)T use a new formula where the weight is calculated on the following formula: Weight = 4096 / (IP precedence +1).

The illustration shows the mapping between IP precedence values and WFQ weights.

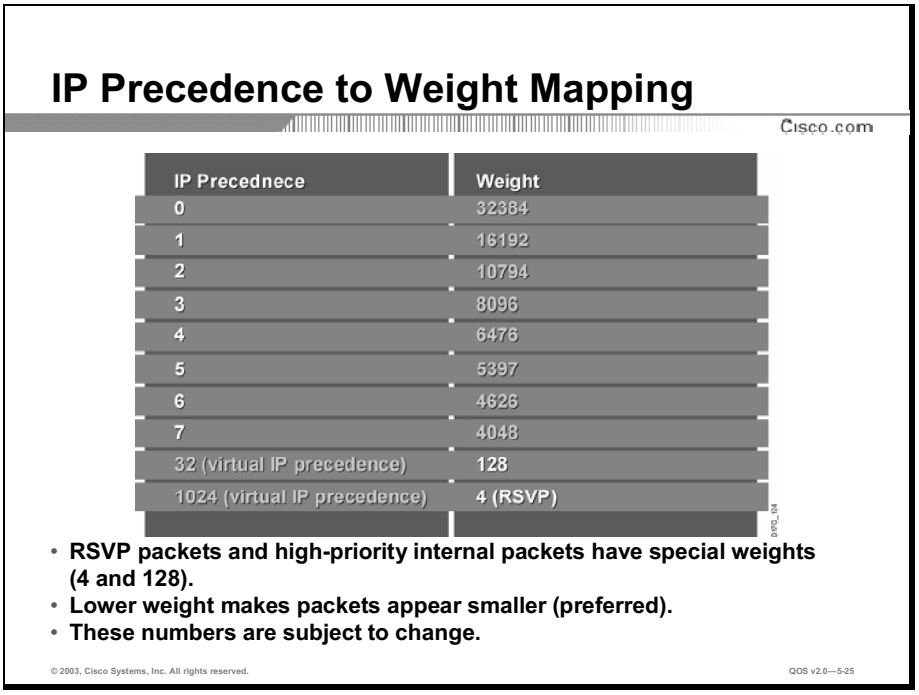

Note: These figures are subject to change. Refer to the Cisco IOS documentation for the latest information on WFQ weights.

The following case study is used to describe how packets are dropped in different situations.

WFQ Case Study

Cisco.com

- **WFQ system can hold a maximum of ten packets (hold-queue limit).**
- **Early dropping (of aggressive flows) should start when there are eight packets (congestive discard threshold) in the WFQ system.**

The WFQ system was reduced to a modest hold-queue limit of ten and a congestive discard threshold of eight.

WFQ Case Study Interface Congestion

Cisco.com

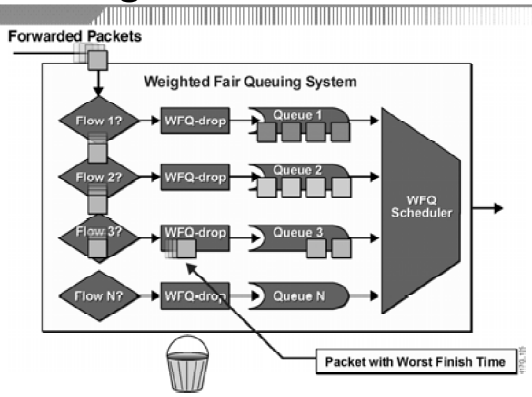

- **HQO (hold-queue out limit) is the maximum number of packets that the WFQ system can hold and HQO = 10.**

There are already ten packets in the WFQ system. The new packet would be the 11th and also the last in the entire WFQ system.

The new packet is dropped.

WFQ Case Study
Interface Congestion (Cont.)

Cisco.com

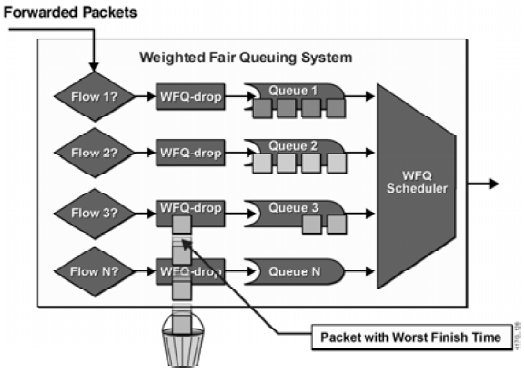

- **HQO (hold-queue out limit) is the maximum number of packets that the WFQ system can hold and HQO = 10.**
- **Absolute maximum (HQO=10) exceeded, new packet is the last in the TDM system and is dropped.**

WFQ Case Study
Flow Congestion

Cisco.com

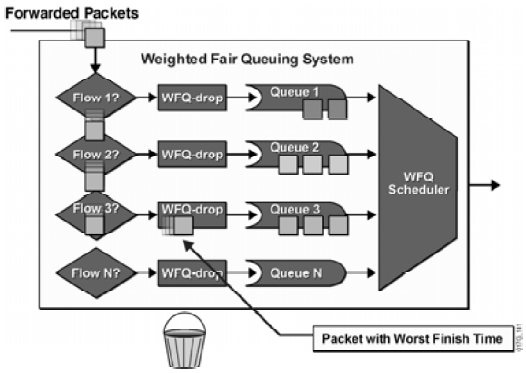

- **Early dropping (of aggressive flows) should start when there are eight packets (congestive discard threshold) in the WFQ system.**

This example illustrates how WFQ can drop packets even if the WFQ system is still within the hold-queue limit. The system, however, is above the CDT limit. In this case a packet can be dropped if it is the last in the system.

WFQ Case Study
Flow Congestion (Cont.)

Cisco.com

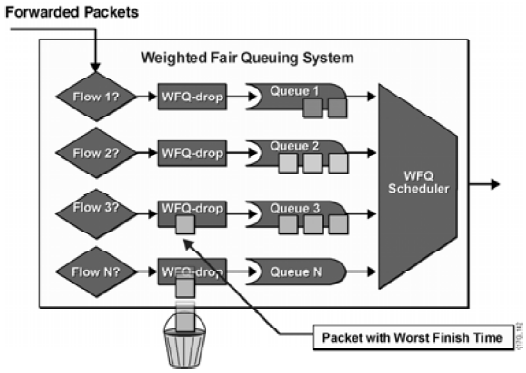

- **Early dropping (of aggressive flows) should start when there are eight packets (congestive discard threshold) in the WFQ system.**
- **CDT exceeded (CDT=8), new packet would be the last in the TDM system and is dropped.**

In this case, a packet can be dropped if it is the last in the system.

Benefits and Drawbacks of WFQ

This topic compares the benefits and drawbacks of WFQ.

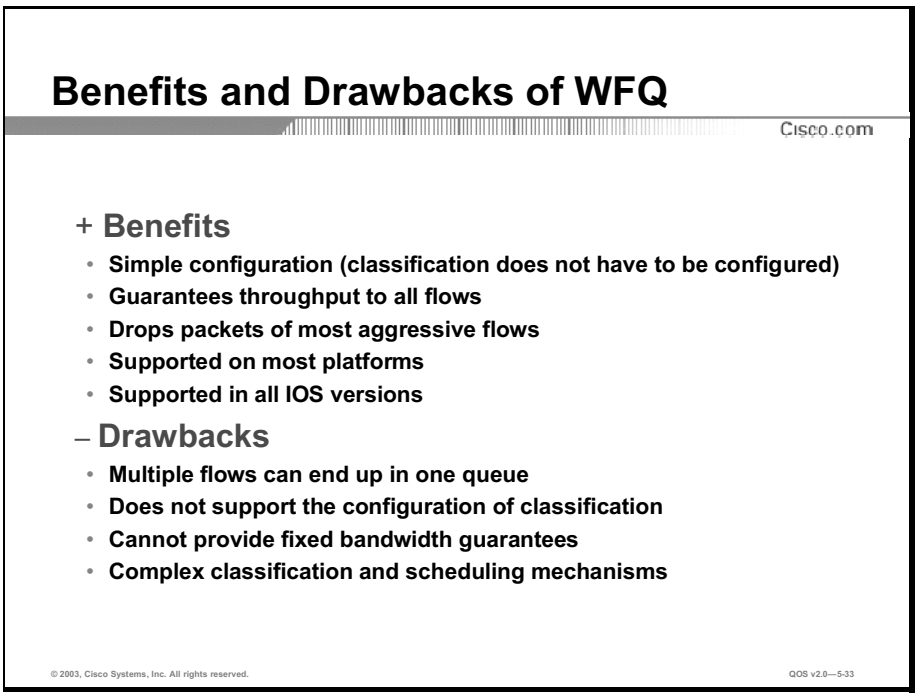

These are the main benefits of WFQ:

- Simple configuration (no manual classification is necessary)
- Drops packets of the most aggressive flows

These are the main drawbacks:

- Not always possible to have one flow per queue
- Does not allow manual classification
- Cannot provide fixed guarantees

Configuring WFQ

This topic explains how to configure WFQ.

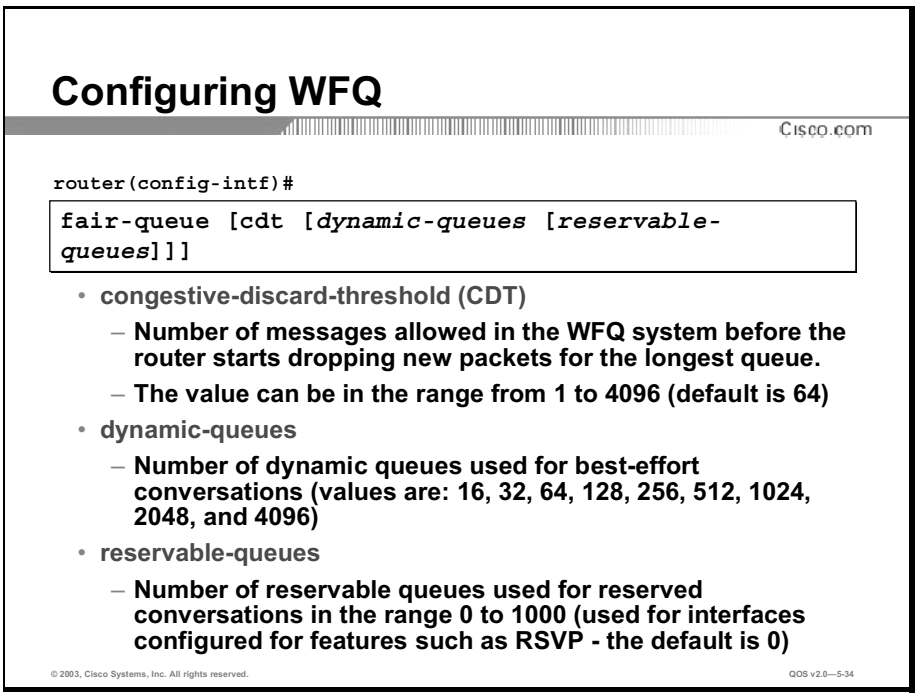

WFQ is automatically enabled on all interfaces that have a default bandwidth of less than 2 Mbps. The **fair-queue** command is used to enable WFQ on interfaces where it is not enabled by default or was previously disabled.

fair-queue [*congestive-discard-threshold* [*dynamic-queues* [*reservable-queues*]]]

Syntax Description

Parameter	Description
congestive-discard-threshold	(Optional) Number of messages allowed in each queue. The default is 64 messages, and a new threshold must be a power of 2 in the range from 16 to 4096. When a conversation reaches this threshold, new message packets are discarded.
dynamic-queues	(Optional) Number of dynamic queues used for best-effort conversations. Values are 16, 32, 64, 128, 256, 512, 1024, 2048, and 4096.
reservable-queues	(Optional) Number of reservable queues used for reserved conversations in the range 0 to 1000. The default is 0. Reservable queues are used for interfaces configured for features such as RSVP.

Additional WFQ Configuration Parameters

Cisco.com

```
router(config-if)#
hold-queue max-limit out
```

- **Specifies the maximum number of packets that can be in all output queues on the interface at any time.**
- **The default value for WFQ is 1000.**
- **Under special circumstances WFQ can consume a lot of buffers which may require lowering this limit.**

The same hold-queue command that can be used with FIFO queuing can also be used with WFQ. The default hold-queue limit with WFQ is 1000 packets.

The WFQ system will generally never reach the hold-queue limit because the CDT limit starts dropping packets of aggressive flows. Under special circumstances it would be possible to fill the WFQ system. For example, a denial-of-service attack that floods the interface with a large number of packets (each different) could fill all queues at the same rate.

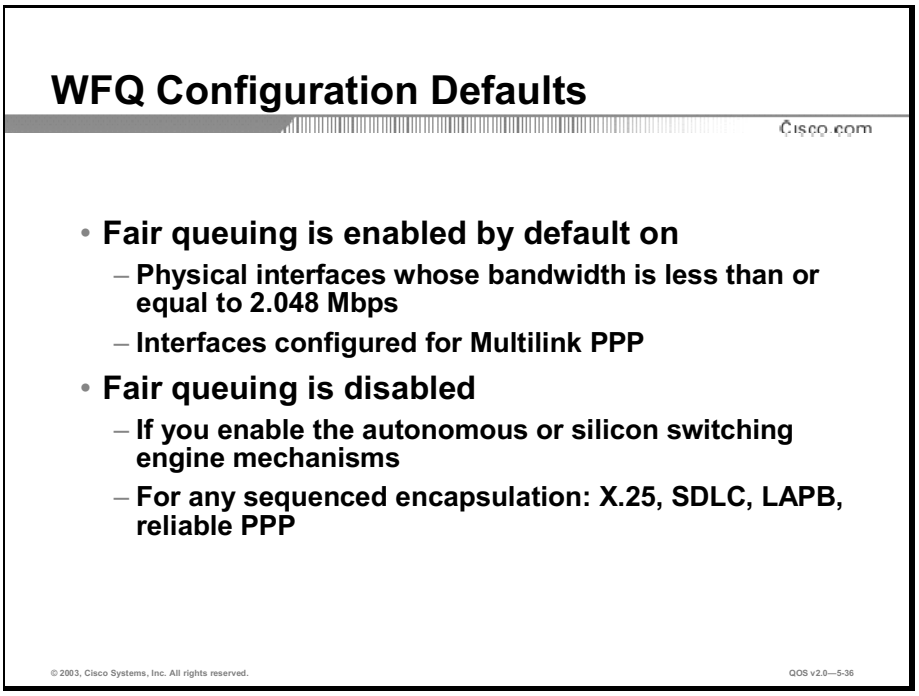

The figure explains the default behavior of WFQ. As mentioned previously, WFQ is automatically enabled on all interfaces slower than 2 Mbps. WFQ is also required on interfaces using Multilink PPP (MLP).

WFQ cannot be used if reordering of frames is not allowed due to sequence numbering of Layer 2 frames or if the switching path does not support WFQ.

Monitoring WFQ

This topic describes the Cisco IOS commands that are used to monitor the operation of WFQ.

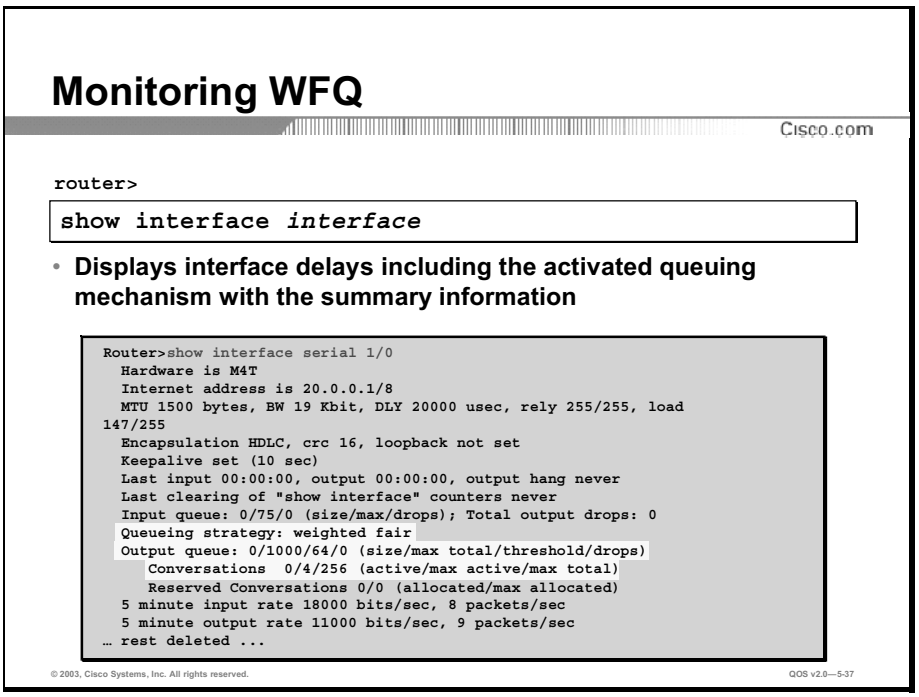

The same **show** commands can be used as with other queuing mechanisms:

- **show interface**
- **show queue**
- **show queueing**

The **show interface** command can be used to determine the queuing strategy. The summary statistics are also displayed.

The sample output in the figure shows that there are currently no packets in the WFQ system that allows up to 1000 packets (hold-queue limit) with CDT 64. WFQ is using 256 queues. The maximum number of concurrent conversations (active queues) was 4.

Monitoring WFQ (Cont.)

Cisco.com

```
router>
show queue interface-name interface-number
```

- **Displays detailed information about the WFQ system of the selected interface**

```
Router>show queue serial 1/0
  Input queue: 0/75/0 (size/max/drops); Total output drops: 0
  Queueing strategy: weighted fair
  Output queue: 2/1000/64/0 (size/max total/threshold/drops)
     Conversations  2/4/256 (active/max active/max total)
     Reserved Conversations 0/0 (allocated/max allocated)

  (depth/weight/discards/tail drops/interleaves) 1/4096/0/0/0
  Conversation 124, linktype: ip, length: 580
  source: 193.77.3.244, destination: 20.0.0.2, id: 0x0166, ttl: 254,
  TOS: 0 prot: 6, source port 23, destination port 11033

  (depth/weight/discards/tail drops/interleaves) 1/4096/0/0/0
  Conversation 127, linktype: ip, length: 585
  source: 193.77.4.111 destination: 40.0.0.2, id: 0x020D, ttl: 252,
  TOS: 0 prot: 6, source port 23, destination port 11013
```


The **show queue** command is used to display the contents of packets inside a queue for a particular interface, including flow (conversation) statistics:

- Queue depth is the number of packets in the queue.
- Weight is 4096/(IP precedence + 1) or 32384/(IP precedence + 1), depending on the Cisco IOS version.
- In the command output, discards are used to represent the number of drops due to the CDT limit.
- In the command output, tail drops are used to represent the number of drops due to the hold-queue limit.

Summary

This topic summarizes the key points discussed in this lesson.

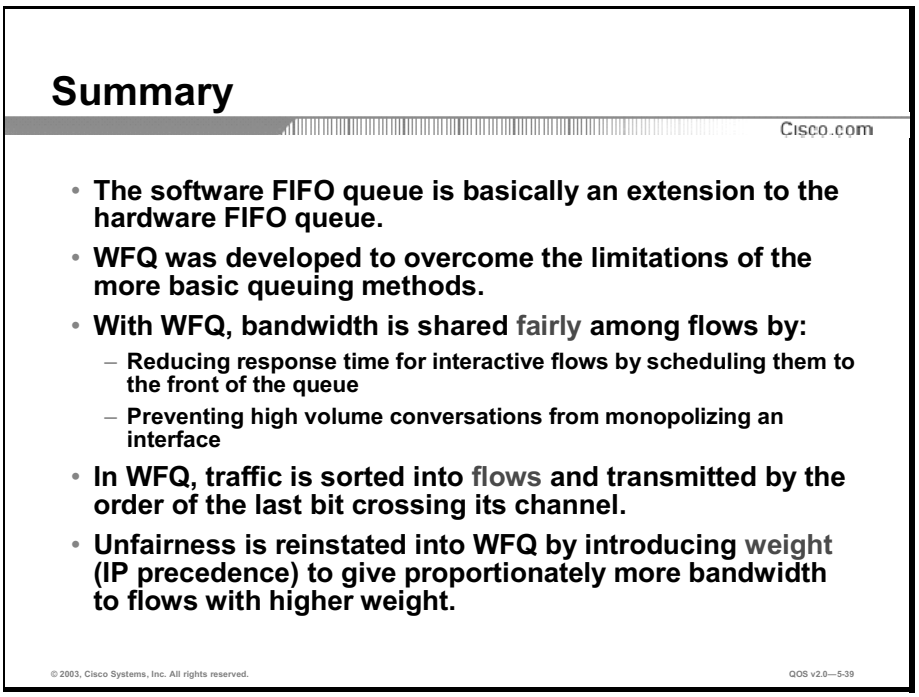

References

For additional information, refer to these resources:

- To learn more about configuring WFQ, refer to “Configuring Weighted Fair Queueing” at the following URL: http://www.cisco.com/en/US/products/sw/iosswrel/ps1828/products_configuration_guide_chapter09186a00800ca597.html

Next Steps

For the associated lab exercise, refer to the following section of the course Lab Guide:

- Lab Exercise 5-1: Configuring Basic Queuing

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) Which three of the following represent the strategy behind WFQ? (Choose three.)

A) dedicated queue for each flow

B) immediately dispatch all voice packets

C) fairly allocate bandwidth among all flows

D) use IP precedence as weight when allocating bandwidth

Q2) Which three of the following would be used to identify a WFQ flow? (Choose three.)

A) destination IP address

B) HTTP application identifier

C) source TCP/UDP port number

D) protocol number (identifying TCP or UDP)

Q3) How many queues are recommended with WFQ?

A) same number of queues as there are flows

B) several times as many queues as there are flows

C) depends upon the QoS requirements of the applications

D) one queue for each predicted three or four flows (on average)

Q4) Which two represent the two modes of dropping used in WFQ? (Choose two.)

A) early dropping when HQO is reached

B) early dropping when CDT is reached

C) aggressive dropping when CDT is reached

D) aggressive dropping when HQO is reached

Q5) Given that the following five packets arrive at a router using WFQ, in what order would the packets be dispatched? (Each packet type—A or B—represents a different flow and, therefore, will go to a different queue.)

Packets Arriving at Router

Packet	Arrival Time	Time to Transmit Packet
A1	T + 0	500 ms
A2	T + 50	100 ms
B1	T + 150	250 ms
B2	T + 150	400 ms
B3	T + 200	400 ms

A) A1, B1, A2, B2, B3

B) A1, A2, B1, B2, B3

C) A1, B1, B2, A2, B3

D) B1, A1, A2, B2, B3

Q6) Which three of the following represent benefits of WFQ? (Choose three.)

A) simple configuration

B) guaranteed throughput to all flows

C) provides fixed bandwidth guarantees

D) drops packets of the most aggressive flows

Quiz Answer Key

Q1) A, C, D

Relates to: Weighted Fair Queuing

Q2) A, C, D

Relates to: WFQ Classification

Q3) B

Relates to: WFQ Classification

Q4) B, D

Relates to: WFQ Insertion and Drop Policy

Q5) A

Relates to: WFQ Scheduling

Q6) A, B, D

Relates to: Benefits and Drawbacks of WFQ

CBWFQ and LLQ

Overview

CBWFQ extends the standard WFQ functionality to provide support for user-defined traffic classes. With CBWFQ, you define traffic classes based on match criteria including protocols, access control lists (ACLs), and input interfaces. 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.

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.

Relevance

These advanced queuing models offer the best QoS congestion management solutions for networks with converged traffic.

Objectives

Upon completing this lesson, you will be able to configure CBWFQ and LLQ to manage congestion. This includes being able to meet these objectives:

- Explain how basic queuing mechanisms can be used to build advanced queuing mechanisms
- Explain the purpose and features of CBWFQ
- Describe CBWFQ features and explain how CBWFQ works using a block diagram
- Describe the benefits of CBWFQ
- Identify the Cisco IOS commands required to configure and monitor CBWFQ on a Cisco router
- Explain the purpose and features of LLQ
- Explain how LLQ works using a block diagram and identify situations in which LLQ is most appropriate for providing QoS
- Describe the benefits of LLQ
- Identify the Cisco IOS commands required to configure and monitor LLQ on a Cisco router

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Basic knowledge of internetworking with TCP/IP concepts and basic knowledge of Cisco IOS commands

Outline

The outline lists the topics included in this lesson.

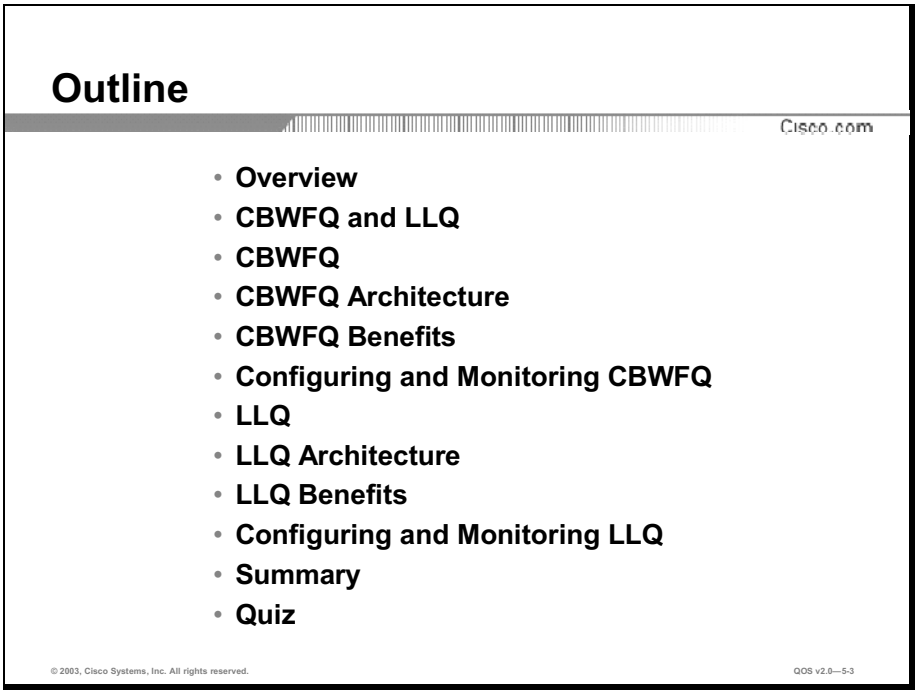

CBWFQ and LLQ

This topic explains how basic queuing mechanisms can be used to build more advanced queuing mechanisms.

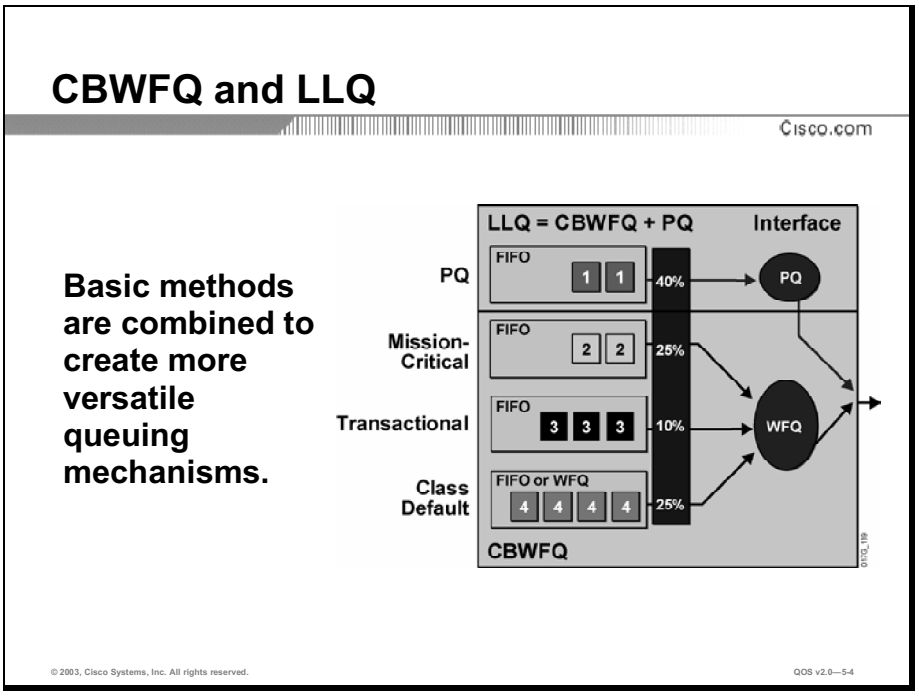

Neither the basic queuing methods nor the more advanced WFQ completely solved the QoS problems resulting from converged network traffic. Some problems remaining were:

- If only priority queuing (PQ) was used for a voice-enabled network, voice would get the priority needed. However, data traffic would suffer.
- If only CQ was used for voice-enabled network, data traffic would be assured of some bandwidth. However, voice traffic would suffer delays.
- If WFQ was used, voice still experienced delay even when treated “fairly” by WFQ.
- All of the classification, marking, and queuing mechanisms were complicated to use and time-consuming when applied on an interface-by-interface basis.

Newer queuing mechanisms were developed, which took the best aspects of existing queuing methods and applied them to give voice the priority it required while still ensuring that data was serviced efficiently on a class basis.

CBWFQ

This topic explains the purpose of CBWFQ.

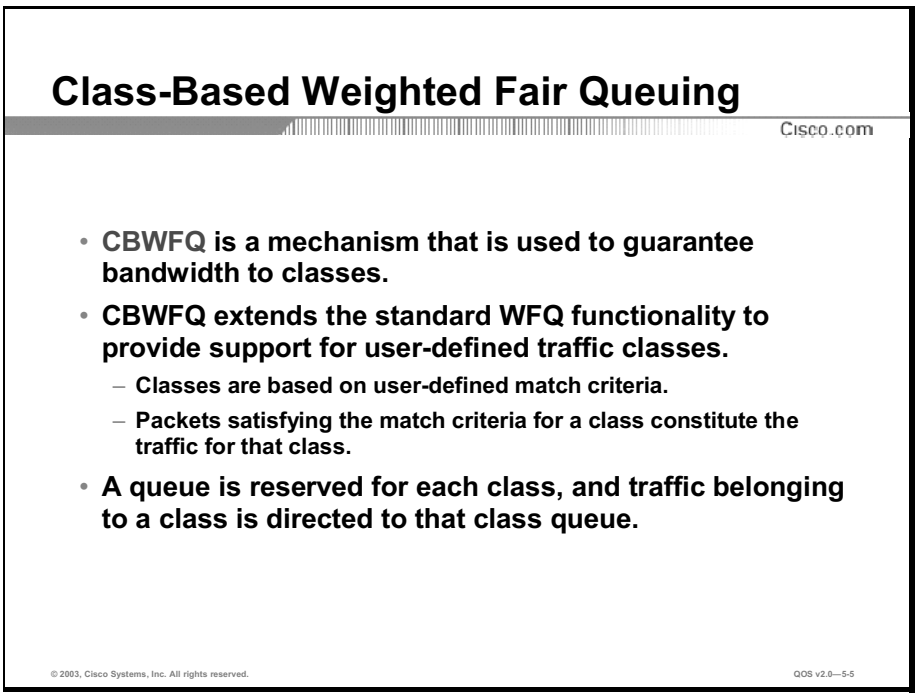

CBWFQ extends the standard WFQ functionality to provide support for user-defined traffic classes. With CBWFQ, the user defines the traffic classes based on match criteria that includes protocols, ACLs, and input interfaces. Packets satisfying the match criteria for a class constitute the traffic for that class. A queue is reserved for each class, and traffic belonging to a class is directed to that class queue.

After a class has been defined according to its match criteria, you can assign it characteristics. To characterize a class, you assign it bandwidth, weight, and maximum packet limit. The bandwidth assigned to a class is the minimum bandwidth delivered to the class during congestion.

To characterize a class, you also specify the queue limit for that class, which is the maximum number of packets allowed to accumulate in the class queue. Packets belonging to a class are subject to the bandwidth and queue limits that characterize the class. After a queue has reached its configured queue limit, enqueuing of additional packets to the class causes tail drop or random packet drop to take effect, depending on how the class policy is configured.

CBWFQ Architecture

This topic explains the features of CBWFQ and how CBWFQ works.

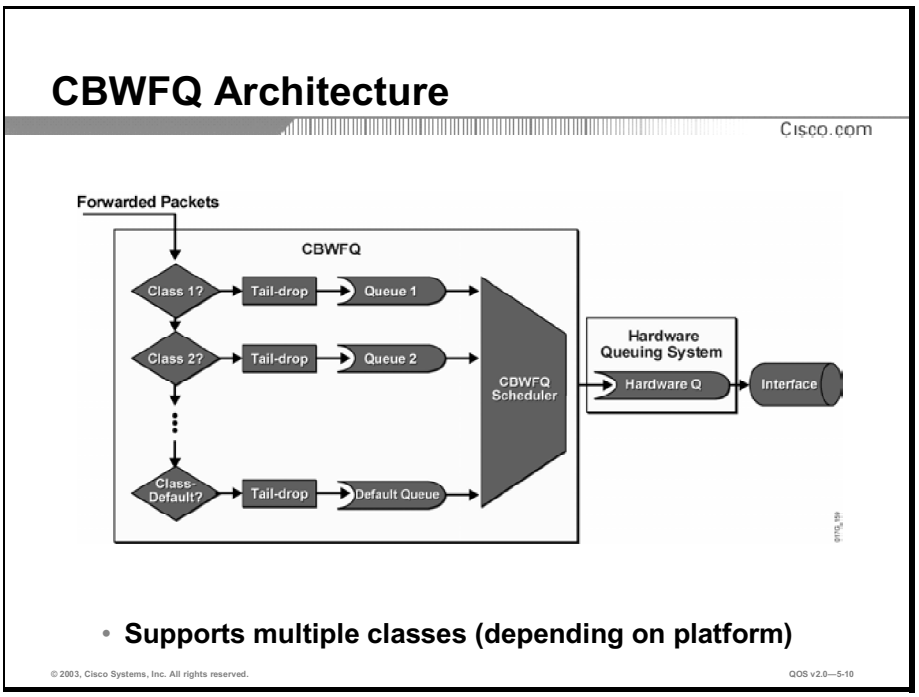

CBWFQ supports multiple class maps (number depends upon platform) to classify traffic into their corresponding FIFO queues. Tail drop is the default dropping scheme of CBWFQ although it can be combined with WRED.

The CBWFQ scheduler is used to guarantee bandwidth that is based on the configured weights.

Note:	Currently, except for the Cisco 7500 series router platform, all traffic classes (default traffic class excluded) only support FIFO queuing within the class. On all platforms, the default traffic class can support either FIFO or WFQ within the class. Check Cisco.com for the latest information for WFQ support within each traffic class.

CBWFQ Architecture: Classification

Cisco.com

- **Classification uses class maps.**
- **Availability of certain classification options depends on the Cisco IOS version.**
- **Some classification options depend on type of interface and encapsulation where service policy is used.**
- **For example:**
  - **Matching on Frame Relay discard eligible (DE) bits can only be used on interfaces with Frame Relay encapsulation.**
  - **Matching on MPLS experimental bits has no effect if MPLS is not enabled.**
  - **Matching on ISL Priority bits has no effect if ISL is not used.**

Any classification option can be used depending on the availability in the Cisco IOS version and the support on the selected interface and encapsulation.

It is important to note that CBWFQ is configured using Modular QoS command-line interface (CLI [MQC]).

CBWFQ Architecture: Insertion Policy

- **Each queue has a maximum number of packets that it can hold (queue size).**
- **The maximum queue size is platform-dependent.**
- **After a packet is classified to one of the queues, the router will enqueue the packet if the queue limit has not been reached (tail-drop within each class).**
- **WRED can be used in combination with CBWFQ to prevent congestion of the class.**

CBWFQ reserves multiple FIFO queues in the WFQ system. The default queue limit is 64 (tail-drop) and can be configured with WRED (random drop).

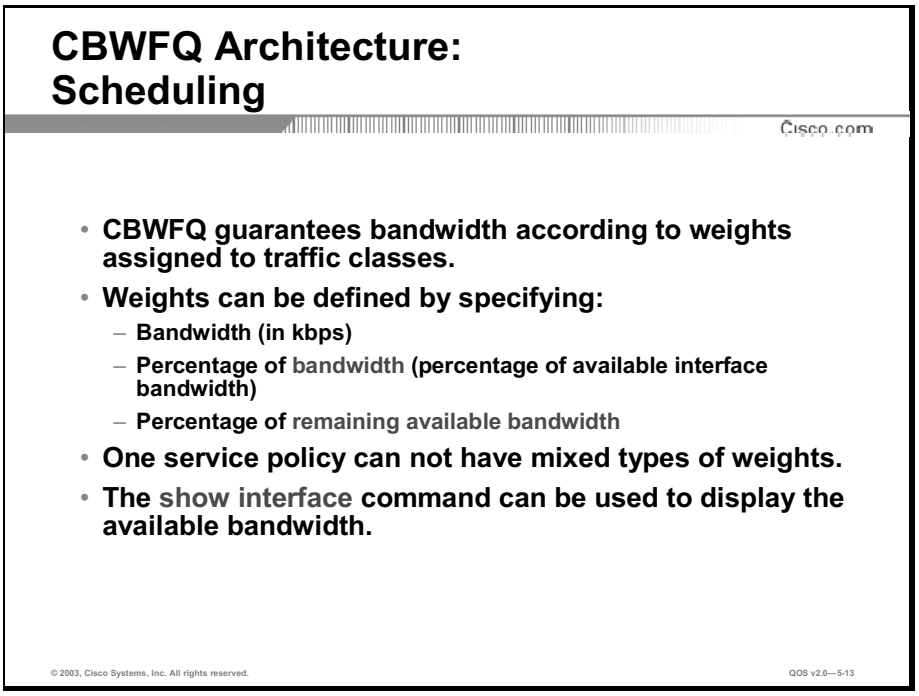

You can configure bandwidth guarantees by using one of the following commands:

- The **bandwidth** command allocates a fixed amount of bandwidth by specifying the amount in kbps. The reserved bandwidth is subtracted from the available bandwidth of the interface where the service policy is used. The allocated bandwidth must also be within the configured reservable limit (75 percent by default).
- The **bandwidth percent** command can be used to allocate a percentage of the default or available bandwidth of an interface. The default bandwidth usually equals the maximum speed of an interface. Sometimes it actually reflects the real speed of an interface (for example, Ethernet or FastEthernet). The default value can be replaced by using the **bandwidth** interface command. It is recommended that the bandwidth reflect the real speed of the link. The allocated bandwidth is subtracted from the available bandwidth of the interface where the service policy is used.
- The **bandwidth remaining percent** command can be used to allocate a portion of the unallocated bandwidth. The bandwidth is not subtracted from the available bandwidth of the interface where the service policy is used.

A single service policy cannot mix the fixed bandwidth (in kbps) and bandwidth percent commands (except with strict priority queues).

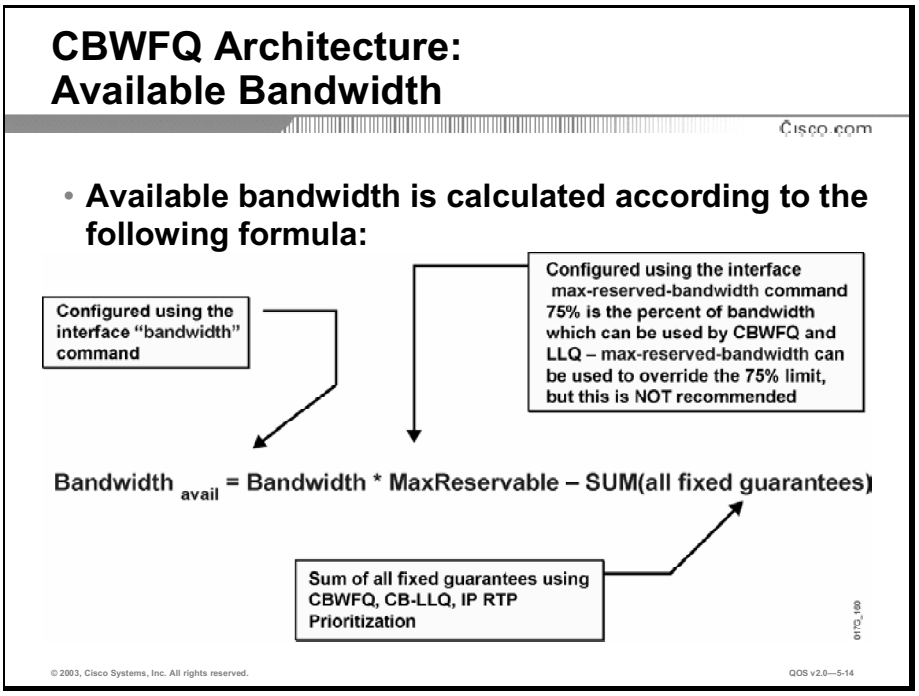

The available bandwidth displayed by the **show interface** command is calculated by subtracting all fixed bandwidth reservations from 75 percent of the configured bandwidth of an interface.

CBWFQ Architecture: 75 Percent Rule

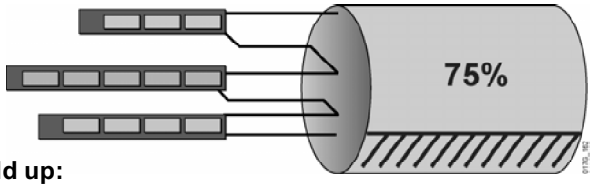

- **Add up:**
  - **Class bandwidths**
  - **RSVP maximum reserved bandwidth**
- **Result must be less than or equal to 75% of interface bandwidth (or Frame Relay, data-link connection identifier [DLCI], committed information rate [CIR])**
  - **Leaves headroom for call signaling, Simple Network Management Protocol (SNMP), Local Management Interface (LMI), and routing traffic**
- **The 75% rule is a conservative rule**
- ***Max-reserved-bandwidth* command overrides 75% limit, but seldom recommended**

Properly provisioning the network bandwidth is a major component of successful network design. You can calculate the required bandwidth by adding the bandwidth requirements for each major application (for example, voice, video, and data). The resulting sum represents the minimum bandwidth requirement for any given link, and it should not exceed approximately 75 percent of the total available bandwidth for the link. This 75 percent rule assumes that some bandwidth is required for overhead traffic, such as routing and Layer 2 keepalive messages, as well as for additional applications such as e-mail and HTTP traffic.

Thus, the total amount of bandwidth allocated for all classes included in a policy map should not exceed 75 percent of the available bandwidth on the interface. The max-reserved bandwidth command overrides the 75 percent limitation, but overriding is recommended only for the most knowledgeable network administrators who have access to precise figures for available, used and required bandwidth. If not all of the bandwidth is allocated, the remaining bandwidth is proportionally allocated among the classes, based on their configured bandwidth.

Note that the 75 percent rule is conservative.

CBWFQ Benefits

This topic describes the benefits of CBWFQ.

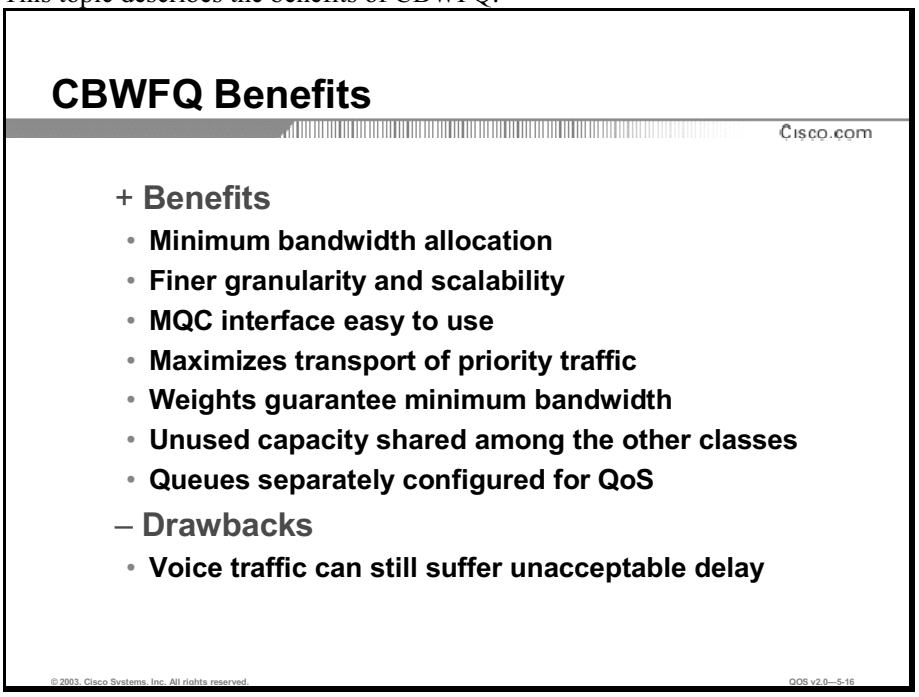

CBWFQ allows the user to define traffic classes based on custom-defined match criteria such as ACLs, input interfaces, protocol, and QoS label. For example, a class might consist of a team working on a certain project or a class can be created for the important mission-critical applications such as enterprise resource planning (ERP). When the traffic classes have been defined, they can be assigned a bandwidth, queue limit, or drop policy such as WRED. Additional benefits of CBWFQ are the following:

- **Bandwidth allocation:** CBWFQ allows you to specify the exact amount of bandwidth to be allocated for a specific class of traffic. Accounting for available bandwidth on the interface, you can configure multiple classes (number depends upon platform).
- **Finer granularity and scalability:** CBWFQ allows you total flexibility to define a class, based on ACLs and protocols or input interfaces, thereby providing finer granularity.
- **Supported by MQC:** CBWFQ is supported by the easy-to-use MQC.

The CBWFQ feature is supported on all platforms that WFQ is supported on; in other words, the Cisco 7200, 4700, 4500, 3600, and 2600 series, and so on.

Configuring and Monitoring CBWFQ

This topic describes the Cisco IOS commands that are used to configure and monitor CBWFQ.

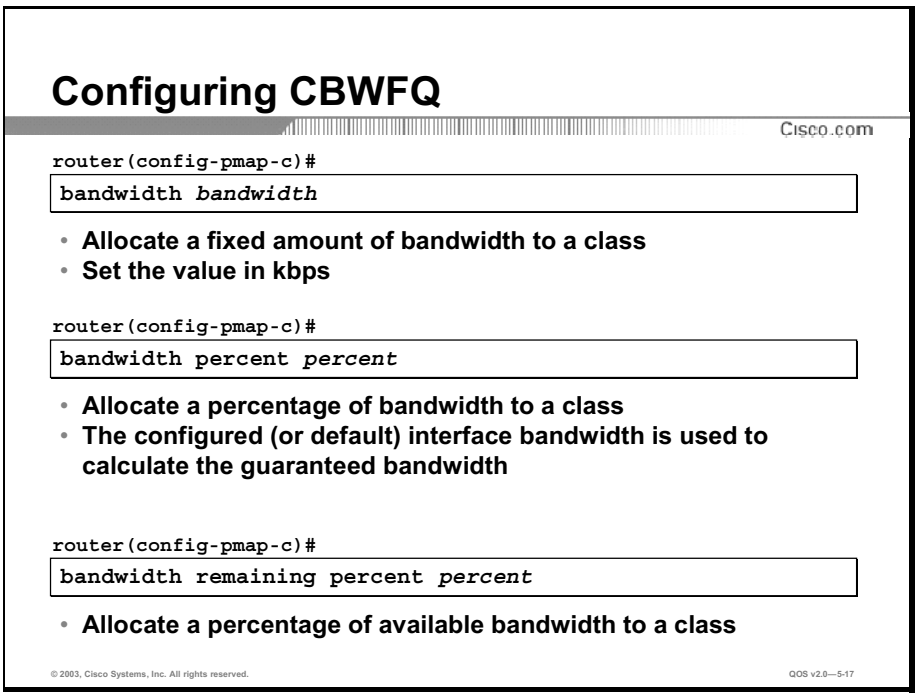

The **bandwidth** policy-map class configuration command is used to specify or modify the bandwidth allocated for a class belonging to a policy map.

All classes belonging to one policy map should use the same type of bandwidth guarantee (fixed in kbps, in percentage of interface bandwidth, in percentage of available bandwidth).

Configuring bandwidth in percentages is most useful when the underlying link bandwidth is unknown or the relative class bandwidth distributions are known.

bandwidth {*bandwidth-kbps* | **remaining percent** *percentage* | **percent** *percentage*}

Syntax Description

Parameter	Description
bandwidth-kbps	Amount of bandwidth, in kbps, to be assigned to the class.
remaining percent *percentage*	Amount of guaranteed bandwidth, based on a relative percent of available bandwidth. The percentage can be a number from 1 to 100.
percent *percentage*	Amount of guaranteed bandwidth, based on an absolute percent of available bandwidth. The percentage can be a number from 1 to 100.

The following restrictions apply to the bandwidth command:

- If the percent keyword is used, the sum of the class bandwidth percentages cannot exceed 100 percent.
- The amount of bandwidth configured should be large enough to also accommodate Layer 2 overhead.
- A policy map can have all the class bandwidths specified in kbps or all the class bandwidths specified in percentages, but not a mix of both. However, the unit for the priority command in the priority class can be different from the bandwidth unit of the low priority class.

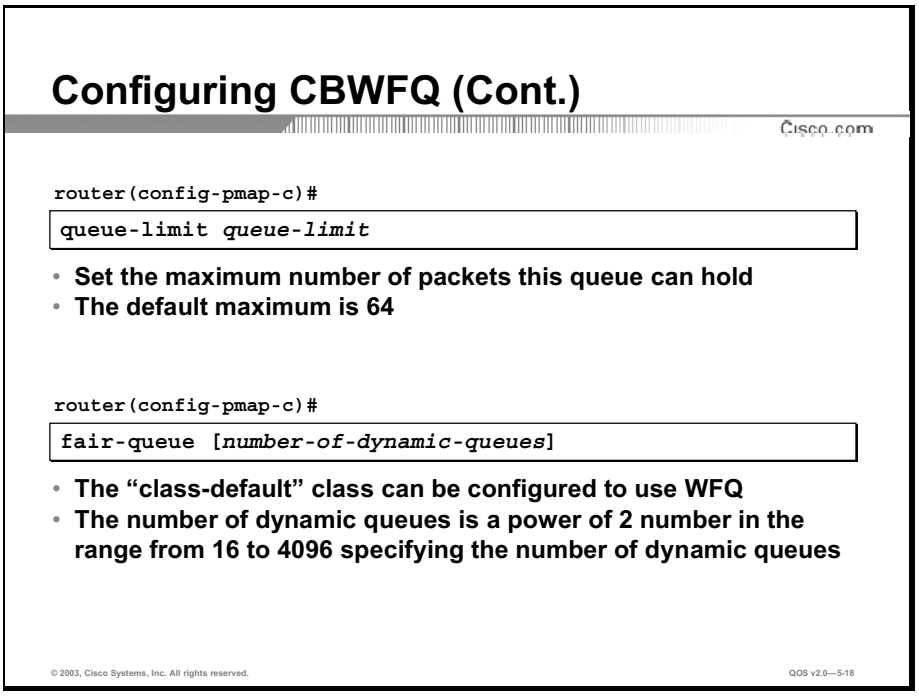

The default queue limit of 64 packets can be changed using the **queue-limit** command. It is recommended not to change the default value.

The default class can be selected by specifying the **class-default** name of the class. The default class supports two types of queuing: one FIFO queue (default) or a flow-based WFQ system. Both types can be combined with WRED. FIFO queue can also get a bandwidth guarantee.

Example: Configuration of FIFO Queuing

The following example shows the configuration of FIFO queuing within the default class. The default class is also guaranteed 1 Mbps of bandwidth and the maximum queue size is limited to 40 packets.

```
policy-map A
 class A
  bandwidth 1000
 class class-default
  bandwidth 1000
  queue-limit 40
```

Example: Configuration of WFQ Queuing

This next example shows the configuration of WFQ queuing within the default class. The number of dynamic queues is set to 1024 and the discard threshold is set to 50.

```
policy-map A
 class A
  bandwidth 1000
 class class-default
  fair-queue 1024
```

Configuring CBWFQ (Cont.)


```
Router(config)# access-list 101 permit udp host 10.10.10.10 host
10.10.10.20 range 16384 20000
Router(config-if)# access-list 101 permit udp host  10.10.10.10
host 10.10.10.20 range 53000 56000
Router(config)# class-map class1
Router(config-cmap)# match access-group 101
Router(config-cmap)# exit
Router(config-cmap)# class-map class2
Router(config-cmap)# match access-group 102
Router(config-cmap)# exit

Router(config)# policy-map policy1
Router(config-pmap)# class class1
Router(config-pmap-c)# bandwidth 3000
Router(config-pmap-c)# queue-limit 30
Router(config-pmap-c)# exit
Router(config-pmap)# class class2
Router(config-pmap-c)# bandwidth 2000
Router(config-pmap-c)# exit
```


The sample configuration shows how CBWFQ is used to guarantee bandwidth to each of the two classes.

Monitoring CBWFQ

Cisco.com

```
router>
show policy-map interface [interface]
```

- **Displays parameters and statistics of CBWFQ**

```
router>show policy-map interface
 FastEthernet0/0

  Service-policy output: Policy1

    Class-map: Class1 (match-any)
      0 packets, 0 bytes
      5 minute offered rate 0 bps, drop rate 0 bps
      Match: any
      Weighted Fair Queueing
        Output Queue: Conversation 265
        Bandwidth remaining 20 (%) Max Threshold 64 (packets)
        (pkts matched/bytes matched) 0/0
        (depth/total drops/no-buffer drops) 0/0/0
    Class-map: class-default (match-any)
      42 packets, 4439 bytes
      5 minute offered rate 0 bps, drop rate 0 bps
      Match: any
```


The **show policy-map interface** command displays all service policies applied to the interface. Among the settings, policing parameters and statistics are displayed.

LLQ

This topic describes the purpose and features of LLQ.

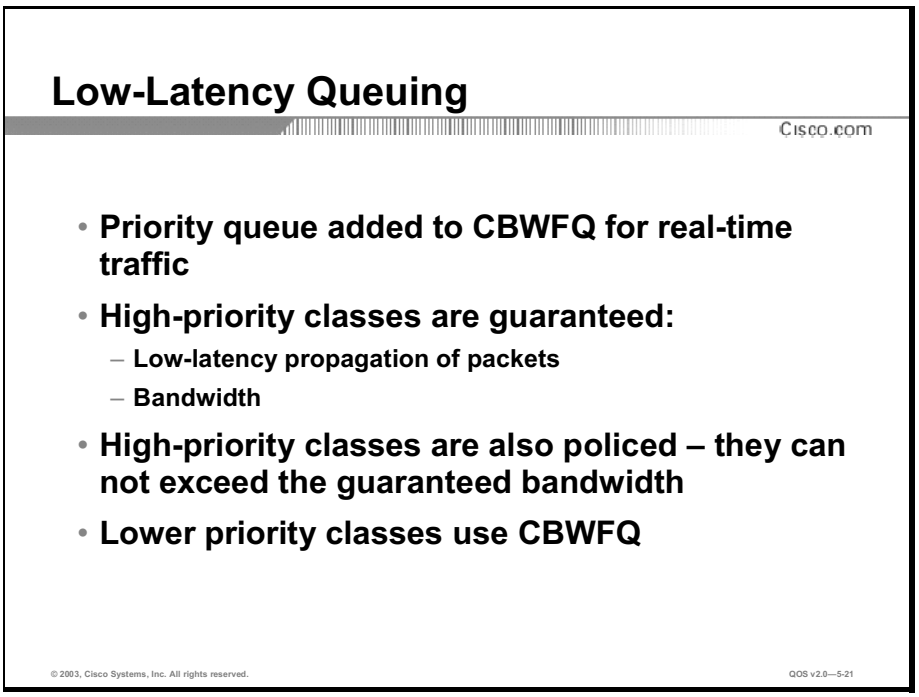

While WFQ provides a fair share of bandwidth to every flow, and provides fair scheduling of its queues, it cannot provide guaranteed bandwidth and low delay to select applications. For example, voice traffic may still compete with other aggressive flows in the WFQ queuing system because the WFQ system lacks priority scheduling for time-critical traffic classes.

The LLQ feature 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.

Without LLQ, CBWFQ provides weighted fair queuing based on defined classes with no strict priority queue available for real-time traffic. CBWFQ allows you to define traffic classes and then assign characteristics to that class. For example, you can designate the minimum bandwidth delivered to the class during congestion.

For CBWFQ, the weight for a packet belonging to a specific class is derived from the bandwidth that you assigned to the class when you configured it. Therefore, the bandwidth assigned to the packets of a class determines the order in which packets are sent. All packets are serviced fairly based on weight; no class of packets may be granted strict priority. This scheme poses problems for voice traffic, which is largely intolerant of delay, especially variation in delay. For voice traffic, variations in delay introduce irregularities of transmission manifesting as jitter in the heard conversation.

The LLQ feature provides strict priority queuing for CBWFQ, reducing jitter in voice conversations. Configured by the **priority** command, LLQ enables use of a single, strict priority queue within CBWFQ at the class level, allowing you to direct traffic belonging to a class to the CBWFQ strict priority queue. To enqueue class traffic to the strict priority queue, you configure the **priority** command for the class after you specify the named class within a

policy map. (Classes to which the priority command is applied are considered priority classes.) Within a policy map, you can give one or more classes priority status. When multiple classes within a single policy map are configured as priority classes, all traffic from these classes is enqueued to the same, single, strict priority queue.

LLQ Architecture

This topic explains how LLQ works and identifies situations in which LLQ is most appropriate for providing QoS.

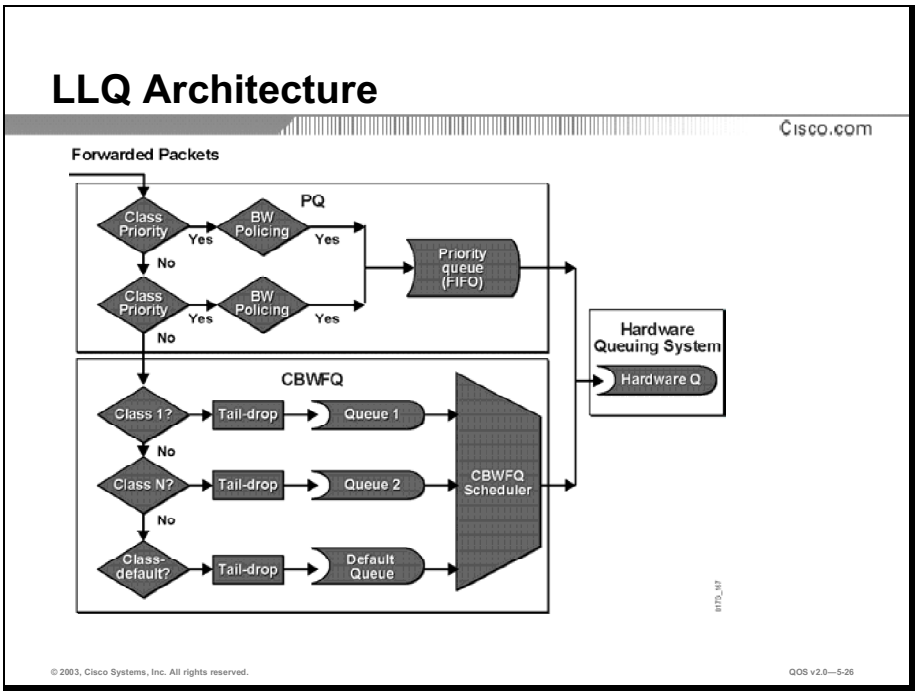

When CBWFQ is configured as the queuing system, it creates a number of queues, into which it classifies traffic classes. These queues are then scheduled with a WFQ-like scheduler, which can guarantee bandwidth to each class.

If LLQ is used within the CBWFQ system, it creates an additional priority queue in the WFQ system, which is serviced by a strict priority scheduler. Any class of traffic can therefore be attached to a service policy, which uses priority scheduling, and hence can be prioritized over other classes.

LLQ Benefits

This topic describes the benefits of LLQ.

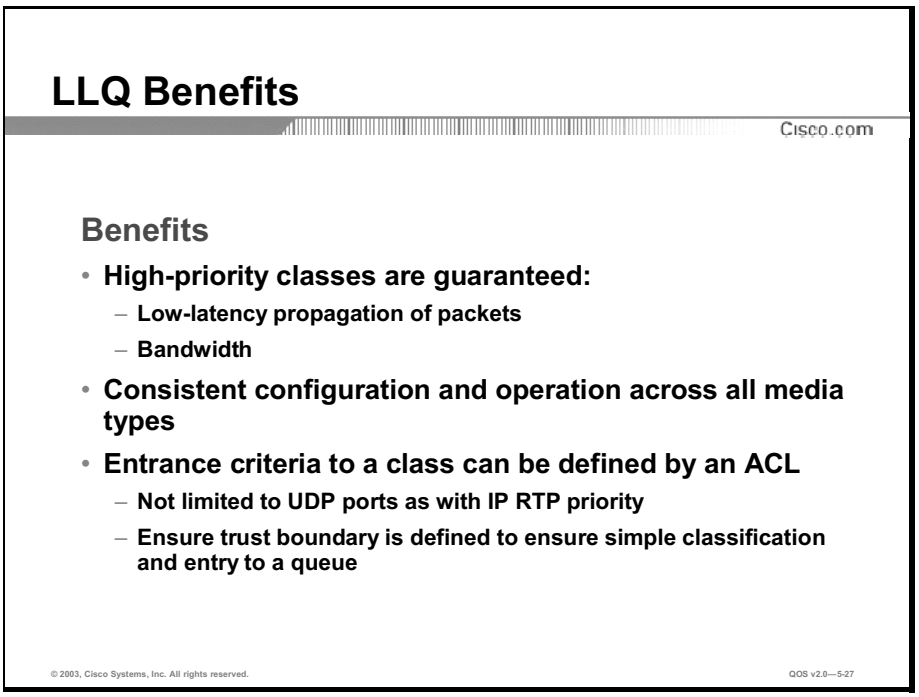

The LLQ priority scheduler guarantees both low-latency propagation of packets and bandwidth to high-priority classes. Low-latency is achieved by expediting traffic using a priority scheduler. Bandwidth is also guaranteed by the nature of priority scheduling, but is policed to a user-configurable value. The strict PQ scheme allows delay-sensitive data such as voice to be dequeued and sent first—that is, before packets in other queues are dequeued. Delay-sensitive data is given preferential treatment over other traffic.

This feature provides strict PQ on ATM virtual circuits (VCs); the IP Real-Time Transport Protocol (RTP) priority feature only allows PQ on interfaces. Because you can configure the priority status for a class within CBWFQ, you are not limited to UDP port numbers to stipulate priority flows (which were necessary with IP RTP). Instead, all of the valid match criteria used to specify traffic for a class now applies to priority traffic.

Policing of priority queues also prevents the priority scheduler from monopolizing the CBWFQ scheduler and starving non-priority classes, like legacy PQ does. By configuring the maximum amount of bandwidth allocated for packets belonging to a class, you can avoid starving non-priority traffic.

Configuring and Monitoring LLQ

This topic describes the Cisco IOS commands that are used to configure and monitor LLQ.

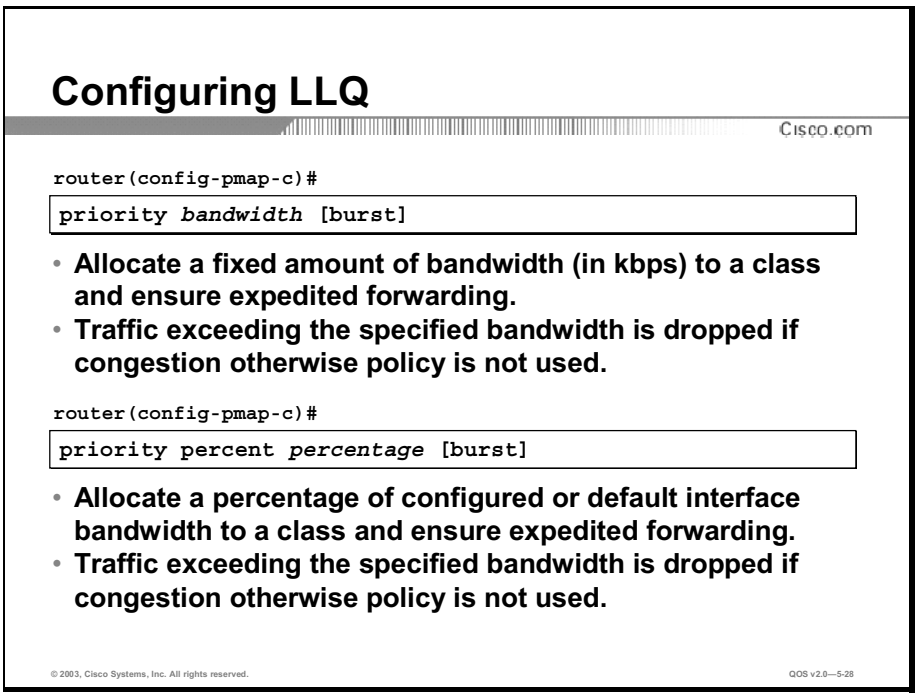

When you specify the **priority** command for a class, it takes a bandwidth argument that gives maximum bandwidth in kbps. You use this parameter to specify the maximum amount of bandwidth allocated for packets belonging to the class configured with the **priority** command. The bandwidth parameter both guarantees bandwidth to the priority class and restrains the flow of packets from the priority class.

priority {***bandwidth-kbps*** | **percent** ***percentage***} [**burst**]

Parameter	Description
bandwidth-kbps	Guaranteed allowed bandwidth, in kbps, for the priority traffic. The amount of guaranteed bandwidth varies according to the interface and platform in use. Beyond the guaranteed bandwidth, the priority traffic will be dropped in the event of congestion to ensure that the non-priority traffic is not starved.
percent	Specifies that the amount of guaranteed bandwidth will be specified by the percent of available bandwidth.
percentage	Used in conjunction with the percent keyword, specifies the percentage of the total available bandwidth to be set aside for the priority class. The percentage can be a number from 1 to 100.
burst	(Optional) Specifies the burst size, in bytes. The range of the burst is 32 to 2,000,000 bytes.

In the event of congestion, when the bandwidth is exceeded policing is used to drop packets. Voice traffic enqueued to the priority queue is UDP-based and therefore not adaptive to the early packet drop characteristic of WRED. Because WRED is ineffective, you cannot use the WRED **random-detect** command with the **priority** command. In addition, because policing is used to drop packets and queue limit is not imposed, the **queue-limit** command cannot be used with the **priority** command.

When congestion occurs, traffic destined for the priority queue is metered to ensure that the bandwidth allocation configured for the class to which the traffic belongs is not exceeded.

Priority traffic metering has the following qualities:

- It is much like Committed Access Rate (CAR) rate limiting, except that priority traffic metering is only performed under congestion conditions. When the device is not congested, the priority class traffic is allowed to exceed its allocated bandwidth. When the device is congested, the priority class traffic above the allocated bandwidth is discarded.
- Performed metering on a per-packet basis, and tokens are replenished as packets are sent. If not enough tokens are available to send the packet, the packet is dropped.
- Restrains priority traffic to its allocated bandwidth to ensure that non-priority traffic, such as routing packets and other data, is not starved.

With metering, the classes are policed and rate-limited individually. That is, although a single policy map might contain four priority classes, all of which are enqueued in a single priority queue, they are each treated as separate flows with separate bandwidth allocations and constraints.

Note: It is important to note that because bandwidth for the priority class is specified as a parameter to the **priority** command, you cannot also configure the **bandwidth** command for a priority class. To do so is a configuration violation that would only introduce confusion in relation to the amount of bandwidth to allocate.

Keep the following guidelines in mind when using the **priority** command:

- Layer 2 encapsulations are accounted for in the amount of bandwidth specified with the **priority** command. However, ensure that a bandwidth is configured with room for the cell-tax overhead.
- Use the **priority** command for Voice over IP (VoIP) on serial links and ATM PVCs.

Note: An exception to these guidelines for LLQ is Frame Relay on the Cisco 7200 router and other non-Route Switch Processor (RSP) platforms. The original implementation of LLQ over Frame Relay on these platforms did not allow the priority classes to exceed the configured rate during periods of non-congestion. Cisco IOS Software release 12.2 removes this exception and ensures that non-conforming packets are only dropped if there is congestion. In addition, packets smaller than an FRF.12 fragmentation size are no longer sent through the fragmenting process, resulting in reduced CPU utilization.

- Use the **priority** command in conjunction with the **set** command. You cannot use the **priority** command in conjunction with any other command, including the **random-detect**, **queue-limit**, and **bandwidth** commands.
- You can configure the **priority** command in multiple classes, but you should only use it for voice-like, constant bit rate (CBR) traffic. If the traffic is not CBR, you must configure a large enough bandwidth parameter to absorb the data bursts.

Warning Although it is possible to enqueue various types of real-time traffic to the strict priority queue, it is strongly recommended that you direct only voice traffic to it. This recommendation is made because voice traffic is well behaved, whereas other types of real-time traffic are not. Moreover, voice traffic requires that delay is non-variable in order to avoid jitter. Real-time traffic such as video could introduce variation in delay, thereby thwarting the steadiness of delay required for successful voice traffic transmission.

Example: Calculating LLQ BW Required for VoIP

The bandwidth consumed by VoIP streams is calculated by adding the packet payload and all the headers, then multiplying that total number by the per-second packet rate.

The following example shows how to calculate the VoIP bearer bandwidth requirement for a single VoIP call using a G.711 codec:

```
G.711              = 160 bytes payload size

Packet size        = payload size + IP/UDP/RTP headers
             = 160 bytes + 20 bytes + 8 bytes + 12 bytes
             = 200 bytes

Sampling Rate = 20 msec per sample = 50 samples per second

Bandwidth (bytes/sec) without Layer 2 overhead
= 200 bytes/packet x 50 packets/second
= 10000 bytes/second

Bandwidth (bits/sec) without Layer 2 overhead
= 10000 bytes/second  * 8 bits/byte
= 80000 bytes/second (80 kbps)

Bandwidth (bits/sec) with Layer 2 overhead

                   = 80000 bytes/second + L2 overhead
bytes/second
```

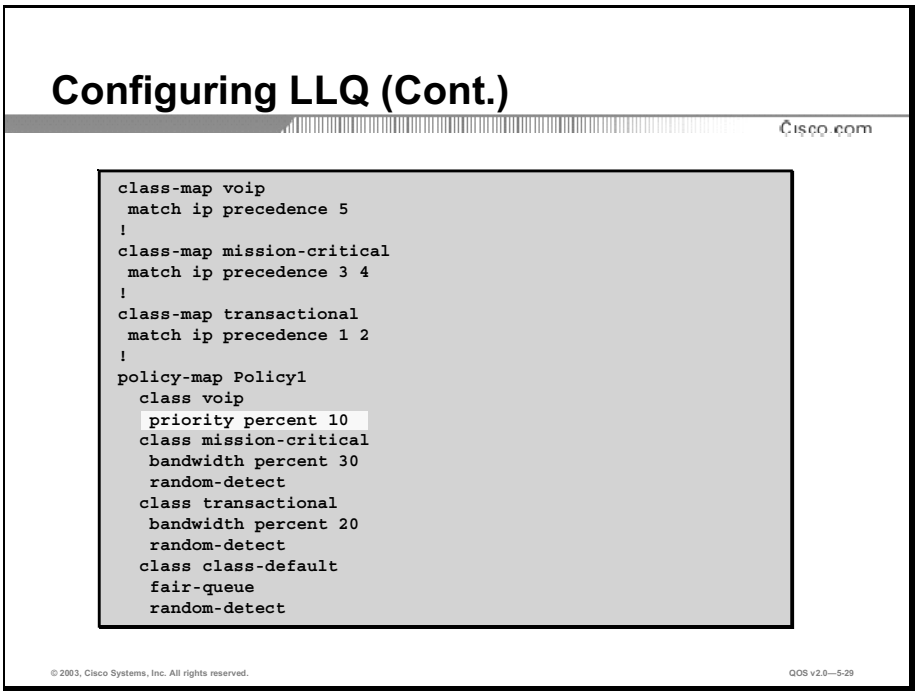

This figure shows a configuration example where the VoIP traffic class, classified by the IP precedence of 5, is queued in a priority queue within the CBWFQ system. The priority class received priority scheduling compared to other classes queues: it is guaranteed but limited to 10 percent of bandwidth.

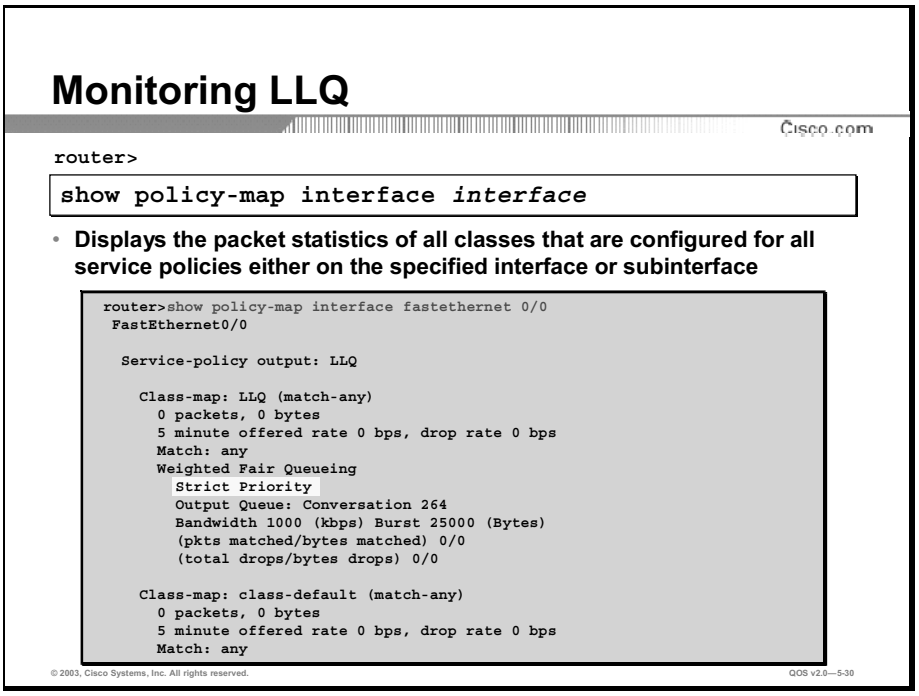

The **show policy-map interface** command displays the packet statistics of all classes that are configured for all service policies on the specified interface. Some of the key fields in the command output are described as follows:

Parameter	Description
Class-map	Class of traffic being displayed. Output is displayed for each configured class in the policy.
offered rate	Rate, in kbps, of packets coming in to the class.
drop rate	Rate, in kbps, at which packets are dropped from the class. The drop rate is calculated by subtracting the number of successfully transmitted packets from the offered rate.
Match	Match criteria specified for the class of traffic.
pkts matched/bytes matched	Number of packets (also shown in bytes) matching this class that were placed in the queue.
depth/total drops/no-buffer drops	Number of packets discarded for this class. No-buffer indicates that no memory buffer exists to service the packet.

Summary

This topic summarizes the key points discussed in this lesson.

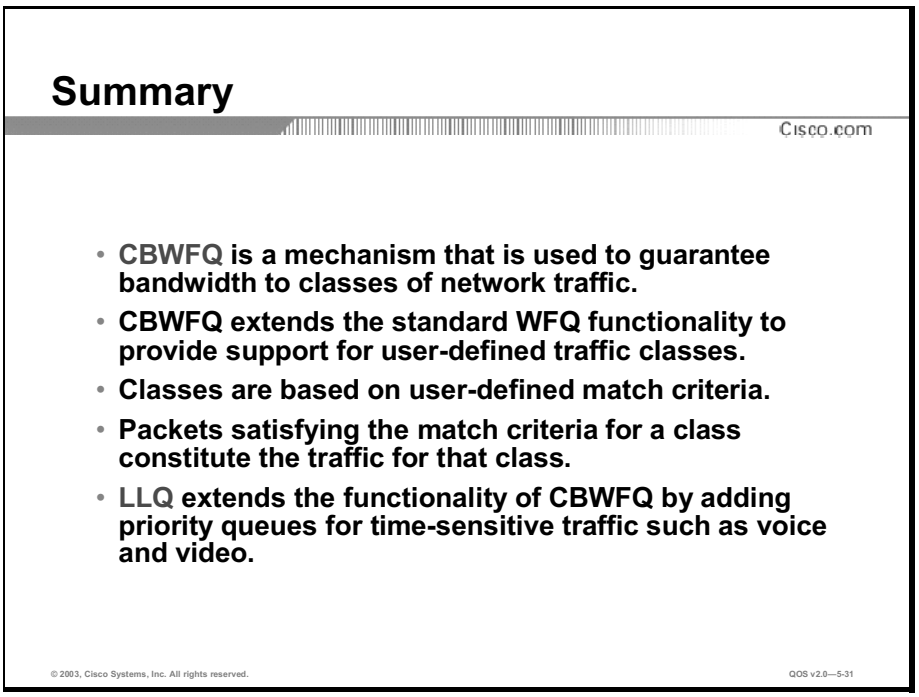

References

For additional information, refer to these resources:

- To learn more about CBWFQ, refer to “Class-Based Weighted Fair Queueing” at the following URL: http://www.cisco.com/en/US/products/sw/iosswrel/ps1830/products_feature_guide09186a0080087a84.html
- To learn more about LLQ, refer to “Low Latency Queueing” at the following URL: http://www.cisco.com/en/US/products/sw/iosswrel/ps1830/products_feature_guide09186a0080087b13.html
- To learn more about WFQ and DWFQ, refer to “Configuring Weighted Fair Queueing” at the following URL: http://www.cisco.com/en/US/products/sw/iosswrel/ps1828/products_configuration_guide_chapter09186a00800ca597.html
- To learn more about configuring CBWRQ with FRTS, refer to “Configuring Class Based Weighted Fair Queueing with FRTS” at the following URL: http://www.cisco.com/en/US/tech/tk713/tk237/technologies_configuration_example09186a008009486b.shtml

Next Steps

For the associated lab exercise, refer to the following section of the course Lab Guide:

- Lab Exercise 5-2: Configuring LLQ

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) How does CBWFQ extend the functionality of standard WFQ?

A) ensures that all flows go to a single queue

B) offers significant new classification options

C) provides support for user-defined traffic classes

D) provides a low-latency queue to ensure instant dispatch of voice packets

Q2) Which three of the following options can be used to define weights in CBWFQ? (Choose three.)

A) IP precedence

B) bandwidth (in kbps)

C) bandwidth percentage

D) percentage of available bandwidth

Q3) Which of the following is used when configuring bandwidth guarantees for CBWFQ?

A) class-map

B) policy-map

C) queue-policy

D) service-policy

Q4) Which of the following statements best represents the LLQ policy for high-priority classes?

A) They are shaped and cannot exceed bandwidth guarantee.

B) They are policed and cannot exceed bandwidth guarantee.

C) They are shaped to provide minimal delay with no packet loss.

D) They are not policed to ensure low-latency packets are always dispatched.

Q5) How does LLQ prevent starvation of non-priority queues?

A) alternate priority queues

B) pre-emptive non-priority queues

C) bandwidth policing of priority queues

D) Modified Round Robin servicing of all queues

Q6) Which header should be included when calculating bandwidth for LLQ for ATM?

A) cell tax overhead

B) Layer 2 encapsulation

C) RTP headers for voice

D) UTP and RTP headers for voice

Quiz Answer Key

Q1) C

Relates to: CBWFQ

Q2) B, C, D

Relates to: CBWFQ Architecture

Q3) B

Relates to: Configuring and Monitoring CBWFQ

Q4) B

Relates to: LLQ

Q5) C

Relates to: LLQ Architecture

Q6) A

Relates to: Configuring and Monitoring LLQ

LAN Congestion Management

Overview

Different Cisco Catalyst switches offer various mechanisms for providing QoS. This lesson provides an overview of the way several key Catalyst switches provide queuing support for QoS and how to configure those switches for QoS. PQ, WRR queuing, and WRR with an expedite queue (as used on the Cisco Catalyst 2950 switches) are explained in this lesson.

Relevance

The most effective QoS implementation begins at the edge of the network where Cisco Catalyst switches are normally deployed. Understanding how queuing works at the earliest stage of QoS deployment is key to designing and building an effective QoS-aware network.

Objectives

Upon completing this lesson you will be able to configure WRR on a Catalyst switch to manage LAN congestion. This includes being able to meet these objectives:

- Describe the different queuing capabilities available on Cisco Catalyst switches
- Explain how WRR works on a Catalyst 2950 switch
- Describe the commands required to configure PQ on Catalyst 2950 switches
- Describe the commands required to configure WRR on Catalyst 2950 switches
- Describe the commands required to monitor queuing on Catalyst 2950 switches

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Basic knowledge of internetworking with TCP/IP concepts and basic Catalyst IOS commands

Outline

The outline lists the topics included in this lesson.

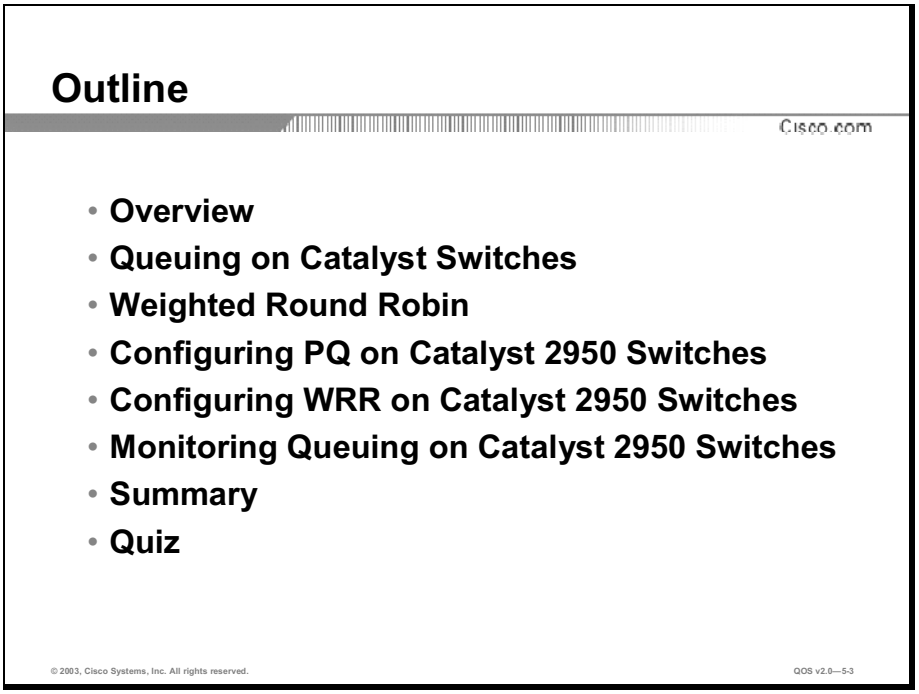

Queuing on Catalyst Switches

This topic explains the different methods of queuing that are available on key Catalyst switches.

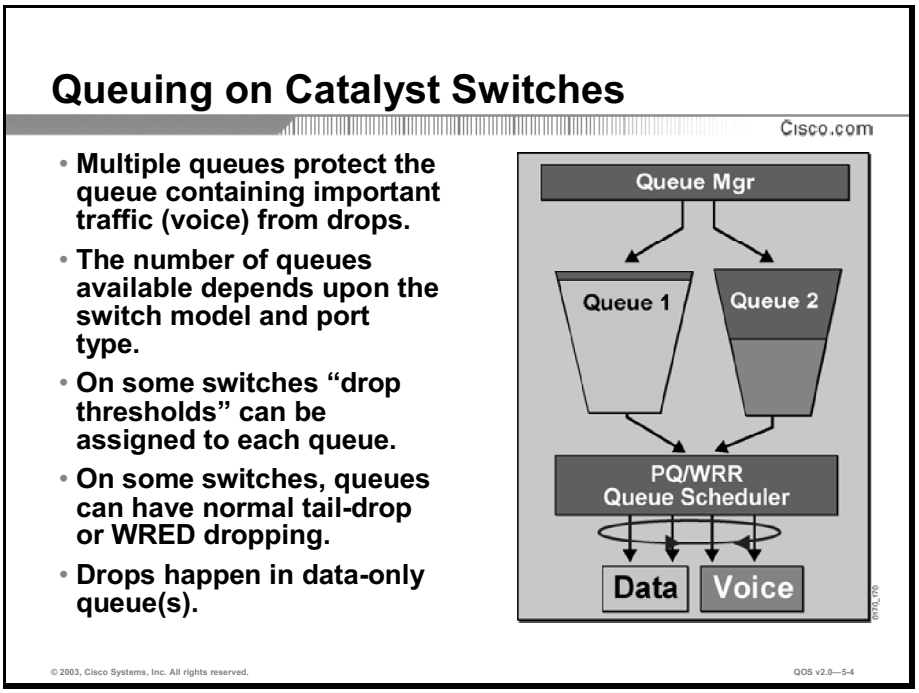

In a converged network, it is vital to ensure that voice traffic is not dropped. The use of multiple queues in Catalyst switches protects the queue containing important traffic (voice) from being dropped. Cisco Catalyst switches offer a variety of queuing capabilities depending upon the switch model and port type.

One of the key options that can be assigned to queues in most Catalyst switches is "drop thresholds." A queue can be assigned one or more drop thresholds. Packets are queued until the thresholds are exceeded.

For example, all packets with differentiated services code points (DSCPs) that are assigned to the first threshold are dropped until the threshold is no longer exceeded. However, packets assigned to a second threshold continue to be queued and sent as long as the second threshold is not exceeded. The thresholds are all specified as percentages ranging from 1 to 100. A value of 10 indicates a threshold when the buffer is 10 percent full.

On some switches, queues can have normal tail drop or WRED dropping. WRED is explained further in the next module of this course.

Drops will happen only in data-only queues. The purpose of using multiple queues is to prevent voice traffic from being dropped or delayed.

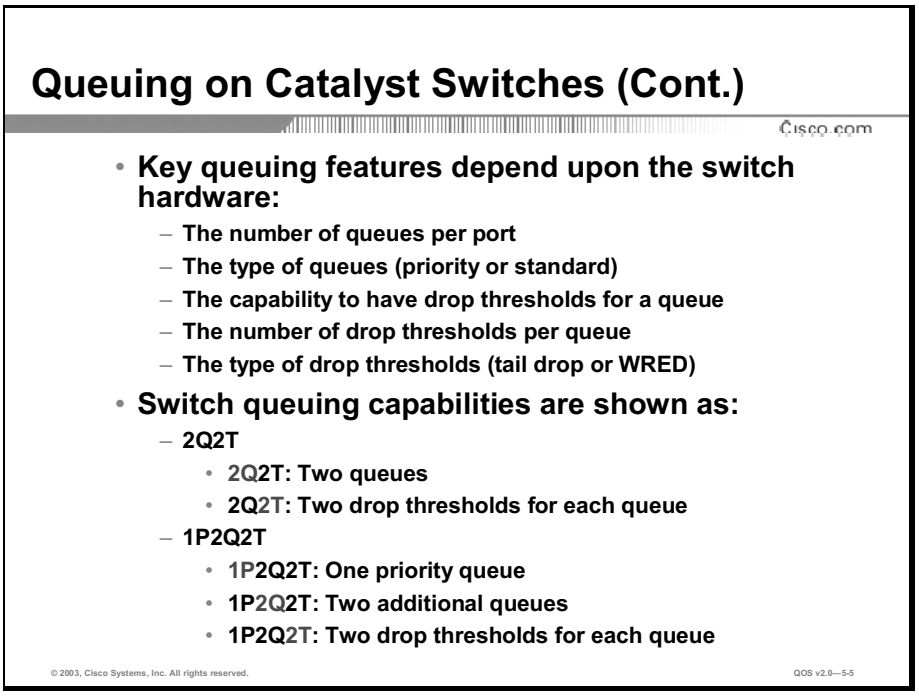

Different Cisco Catalyst switches offer different queuing capabilities. The queuing capabilities include:

- The number of queues per port
- The type of queues (priority or standard)
- The capability to have drop thresholds for a queue
- The number of drop thresholds per queue
- The type of drop thresholds (tail drop or WRED).When you view information on Cisco Catalyst switches, queuing information is displayed in an abbreviated format. For example:
  - 2Q2T indicates that the switch supports two standard queues and two drop thresholds per queue.
  - 1P2Q2T indicates that the switch supports one priority queue, two standard queues, and two drop thresholds per queue.

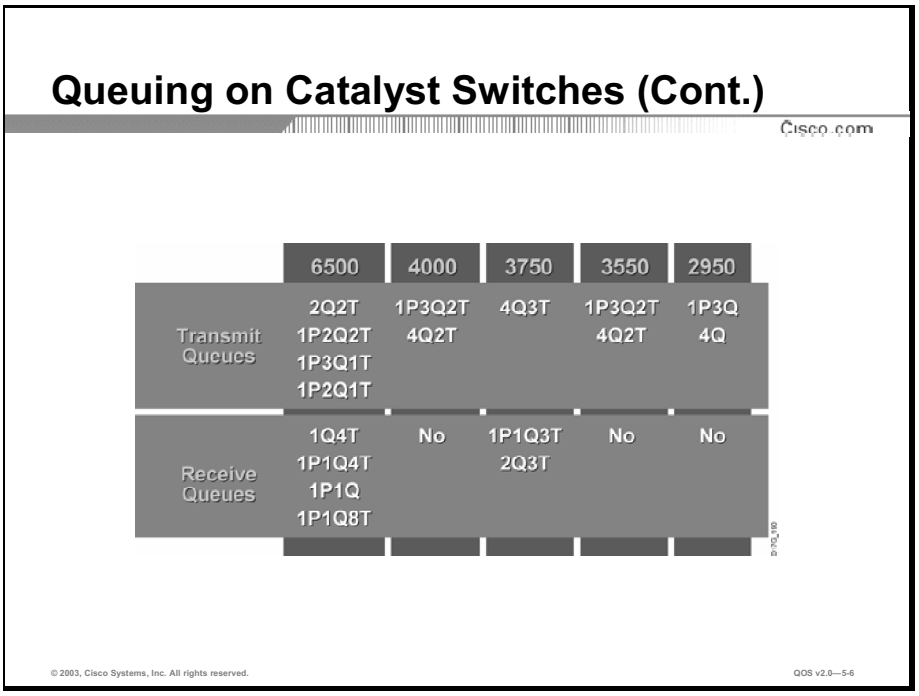

The chart shows the capabilities of Catalyst switch models.

6500 Series Catalyst Switches

The Catalyst 6500 provides both receive (Rx) and transmit (Tx) queues. The number and type of queues is dependent upon the line card.

The Rx queues are designed to protect voice traffic from delays or drops. An example of the implementation of an Rx queue with a priority queue and a standard queue with drop thresholds (1p1q4t) is:

- Frames with class of service (CoS) 5 go to the priority queue.
- Frames with CoS 0, 1, 2, 3, 4, 6, or 7 go to the standard Rx queue as follows:
  - Using standard receive-queue tail-drop threshold 1, the switch drops incoming frames with CoS 0 or 1 when the receive-queue buffer is 50 percent or more full.
  - Using standard receive-queue tail-drop threshold 2, the switch drops incoming frames with CoS 2 or 3 when the receive-queue buffer is 60 percent or more full.
  - Using standard receive-queue tail-drop threshold 3, the switch drops incoming frames with CoS 4 when the receive-queue buffer is 80 percent or more full.
  - Using standard receive-queue tail-drop threshold 4, the switch drops incoming frames with CoS 6 or 7 when the receive-queue buffer is 100 percent full.

An example of the implementation of a Tx queue with 2 queues and 2 drop thresholds (2q2t) on the Catalyst 6500 is:

- For 2q2t ports, each Tx queue has two tail-drop thresholds that function as follows:
  - Frames with CoS 0, 1, 2, or 3 go to the low-priority Tx queue (queue 1):
    - Using Tx queue 1, tail-drop threshold 1, the switch drops frames with CoS 0 or 1 when the low-priority transmit-queue buffer is 80 percent full.
    - Using Tx queue 1, tail-drop threshold 2, the switch drops frames with CoS 2 or 3 when the low-priority transmit-queue buffer is 100 percent full.
  - Frames with CoS 4, 5, 6, or 7 go to the high-priority Tx queue (queue 2):
    - Using Tx queue 2, tail-drop threshold 1, the switch drops frames with CoS 4 or 5 when the high-priority transmit-queue buffer is 80 percent full.
    - Using Tx queue 2, tail-drop threshold 2, the switch drops frames with CoS 6 or 7 when the high-priority transmit-queue buffer is 100 percent full.

4000 Series Catalyst Switches

On the Catalyst 4000 with a Supervisor III engine, each physical port has four Tx queues (egress queues). Each packet that needs to be transmitted is enqueued to one of the Tx queues. The Tx queues are then serviced based on the Tx queue scheduling algorithm.

When the final transmit DSCP is computed (including any markdown of DSCP), the transmit DSCP to Tx queue mapping configuration determines the Tx queue. The packet is placed in the Tx queue of the transmit port, determined from the transmit DSCP.

The four Tx queues for a transmit port share the available link bandwidth of that transmit port. You can set the link bandwidth to be shared differently among the Tx queues using bandwidth command in interface Tx queue configuration mode. With this command, you assign the minimum guaranteed bandwidth for each Tx queue.

By default, all queues are scheduled in a round robin manner.

You can configure Tx queue 3 on each port with higher priority. When Tx queue 3 is configured with higher priority, packets in Tx queue 3 are scheduled ahead of packets in other queues.

When Tx queue 3 is configured at a higher priority, the packets are scheduled for transmission before the other Tx queues only if it has not met the allocated bandwidth sharing configuration. Any traffic that exceeds the configured shape rate will be queued and transmitted at the configured rate. If the burst of traffic exceeds the size of the queue, packets will be dropped to maintain transmission at the configured shape rate.

Drop thresholds can be configured as tail drop or WRED.

3550 Catalyst Switches

On the 3550 Catalyst switches, the default scheduling method is strict priority. Strict priority scheduling is based on the priority of queues. Packets in the high-priority queue always transmit first, and packets in the low-priority queue do not transmit until all the higher-priority queues become empty.

CoS values can be assigned to queues during configuration. The default CoS to queue assignment is:

- CoS 6 to 7 placed in queue 4
- CoS 4 to 5 placed in queue 3
- CoS 2 to 3 placed in queue 2
- CoS 0 to 1 placed in queue 1

The switches support PQ, WRR scheduling, and WRR with a priority queue.

- The WRR scheduling algorithm ensures that lower-priority packets are not entirely starved for bandwidth and are serviced without compromising the priority settings administered by the network manager.
- WRR with a priority queue ensures that higher-priority packets will always get serviced first, ahead of other traffic in lower-priority queues. The priority queue is defined as queue 4.

Queue weights and queue depths are configurable.

Drop thresholds can be configured as tail drop or WRED.

Queuing on Catalyst Switches (Cont.)

Cisco.com

Catalyst 2950 Switches

- 4 Transmit Queues (1P3Q or 4Q)
- Need to configure PQ and ensure that CoS 5 traffic assigned to the PQ
  - Configurable PQ for queue 4
  - Configurable CoS to specific queue
  - Configurable queue weight

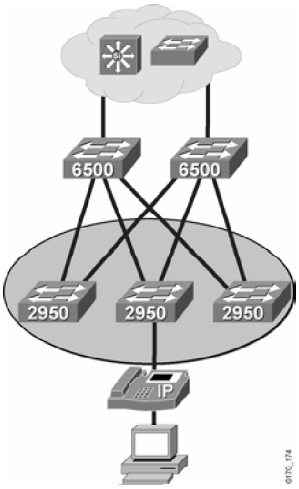

On the Catalyst 2950 series switches, the default scheduling method is strict priority. Strict priority scheduling is based on the priority of queues. Packets in the high-priority queue always transmit first; packets in the low-priority queue do not transmit until all the high-priority queues become empty.

CoS values can be assigned to queues during configuration. The default CoS to queue assignment is:

- CoS 6 to 7 placed in queue 4
- CoS 4 to 5 placed in queue 3
- CoS 2 to 3 placed in queue 2
- CoS 0 to 1 placed in queue 1

Catalyst 2950 switches support PQ, WRR scheduling, and WRR with a priority queue.

- The WRR scheduling algorithm ensures that lower-priority packets are not entirely starved for bandwidth and are serviced without compromising the priority settings administered by the network manager.
- WRR with a priority queue ensures that higher-priority packets will always get serviced first, ahead of other traffic in lower priority queues. The priority queue is defined as queue 4.

Queue weights are configurable.

Weighted Round Robin

This topic explains how WRR works on a Catalyst 2950 switch.

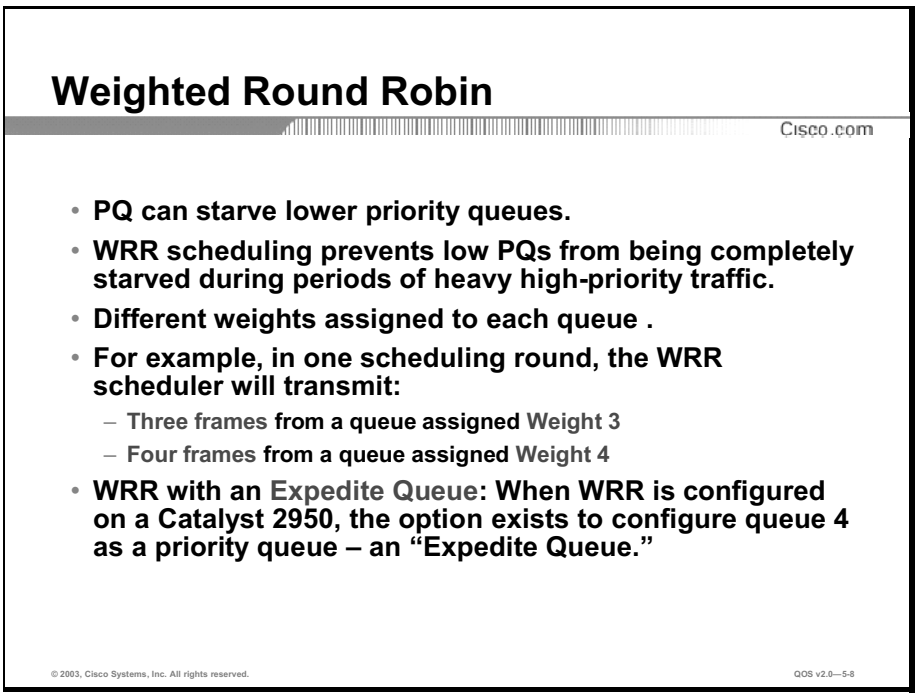

WRR scheduling requires that you specify a number that indicates the importance (weight) of the queue relative to the other CoS queues. WRR scheduling prevents the low-priority queues from being completely neglected during periods of high-priority traffic. The WRR scheduler transmits some packets from each queue in turn. The number of packets it sends corresponds to the relative importance of the queue. For example, if one queue has a weight of 3 and another has a weight of 4, three packets are sent from the first queue for every four that are sent from the second queue. By using this scheduling, low-priority queues have the opportunity to send packets even though the high-priority queues are not empty.

WRR scheduling with an expedite priority queue (also referred to as strict PQ) uses one of the egress queues as an expedite queue (queue 4 on a Catalyst 2950). The remaining queues participate in WRR. When the expedite queue is configured, it is a priority queue and is serviced until it is empty before the other queues are serviced by WRR scheduling. Actions at the egress interface include queuing and scheduling:

- Queuing evaluates the internal DSCP and determines which of the four egress queues in which to place the packet. The DSCP value is mapped to a CoS value, which selects one of the queues.
- Scheduling services the four egress queues based on their configured WRR weights and thresholds. One of the queues can be the expedite queue, which is serviced until empty before the other queues are serviced.
- Congestion avoidance techniques include tail drop and WRED on Gigabit-capable Ethernet ports and tail drop (with only one threshold) on 10/100 Ethernet ports.

With WRR, lower-priority queues have the opportunity to transmit packets even though the high-priority queues have not been emptied. With WRR with an expedite queue, one queue (queue 4 on the Catalyst 2950 and 3550) can be configured as an expedite priority queue All traffic from the expedite queue must be serviced before the remaining three queues are serviced. The expedite queue can be used to ensure that voice traffic incurs minimal delay and no drops.

Configuring PQ on Catalyst 2950 Switches

This topic explains how to configure PQ on the Cisco Catalyst 2950 switch.

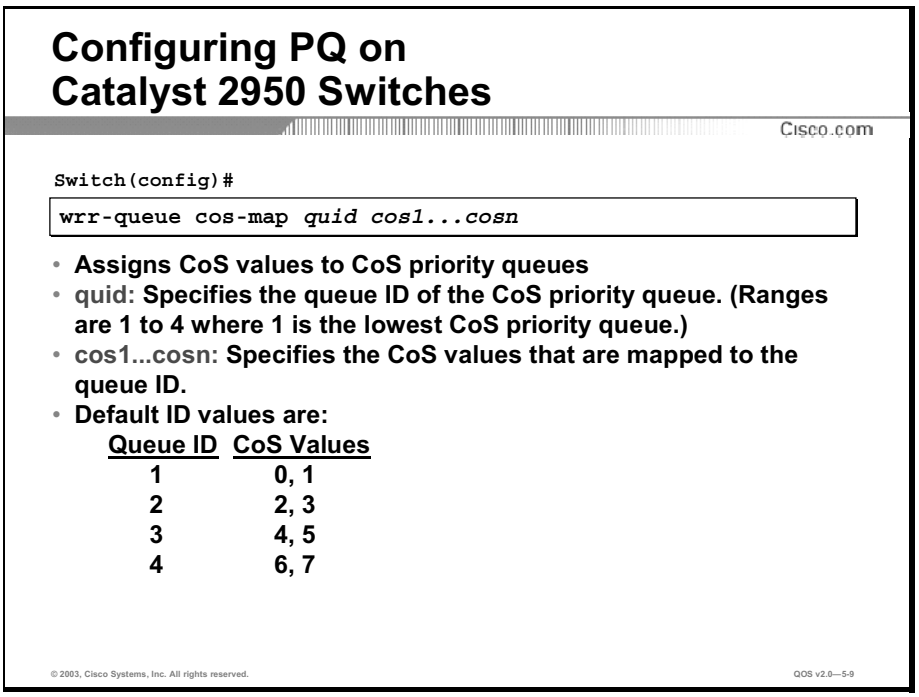

To configure PQ on the Catalyst 2950 switch, specify the queue ID of the CoS priority queue. (Ranges are 1 to 4 where 1 is the lowest CoS priority queue.) Then, specify the CoS values that are mapped to the queue ID.

wrr-queue cos-map *quid cos1...cosn*

Syntax Description

Parameter	Description
quid	The queue ID of the CoS priority queue. Ranges are 1 to 4 where 1 is the lowest CoS priority queue.
cos1...cosn	The CoS values that are mapped to the queue ID.

The default CoS to priority queue assignments are shown in the table.

Queue	1	2	3	4
CoS Values	0, 1	2, 2	4, 5	6, 7

Configuring WRR on Catalyst 2950 Switches

This topic explains how to configure WRR on the Catalyst 2950 switch.

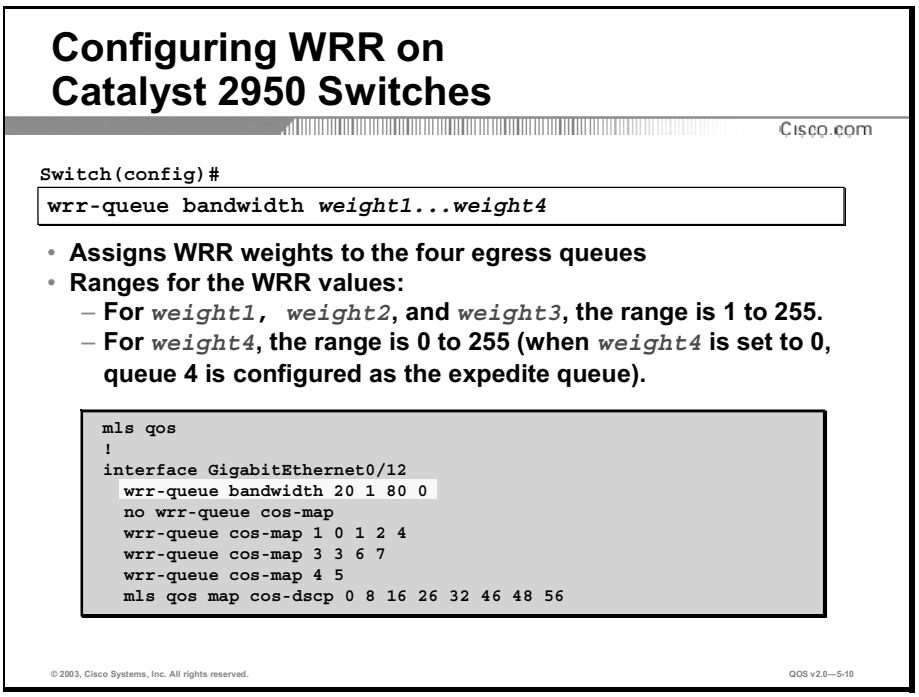

The **wrr-queue bandwidth** global configuration command is used to assign WRR weights to the four CoS priority queues on the Catalyst 2950 switch. Using the **no** form of this command will disable the WRR scheduler and enable the strict priority scheduler.

For weight 1, weight 2, and weight 3, the range is 1 to 255. The range for weight 4 is 0 to 255. Queues 1, 2, and 3 can be configured for WRR scheduling and queue 4 can be configured for strict priority scheduling. To configure queue 4 as the expedite queue, set weight4 to 0. When queue 4 is empty, packets from queues 1, 2, and 3 are sent according to the assigned WRR weights.

wrr-queue bandwidth *weight1...weight4*

Syntax Description

Parameter	Description
weight1...weight4	The ratio of weight 1, weight 2, weight 3, and weight 4 determines the weights of the WRR scheduler.

Note: In Cisco IOS software releases earlier than Release 12.1(12c)EA1, the ranges for all queues is 1 to 255.

In this example for the Catalyst 2950 switch, the following configuration has been made:

- The interface is set to GigabitEthernet 0/12
- Queue bandwidth (queue weight) is set to these weights:
  - Queue 1: 20
  - Queue 2: 1
  - Queue 3: 80
  - Queue 4: 0 (because this is 0, this is the expedite queue)
- The CoS map is set to its default settings.
- CoS is mapped to queues according to the following:
  - Queue 1: CoS 0, 1, 2, 4
  - Queue 2: No CoS assigned
  - Queue 3: CoS 3, 6, 7
  - Queue 4: CoS 5 (voice traffic goes to the expedite queue)

Note: This is the AutoQoS configuration for the Catalyst 2950.

Monitoring Queuing on Catalyst 2950 Switches

This topic lists the commands used for monitoring queuing on Catalyst 2950 switches.

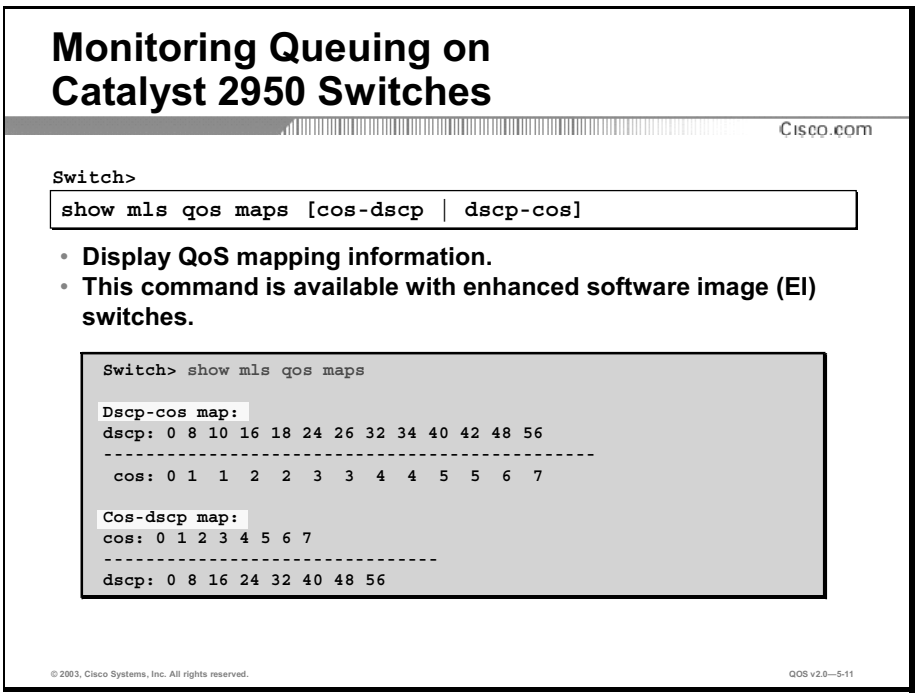

The **show mls qos maps** command is used to display QoS mapping information on the Catalyst 2950 switch. Maps are used to generate an internal DSCP value, which represents the priority of the traffic.

The **show mls qos maps** command is available only if the Catalyst 2950 switch is running the enhanced image (EI) software.

If the **show mls qos maps** command is used without any keywords, it will display all maps.

show mls qos maps [cos-dscp | dscp-cos]

Syntax Description

Parameter	Description
`cos-dscp`	(Optional) Display CoS-to-DSCP map.
`dscp-cos`	(Optional) Display DSCP-to-CoS map.

Monitoring Queuing on Catalyst 2950 Switches (Cont.)

Cisco.com

```
Switch>
show wrr-queue bandwidth
```

- **Display the WRR bandwidth allocation for the CoS priority queues**

```
Switch> show wrr-queue bandwidth
WRR Queue : 1 2 3 4
Bandwidth : 10 20 30 40
```

```
Switch>
show wrr-queue cos-map
```

- **Display the mapping of the CoS priority queues**

```
Switch> show wrr-queue cos-map
CoS Value      : 0 1 2 3 4 5 6 7
Priority Queue : 1 1 2 2 3 3 4 4
```


The **show wrr-queue bandwidth** command is used to display the WRR bandwidth allocation for the four CoS priority queues.

The **show wrr-queue cos-map** command is used to display the mapping of the CoS priority queues.

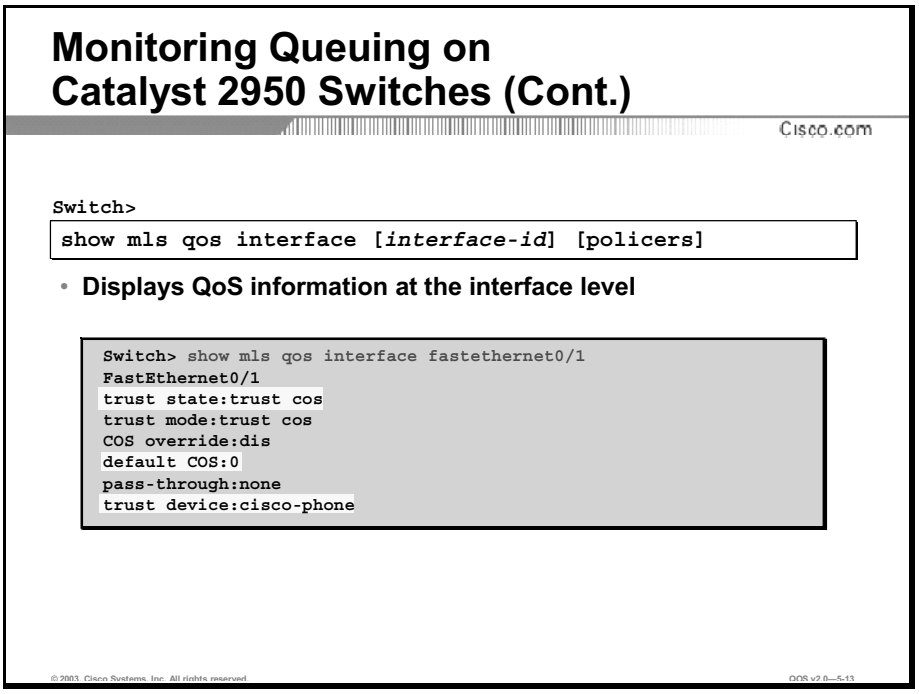

The **show mls qos interface** command is used to display QoS information at the interface level. Although it will be visible in CLI help strings, the **policers** keyword is available only when the Catalyst 2950 switch is running the enhanced software image.

show mls qos interface [*interface-id***] [policers]**

Syntax Description

Parameter	Description
interface-id	(Optional) Display QoS information for the specified interface.
policers	(Optional) Display all the policers configured on the interface, their settings, and the number of policers unassigned. Available only when the switch is running the EI software.

Summary

This topic summarizes the key points discussed in this lesson.

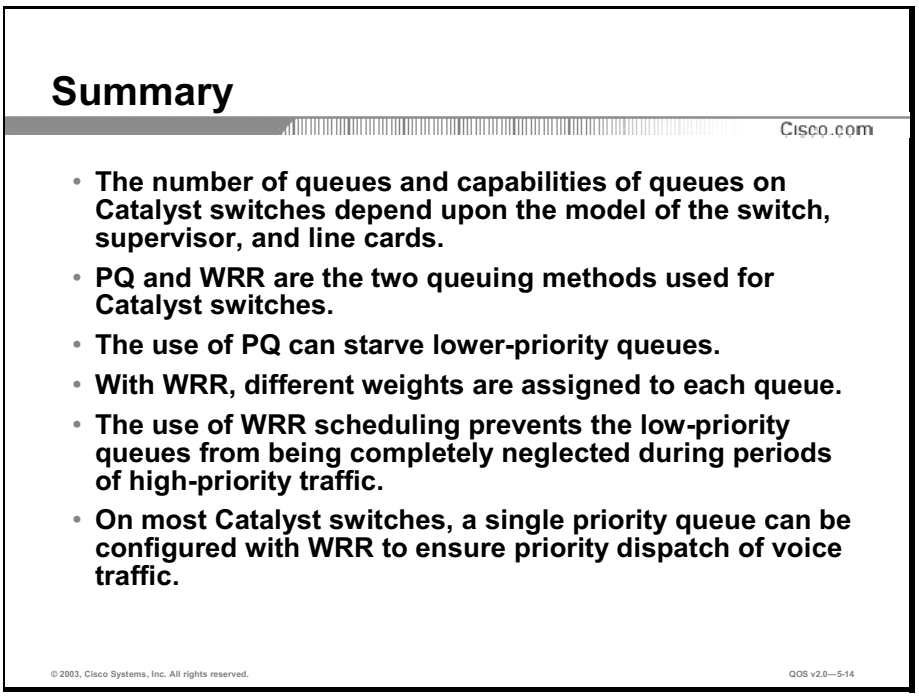

References

For additional information, refer to these resources:

- To learn more about queuing on the Cisco Catalyst 3550 series switches, refer to "Configuring QoS" at the following URL: http://www.cisco.com/en/US/products/hw/switches/ps646/products_configuration_guide_chapter09186a008014f36e.html#1127419
- To learn more about queuing on the Cisco Catalyst 2950 series switches, refer to "Configuring QoS" at the following URL: http://www.cisco.com/en/US/products/hw/switches/ps628/products_configuration_guide_chapter09186a008014f2c0.html#1025310
- To learn more about queuing on the Cisco Catalyst 4000 series switches, refer to "Configuring QoS" at the following URL: http://www.cisco.com/en/US/products/hw/switches/ps663/products_configuration_guide_chapter09186a008007eddd.html
- To learn more about queuing on the Cisco Catalyst 6500 series switches, refer to "Configuring QoS" at the following URL: http://www.cisco.com/en/US/products/hw/switches/ps708/products_configuration_guide_chapter09186a0080121d31.html#36454

Next Steps

For the associated lab exercise, refer to the following section of the course Lab Guide:

- Lab Exercise 5-3: Queuing on Catalyst Switches

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) Which three of the following are queuing features that differ among Catalyst switch models? (Choose three.)

A) number of queues per port

B) the speed at which the queue transmits

C) the type of queues (priority or standard)

D) capability to have drop thresholds for a queue

Q2) Which of the following names would represent the queuing features of a switch with 4 queues, one priority queue, and 2 standard queues, each with 2 drop thresholds?

A) 1P3Q2T

B) 1P4S2T

C) 1P4Q2T

D) 1P3S2T

Q3) When using WRR on a Cisco Catalyst 2950, what is the most certain way to ensure minimal delay for voice traffic?

A) configure an expedite queue

B) configure LLQ

C) use MDRR

D) configure the high queue with a very high weight factor

Q4) When WRR is configured on a switch with four Tx queues, given that weights have been assigned to each of the queues as follows and that all queues are full, how many packets from queue 4 would be dispatched every time a packet from queue 2 is dispatched?

Queue	Weight
4	8
3	4
2	2
1	1

A) 2

B) 4

C) 16

D) 24

Q5) Which three of the following queuing algorithms are available on the Cisco Catalyst 2950 switch? (Choose three.)

A) Priority Queuing

B) Weighted Round Robin

C) Modified Deficit Round Robin

D) Weighted Round Robin with an expedite queue

Q6) Which two represent options for configuring drop thresholds? (Choose two.)

A) WRED

B) tail drop

C) expedite drop

D) priority drop

Q7) On the Cisco Catalyst 2950 switch, what is the effect of configuring the `wrr-queue bandwidth weight` parameter of queue 4 as "0"?

A) queue 4 is disabled

B) queue 4 has the lowest weight

C) queue 4 becomes the expedite queue

D) queues 1, 2, and 3 obtain extra bandwidth

Q8) On the Cisco Catalyst 2950 switch, which command would you use to assign CoS values 0, 1, 2, and 3 to queue 1?

E) **wrr-queue cos –map 1 0 1 2 3**

F) **wrr-queue cos –map 0 1 2 3 1**

G) **wrr-queue cos –map 0,1,2,3,1**

H) **wrr-queue cos –map 0 1/ 2/3/1**

Q9) Which command is used to display the trust state of a port on a Cisco Catalyst 2950 switch?

A) show mls qos maps

B) show mls qos interface

C) show wrr-queue bandwidth

D) show mls interface trust-state

Quiz Answer Key

Q1) A, C, D

Relates to: Queuing on Catalyst Switches

Q2) A

Relates to: Queuing on Catalyst Switches

Q3) A

Relates to: Configuring WRR on Catalyst 2950 Switches

Q4) B

Relates to: Weighted Round Robin

Q5) A, B, D

Relates to: Configuring WRR on Catalyst 2950 Switches

Q6) A, B

Relates to: Queuing on Catalyst Switches

Q7) C

Relates to: Configuring WRR on Catalyst 2950 Switches

Q8) A

Relates to: Configuring PQ on Catalyst 2950 Switches

Q9) B

Relates to: Monitoring Queuing on Catalyst 2950 Switches

Module Assessment

Overview

Use this assessment to test what you learned in this module. The correct answers and solutions are found in the Module Assessment Answer Key.

Quiz: Congestion Management

Complete the Quiz to assess what you have learned in the module.

Objectives

This activity tests your knowledge on how to meet these objectives:

- Identify and explain the operation of basic queuing algorithms including FIFO, priority, and round-robin queuing
- Describe the functions of hardware and software queuing
- Configure weighted fair queuing to manage congestion
- Configure CBWFQ and LLQ to manage congestion
- Configure WRR on a Catalyst switch to manage LAN congestion

Instructions

Complete these steps:

Step 1 Answer all questions in this quiz by selecting the best answer(s) to each question

Step 2 Verify your results against the answer key located at the end of this section

Step 3 Review the topics in this module that relate to the questions that you answered incorrectly.

Q1) In FIFO queuing, what happens to packets when the queue is full?

A) they are delayed
B) they are dropped
C) they are retransmitted
D) they are moved to another queue

Q2) What happens when the highest-priority queue becomes congested in priority queuing algorithm?

A) all the other queues starve
B) tail dropping focuses on the highest-priority queue
C) other queues are served on a round-robin basis
D) packets in the highest-priority queue are moved to a lower-priority queue

Q3) In WRR implementation using a byte threshold as a measurement of each queue share of bandwidth, given an MTU of 2000 and a byte-count of 4000, what would the router do with the next packet for queue 2 (800 bytes) if the router had just dispatched two packets from queue 2 (sizes 2000 and 1600) to the hardware queue?

A) tail drop the next packet
B) dispatch the first packet from the next queue
C) dispatch the next packet to the hardware queue
D) split the packet and transmit the first 400 bytes

Q4) On a network device connecting a LAN and a WAN, what type of congestion would be likely for traffic moving from the LAN to the WAN?

A) bursty
B) transient
C) persistent
D) fluctuating

Q5) Given that the hardware queue is NOT full, how will the next packet be serviced by the software queue?

A) software queue will be bypassed
B) software queue will enqueue the packet
C) software queue will expedite the packet
D) software queue will only meter the packet

Q6) How does WFQ implement tail dropping?

A) drops the last packet to arrive
B) drops all non-voice packets first
C) drops the lowest-priority packets first
D) drops packets from the most aggressive flows

Q7) Consider that a WFQ system has a modest hold-queue limit of ten (HQO=10) and a congestive discard threshold of eight (CDT=8), and that there are already eight (8) packets in the system. If a newly arriving packet had the worst finish time of all packets in the system, what would happen to the packet?

A) it would be dropped
B) it would be enqueued
C) it would be buffered until a spot in a queue came open
D) it would be dispatched

Q8) Which of the following is the default dropping scheme for CBWFQ?

A) RED
B) WRED
C) tail-drop
D) class-based policing

Q9) What does LLQ bring to CBWFQ?

A) strict priority scheduling
B) alternate priority scheduling
C) non-policed queues for low latency traffic
D) special voice traffic classification and dispatch

Q10) What type of traffic should you limit the use of the priority command to?

A) low-latency data traffic
B) voice-like, CBR traffic
C) high volume, VBR traffic
D) video and teleconferencing, available (ABR) traffic

Q11) When WRR with an expedite queue has been configured on a Cisco Catalyst 2950 switch, which queue is emptied before any other queues are serviced?

A) queue 1
B) queue 2
C) queue 3
D) queue 4

Scoring

You have successfully completed the quiz for this lesson when you earn a score of 80 percent or better.

Module Assessment Answer Key

Q1) A

Relates to: Introduction to Queuing

Q2) A

Relates to: Introduction to Queuing

Q3) C

Relates to: Introduction to Queuing

Q4) C

Relates to: Queuing Implementations

Q5) A

Relates to: Queuing Implementations

Q6) D

Relates to: FIFO and WFQ

Q7) A

Relates to: FIFO and WFQ

Q8) C

Relates to: CBWFQ and LLQ

Q9) A

Relates to: CBWFQ and LLQ

Q10) B

Relates to: CBWFQ and LLQ

Q11) D

Relates to: LAN Congestion Management

Module Summary

This topic summarizes the key points discussed in this module.

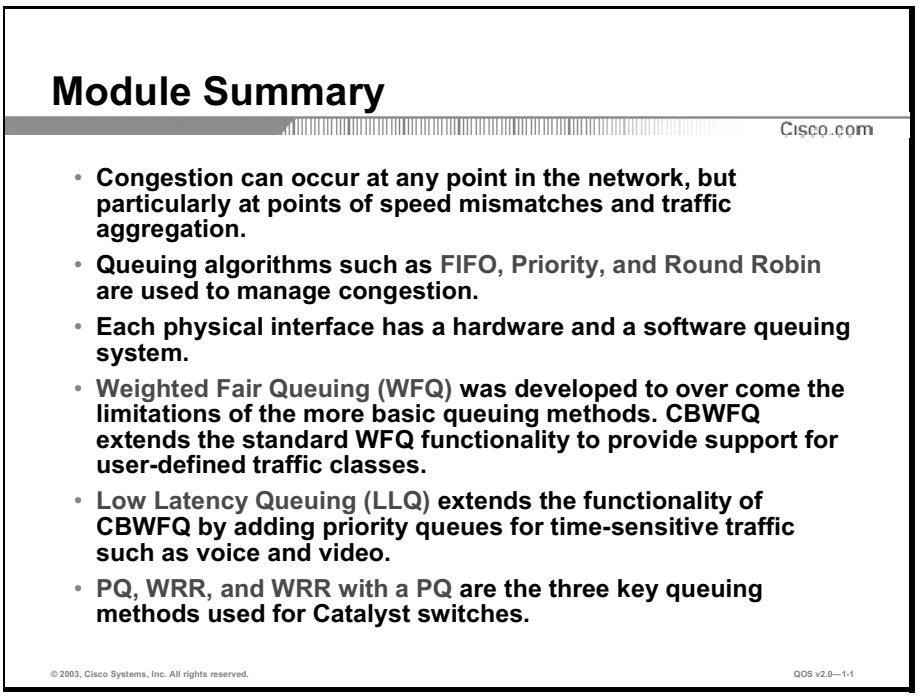

Effective congestion management is key to QoS in converged networks. Low latency traffic such as voice and video must be constantly moved to high priority queues in order to ensure reasonable quality.

Cisco routers offer a variety of simple (FIFO, PQ, and CQ) and sophisticated (WFQ, CBWFQ, and LLQ) queuing algorithms to provide effective congestion management on converged networks. LLQ, the most sophisticated, was specifically designed to provide the highest QoS to voice traffic.

Cisco switches offer a variety of queuing capabilities depending upon the model of switch being used. On most Catalyst switches, three queuing methods are available for use: PQ, WRR, and WRR with a priority queue.

Module 6

Congestion Avoidance

Overview

Congestion is a normal occurrence in networks today. Whether congestion occurs as a result of a lack of buffer space, network aggregation points, or a low-speed wide area link, many congestion management techniques exist to ensure specific applications and traffic classes are given their share of available bandwidth when congestion occurs. Congestion management does, however, come at a price. When congestion occurs, some traffic is delayed or even dropped at the expense of other traffic. When drops occur, different problems may arise, which can exacerbate the congestion such as retransmissions and TCP global synchronization in TCP/IP networks.

Congestion avoidance mechanisms are designed to reduce the negative effects of congestion by penalizing the most aggressive traffic streams as software queues begin to fill. This module discusses the problems with TCP congestion management and the benefits of deploying congestion avoidance mechanisms in a network.

Module Objectives

Upon completing this module, you will be able to use Cisco QoS congestion avoidance mechanisms to reduce the effects of congestion on the network.

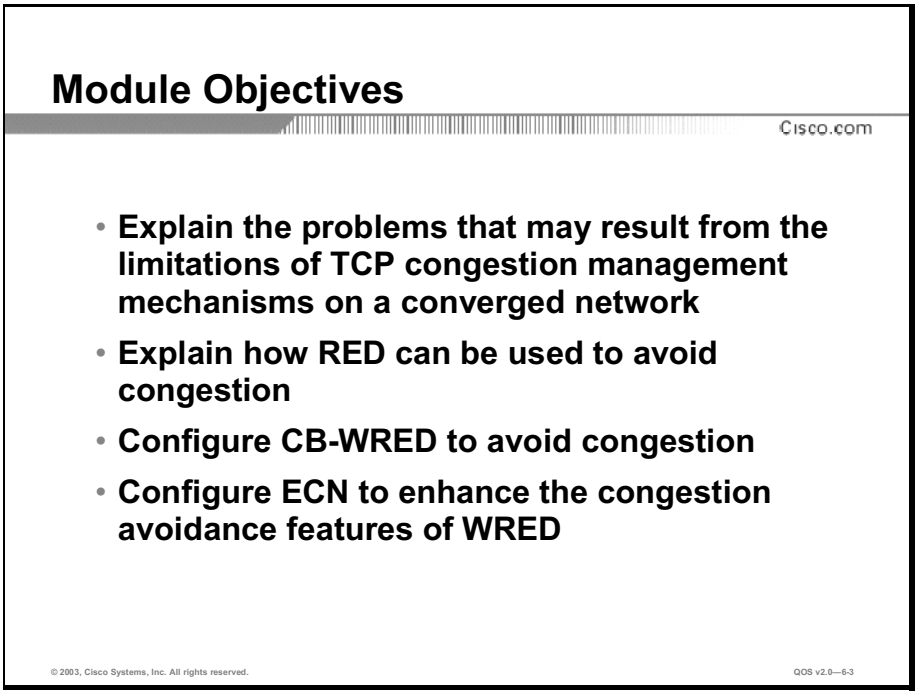

Module Outline

The outline lists the components of this module.

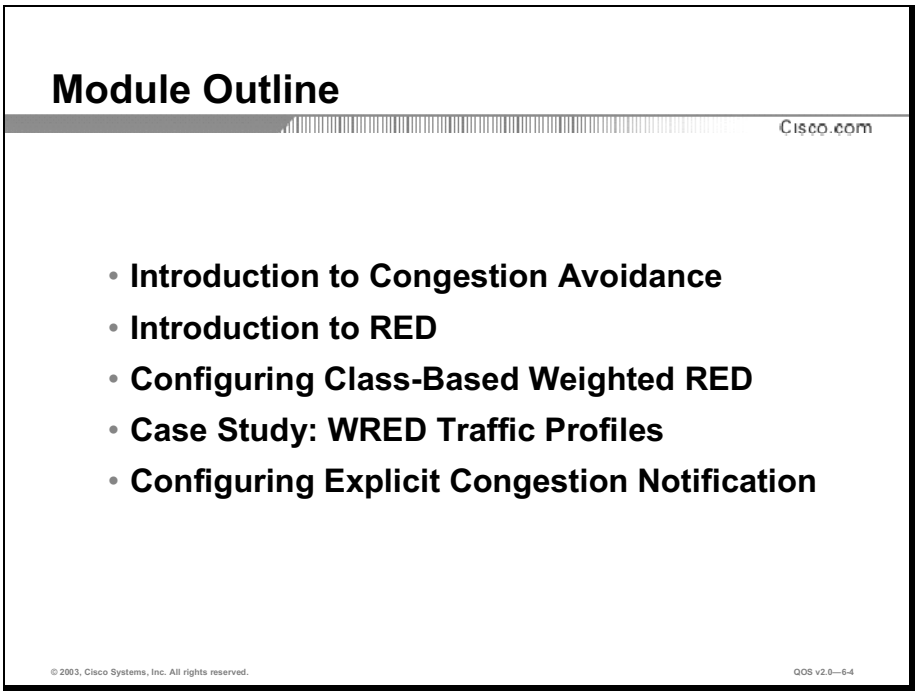

Introduction to Congestion Avoidance

Overview

This lesson explains the behavior of the TCP when hosts send and receive packets. This lesson also explains traffic management mechanisms that are used by TCP during periods of congestion and the effects of packet loss on TCP sessions.

Relevance

Congestion avoidance mechanisms are important tools for reducing the effects of congestion on networks. To understand the problems these mechanisms solve, and the benefits they bring, it is first helpful to understand how TCP recognizes and responds to congestion.

Objectives

Upon completing this lesson, you will be able to explain the problems that may result from the limitations of TCP congestion management mechanisms on a converged network. This includes being able to meet these objectives:

- Explain the behavior of TCP senders and receivers
- Explain how TCP responds to congestion
- Describe tail drop as a congestion control mechanism
- Explain the drawbacks of tail drop

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Basic knowledge of internetworking with TCP/IP concepts

Outline

The outline lists the topics included in this lesson.

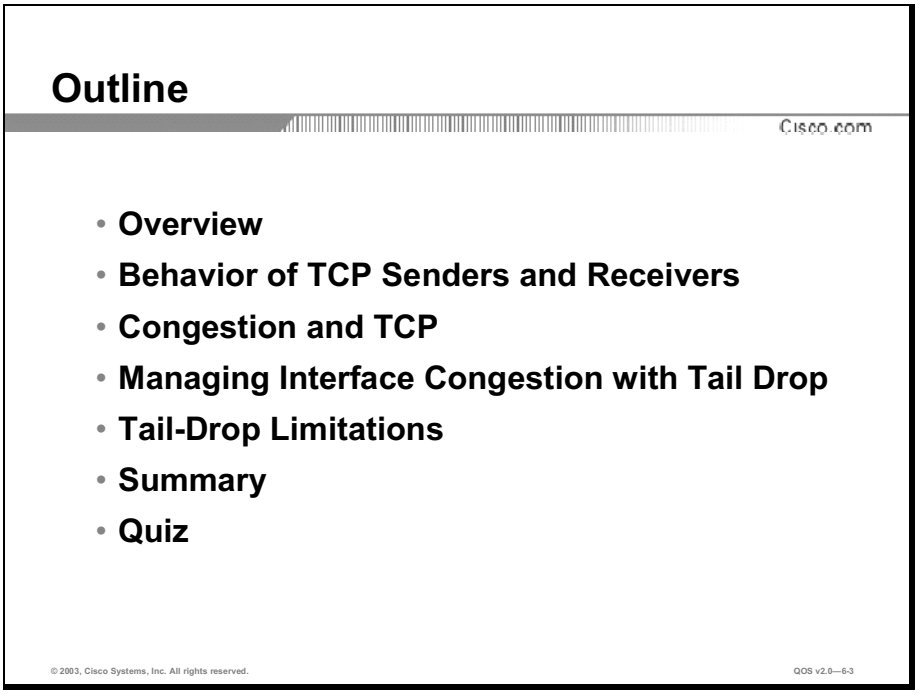

Behavior of TCP Senders and Receivers

This topic describes the behavior of TCP senders and receivers when sending packets.

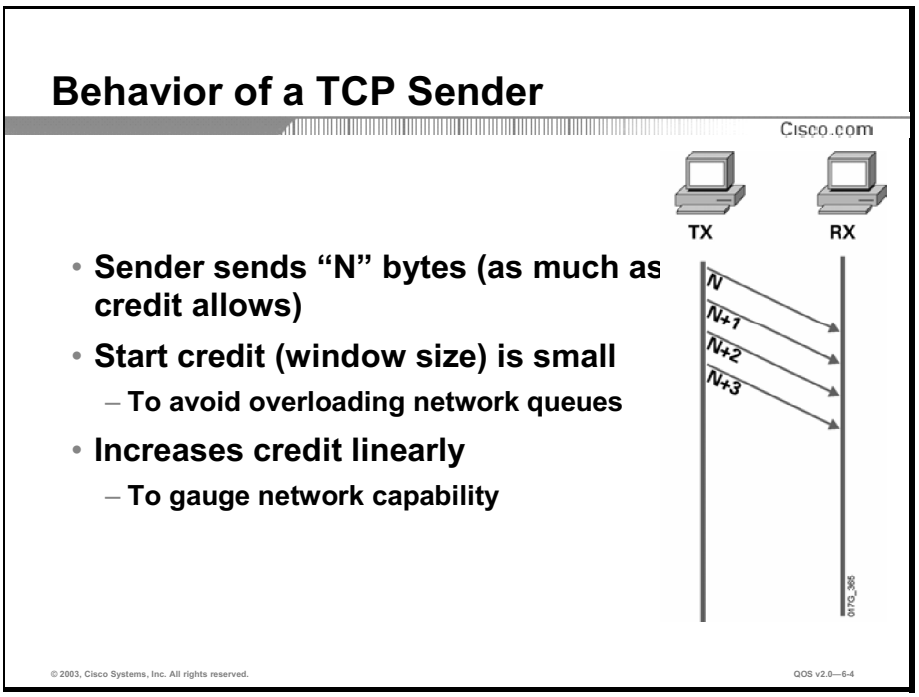

Before any data is transmitted using TCP, a connection must first be established between the transmitting and receiving hosts. When the connection is initially established, the two hosts must agree on certain parameters that will be used during the communication session. One of the parameters that must be decided is called the window size, or how many data bytes to transmit at a time. Initially, TCP sends a small number of data bytes, and then exponentially increases the number sent. For example, a TCP session originating from host A begins with a window size of 1 and therefore sends one packet. When host A receives a positive ACK from the receiver, it increases its window size to 2. Host A then sends 2 packets, receives a positive ACK, and increases its window size to 4, and so on.

Note: TCP tracks window size by byte count. For purposes of illustration packets (N) is used.

In traditional TCP, the maximum window size is 64 Kb (65,535 bytes). Extensions to TCP specified in RFC 1323, allow for tuning TCP by extending the maximum TCP window size to 2^{30} bytes. TCP extensions for high performance, although supported on most operating systems, may not be supported on your system.

Example: Example of Windowing in TCP

After connecting to an Internet website, a file transfer using the FTP download is initiated. Watching the progress of the transfer, it is noticed that the bytes per second counter steadily increases during the file transfer. This is an example of TCP windowing in action.

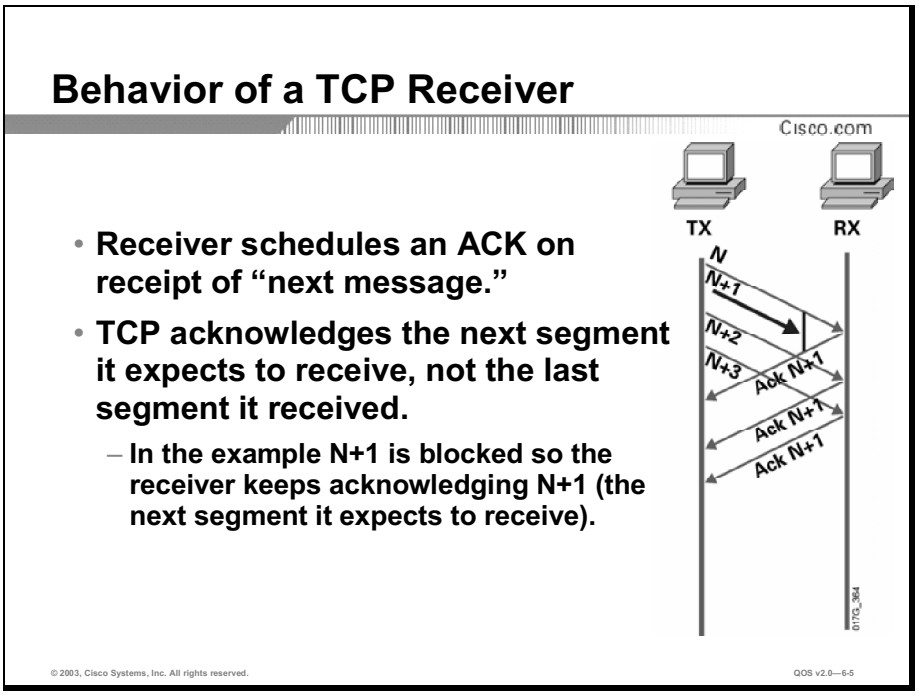

When the receiver receives a data segment, it checks that data segment sequence number (byte count). If the data received fills in the next sequence of numbers expected, it indicates that the data segment was received in order. The receiver then:

- Delivers all the data that it holds to the target application
- Updates the sequence number to reflect the next byte number in expected order

When this process is complete, it performs one of the following actions:

- Immediately transmits an acknowledgment (ACK) to the sender
- Schedules an ACK to be transmitted to the sender after a short delay

The ACK notifies the sender that the receiver received all data segments up to but not including the byte number in the new sequence number. Receivers usually try to send an ACK in response to alternating data segments they receive. They send the ACK because for many applications, if the receiver waits out a small delay, it can efficiently piggyback its reply acknowledgment on a normal response to the sender. However, when the receiver receives a data segment out of order, it immediately responds with an ACK to direct the sender to retransmit the lost data segment.

Note:	TCP tracks window size by byte count. For purposes of illustration packets (N) is used.

Congestion and TCP

This topic describes the TCP response to lost data packets.

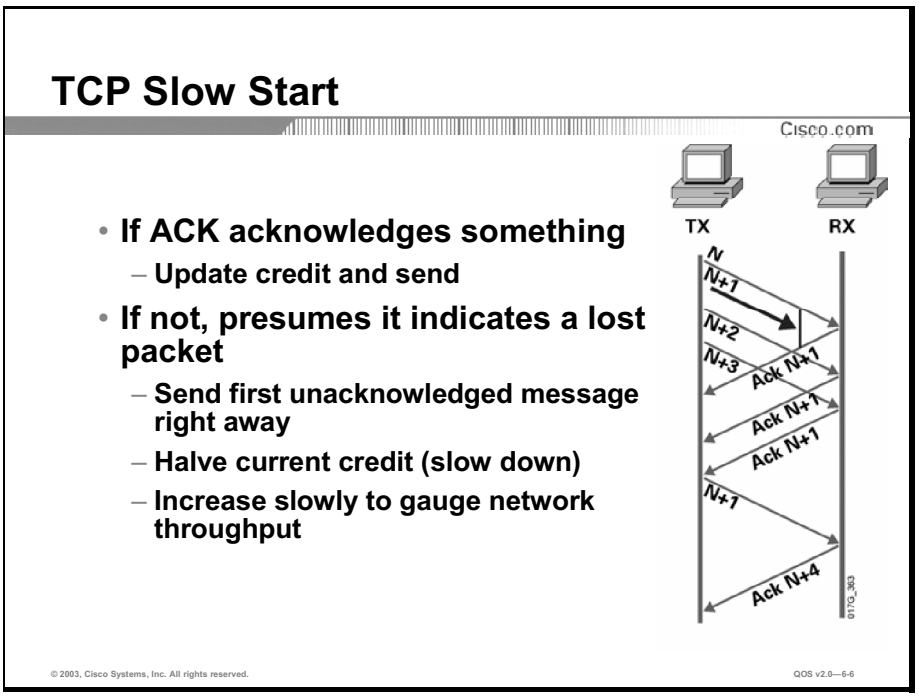

When the sender receives an ACK, it determines if any data is outstanding:

- If no data is outstanding, the sender determines that the ACK is a keepalive, meant to keep the line active, and it does nothing.
- If data is outstanding, the sender determines whether the ACK indicates that the receiver has received some or none of the data.
  - If the ACK acknowledges receipt of some data sent, the sender determines if new credit has been granted to allow it to send more data.
  - When the ACK acknowledges receipt of none of the sent data and there is outstanding data, the sender interprets the ACK to be a repeatedly sent ACK. This condition indicates that some data was received out of order, forcing the receiver to remit the first ACK, and that a second data segment was received out of order, forcing the receiver to remit the second ACK. In most cases, the receiver would receive two segments out of order because one of the data segments had been dropped.

When a TCP sender detects a dropped data segment, it retransmits the segment. Then it slows its transmission rate so that the rate is half of what it was before the drop was detected. This is known as the TCP slow-start mechanism.

In the figure, a station transmits three packets to the receiving station. Unfortunately, the first packet is dropped somewhere in the network. Therefore the receiver sends an ACK 1, to request the missing packet. Because the transmitter does not know if the ACK was just a duplicate ACK, it will wait for three ACK 1 packets from the receiver. Upon receipt of the third ACK, the missing packet, packet 1 is resent to the receiver. The receiver now sends an ACK 4 indicating it has already received packets 2 and 3 and is ready for the next packet.

Multiple Drops in TCP

- **If multiple drops occur in the same session:**
  - **Current TCPs wait for time-out**
  - **Selective acknowledge (SACK) may be a work around**
  - **New "fast retransmit phase" takes several round-trip times (RTTs) to recover**

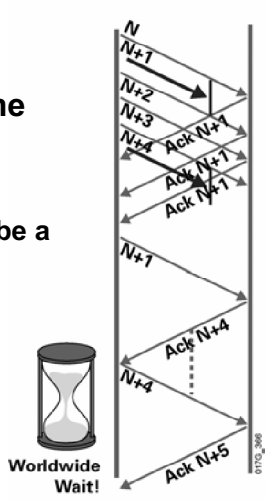

Although the TCP slow-start behavior is appropriately responsive to congestion, problems can arise when multiple TCP sessions are concurrently carried on the same router and all TCP senders slow down transmission of packets at the same time.

If a TCP sender does not receive acknowledgement for sent segments, it cannot wait indefinitely before it assumes that the data segment it sent never arrived at the receiver. TCP senders maintain the retransmission timer as a means of signaling a segment retransmission. The retransmission timer can impact TCP performance. If the retransmission timer is too short, duplicate data will be sent into the network unnecessarily. If the retransmission timer is too long, the sender will wait (remain idle) for too long, slowing down the flow of data.

The selective acknowledgment (SACK) mechanism, as proposed in RFC 2018, can improve the time it takes for the sender to recover from multiple packet losses as non-contiguous blocks of data can be acknowledged, and the sender only has to retransmit data that is actually lost. SACK is used to convey extended acknowledgement information from the receiver to the sender to inform the sender of non-contiguous blocks of data that have been received. Using the example in the slide, instead of sending back an ACK N + 1, the receiver can send a SACK N + 1 and also indicate back to the sender that N + 3 has been correctly received with the SACK option.

In standard TCP implementations, a TCP sender can only discover that a single packet has been lost each round-trip time (RTT), causing poor TCP performance when multiple packets are lost.

Remember, the sender must receive three duplicate ACK packets before it realizes that a packet has been lost. As a result of receiving the third ACK, the sender will immediately send the segment referred to by the ACK. This TCP behavior is called fast retransmit.

Managing Interface Congestion with Tail Drop

This topic describes the default mechanism for managing interface congestion, tail drop.

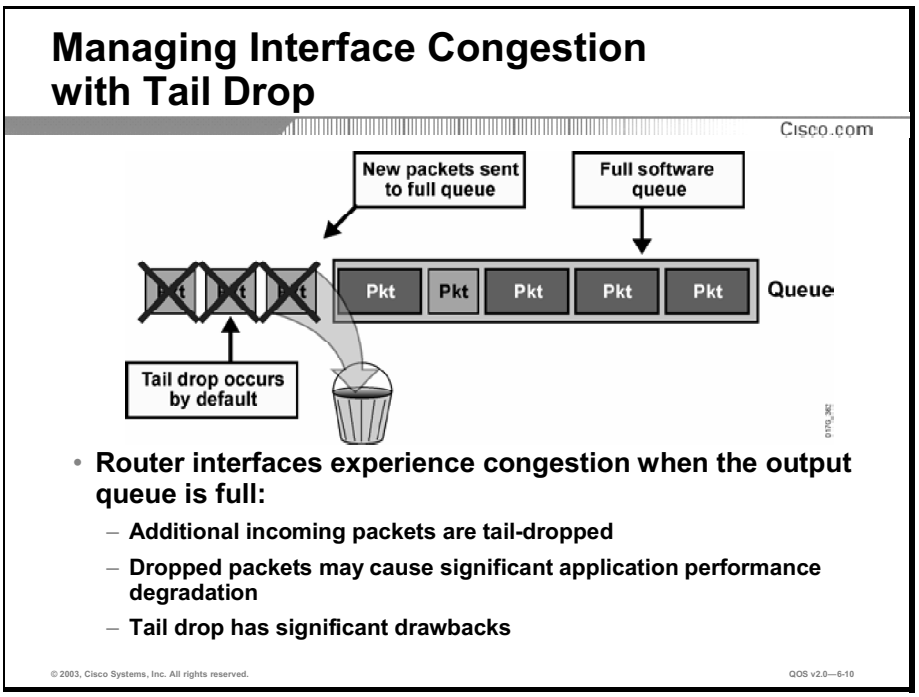

When an interface on a router cannot transmit a packet immediately, the packet is queued, either in an interface transmit (Tx) ring, or the interface output hold queue, depending on the switching path that is used. Packets are then taken out of the queue and eventually transmitted on the interface.

If the arrival rate of packets to the output interface exceeds the router capability to buffer and forward traffic, the queues increase to their maximum length and the interface becomes congested. Tail drop is the default queuing response to congestion. Tail drop treats all traffic equally and does not differentiate between classes of service. Applications may suffer performance degradation due to packet loss caused by tail drop. When the output queue is full and tail drop is in effect, all packets trying to enter (at the tail of) the queue are dropped until the congestion is eliminated and the queue is no longer full.

Weighted fair queuing (WFQ), if configured on an interface, has a more elaborate scheme for dropping traffic, as it is able to punish the most aggressive flows via its congestive discard threshold (CDT)-based dropping algorithm. Unfortunately, WFQ does not scale to backbone speeds.

Tail-Drop Limitations

This topic describes the limitations of using tail drop as a congestion management mechanism.

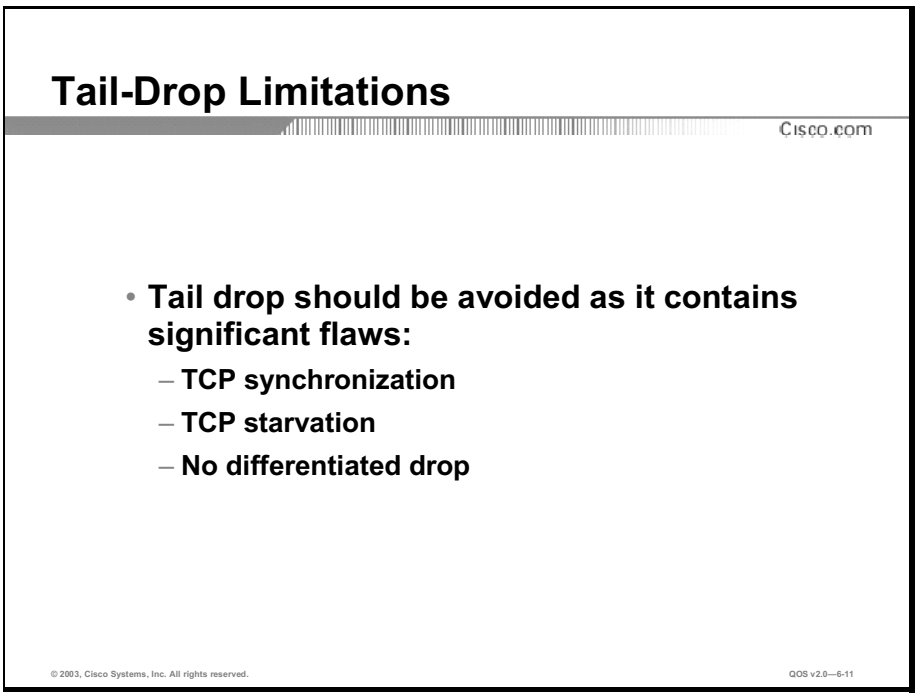

The simple tail-drop scheme unfortunately does not work very well in environments with a large number of TCP flows or in environments in which selective dropping is required. Understanding the network interaction between TCP stack intelligence and dropping is required to implement a more efficient and fair dropping scheme, especially in service provider environments.

Tail drop has the following shortcomings:

- When congestion occurs, dropping affects most of the TCP sessions, which simultaneously back off and then restart again. This causes inefficient link utilization at the congestion point (TCP global synchronization).
- TCP starvation, where all buffers are temporarily seized by aggressive flows, and normal TCP flows experience buffer starvation.
- There is no differentiated drop mechanism, and therefore premium traffic is dropped in the same way as best-effort traffic.

TCP Synchronization

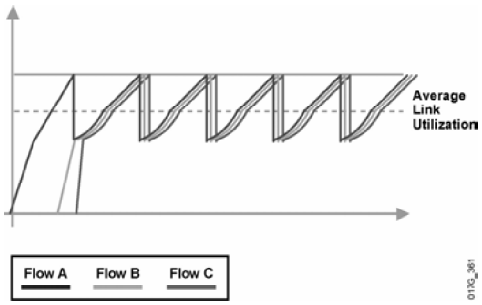

- **Multiple TCP sessions start at different times.**
- **TCP window sizes are increased.**
- **Tail drops cause many packets of many sessions to be dropped at the same time.**
- **TCP sessions restart at the same time (synchronized).**

A router can handle multiple concurrent TCP sessions. There is a high probability that when traffic exceeds the queue limit, it vastly exceeds the limit due to the bursty nature of packet networks. However, there is also a high probability that excessive traffic depth caused by packet bursts are temporary and that traffic does not stay excessively deep except at points where traffic flows merge, or at edge routers.

If the receiving router drops all traffic that exceeds the queue limit, as is done by default (with tail drop), many TCP sessions then simultaneously go into slow start. Consequently, traffic temporarily slows down to the extreme and then all flows slow-start again. This activity creates a condition called global synchronization.

Global synchronization occurs as waves of congestion crest only to be followed by troughs during which the transmission link is not fully utilized. Global synchronization of TCP hosts can occur because packets are dropped all at once. Global synchronization manifests when multiple TCP hosts reduce their transmission rates in response to packet dropping. When congestion is reduced their transmission rates are increased. The most important point is that the waves of transmission known as global synchronization result in significant link underutilization.

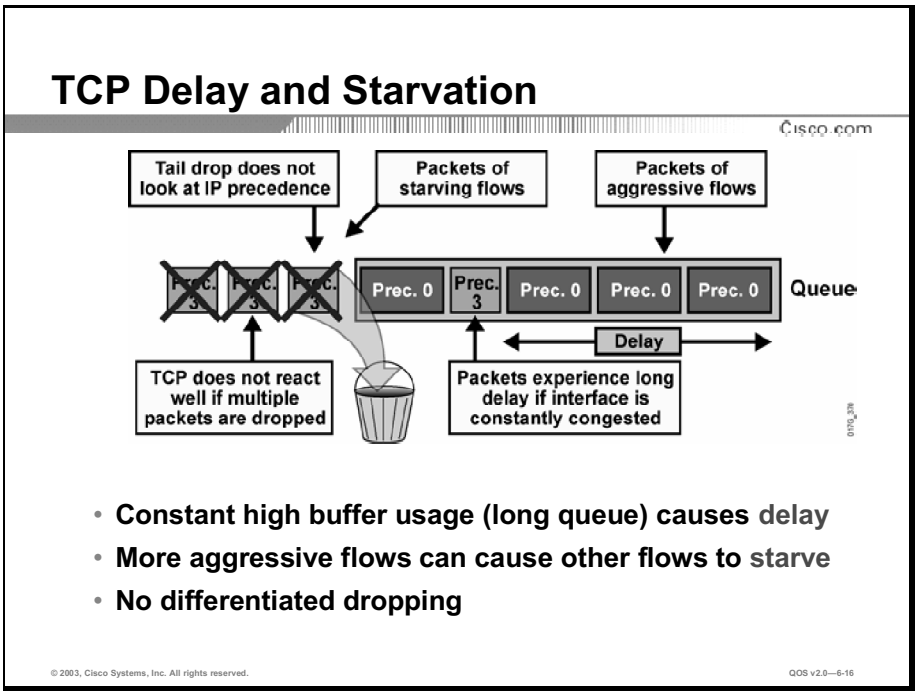

During periods of congestion, packets are queued up to the full queue length, which also causes increased delay for packets that are already in the queue. In addition, queuing, being a probabilistic mechanism, introduces unequal delays for packets of the same flow, thus producing jitter.

Another TCP-related phenomenon that reduces optimal throughput of network applications is TCP starvation. When multiple flows are established over a router, some of these flows may be much more aggressive as compared to others. For instance, when a file transfer application TCP transmit window increases, it can send a number of large packets to its destination. The packets immediately fill the queue on the router, and other, less aggressive flows can be starved because there is no differentiated treatment indicating which packets should be dropped. As a result, these less aggressive flows are tail-dropped at the output interface.

Based on the knowledge of TCP behavior during periods of congestion, it can be concluded that tail drop is not the optimal mechanism for congestion avoidance and therefore should not be used. Instead, more intelligent congestion avoidance mechanisms should be used that slow down traffic before actual congestion occurs.

Summary

This topic summarizes the key points discussed in this lesson.

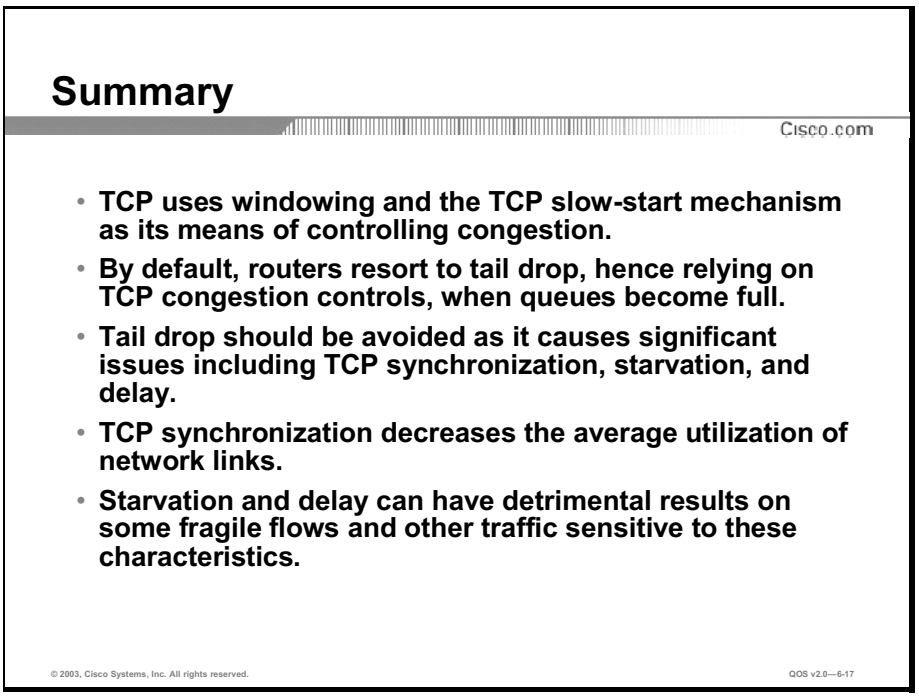

References

For additional information, refer to these resources:

- Further details on TCP slow start can be found in RFC 2001 at the following URL: http://www.faqs.org/rfcs/rfc2001.html
- For more information, and a list of operating systems supporting RFC 1323, refer to: http://www.psc.edu/networking/perf_tune.html#table.
- For a detailed discussion of TCP protocol behavior see Geoff Huston, Telstra, “TCP Performance,” *Internet Protocol Journal*, Vol. 3, No. 2, June 2000, at the following URL: http://www.cisco.com/warp/public/759/ipj_3-2/ipj_3-2_tcp.html
- For a detailed discussion of TCP congestion behavior see Geoff Huston, Telstra, “The Future for TCP,” *Internet Protocol Journal*, Vol. 3, No. 3, September 2000, at the following URL: http://www.cisco.com/warp/public/759/ipj_3-3/ipj_3-3_futureTCP.html

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) What are two main drawbacks of using tail drop as a means of congestion control? (Choose two.)

A) global synchronization

B) small window sizes

C) head of line blocking

D) starvation

Q2) How does a TCP receiver respond to the receipt of an out-of-order segment?

A) It will send an ICMP error to the sender and close the TCP session.

B) It will store all received segments in a buffer and reorder the packets.

C) It immediately sends an ACK to the sender indicating the sequence number of the missing segment.

D) It will wait for the last segment to be transmitted to ensure the missing segment was not delayed by an alternate network path or network congestion.

Q3) Which two reasons confirm why tail drop is inadequate for avoiding congestion? (Choose two.)

A) Tail drop drops packets at the receiver, not the sender, and therefore does not slow the cause of congestion.

B) Tail drop treats all traffic equally and does not differentiate between classes of service.

C) Tail drop can result in many sessions simultaneously utilizing the TCP slow-start mechanism at the same time.

D) Tail drop depends upon TCP to control window sizes as a means of congestion control.

Q4) What are two ways that a TCP sender interprets an unacknowledged packet? (Choose two.)

A) it assumes the packet was dropped due to congestion and retransmits the packet

B) it will wait for the ACK timer to expire and then close the TCP session

C) it will send an ACK to probe the receiver

D) it will reduce the window size to ½ of its value before the unacknowledged packet was detected

Q5) What is the largest negative impact of global synchronization in TCP networks?

A) TCP sessions are closed

B) significant link under-utilization

C) poor TCP response to congestive discard techniques

D) Reduces the ability of TCP to utilize bursting

Q6) What is the default router response to a full queue resulting from congestion?

A) as congestion increases (measured by the average size of the queue), selectively drop packets proportional to the increase in queue size

B) dynamically increase the queue buffer size to accommodate newly arriving packets

C) drop packets from the end of the queue to make room to buffer the incoming packets

D) tail drops all incoming packets until the queue has scheduled packets and is no longer full

Quiz Answer Key

Q1) A, D

Relates to: Tail Drop Limitations

Q2) C

Relates to: Behavior of TCP Senders and Receivers

Q3) C, D

Relates to: Tail-Drop Limitations

Q4) A, D

Relates to: Congestion and TCP

Q5) B

Relates to: Tail-Drop Limitations

Q6) D

Relates to: Managing Interface Congestion with Tail Drop

Introduction to RED

Overview

Congestion avoidance techniques monitor network traffic loads in an effort to anticipate and avoid congestion at common network bottleneck points. Congestion avoidance is achieved through packet dropping using a more complex dropping technique than simple tail drop. This lesson introduces the congestion avoidance technique random early detection (RED) and its scalable dropping method, which is suitable for low- and high-speed networks.

Relevance

Congestion avoidance techniques offer a viable alternative to the default router congestion response, tail drop. RED is one of the most commonly used congestion avoidance techniques used in high-speed transit networks.

Objectives

Upon completing this lesson, you will be able to explain how RED can be used to avoid congestion. This includes being able to meet these objectives:

- Describe RED and how it can be used to prevent congestion
- Describe the elements of a RED traffic profile
- Describe the different drop modes of RED
- Describe the effects of RED on TCP traffic
- Identify the points in a network where congestion avoidance can most effectively be employed

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Basic knowledge of internetworking with TCP/IP concepts

Outline

The outline lists the topics included in this lesson.

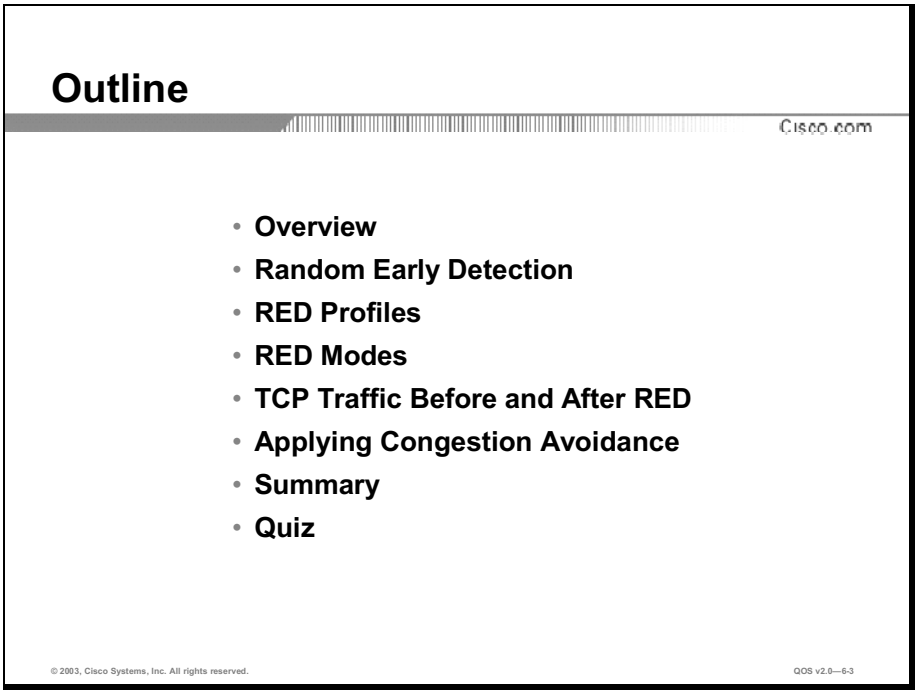

Random Early Detection

This topic describes the purpose and function of RED.

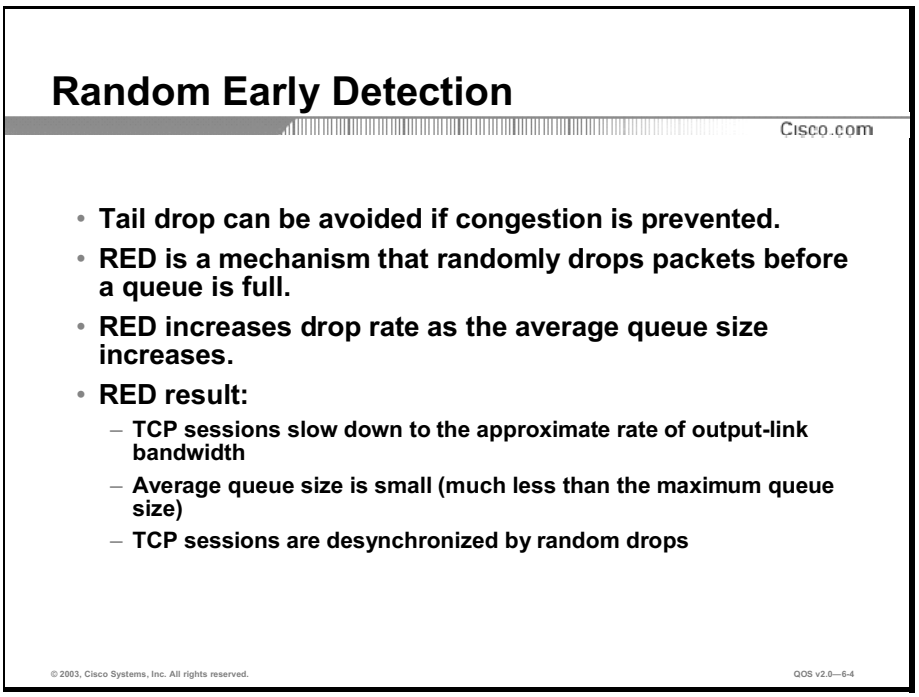

RED is a dropping mechanism that randomly drops packets before a queue is full. The dropping strategy is based primarily on the average queue length—that is, when the average size of the queue increases, RED will be more likely to drop an incoming packet than when the average queue length is shorter.

Because RED drops packets randomly, it has no per-flow intelligence. The rationale is that an aggressive flow will represent most of the arriving traffic and therefore it is more probable that RED will drop a packet of an aggressive session. RED therefore punishes more aggressive sessions with higher statistical probability and is, therefore, able to somewhat selectively slow down the most significant cause of congestion. Directing one TCP session at a time to slow down allows for full utilization of the bandwidth, rather than utilization that manifests itself as crests and troughs of traffic.

As a result of implementing RED, the problem of TCP global synchronization is much less likely to occur and TCP can utilize link bandwidth more efficiently. In RED implementations, the average queue size also decreases significantly, as the possibility of the queue filling up is reduced. This is because of very aggressive dropping in the event of traffic bursts, when the queue is already quite full.

RED distributes losses over time and normally maintains a low queue depth while absorbing traffic spikes. RED can also utilize IP precedence or differentiated services code point (DSCP) bits in packets to establish different drop profiles for different classes of traffic.

RED Profiles

This topic describes the elements of a RED traffic profile that is used to implement the RED packet dropping strategy.

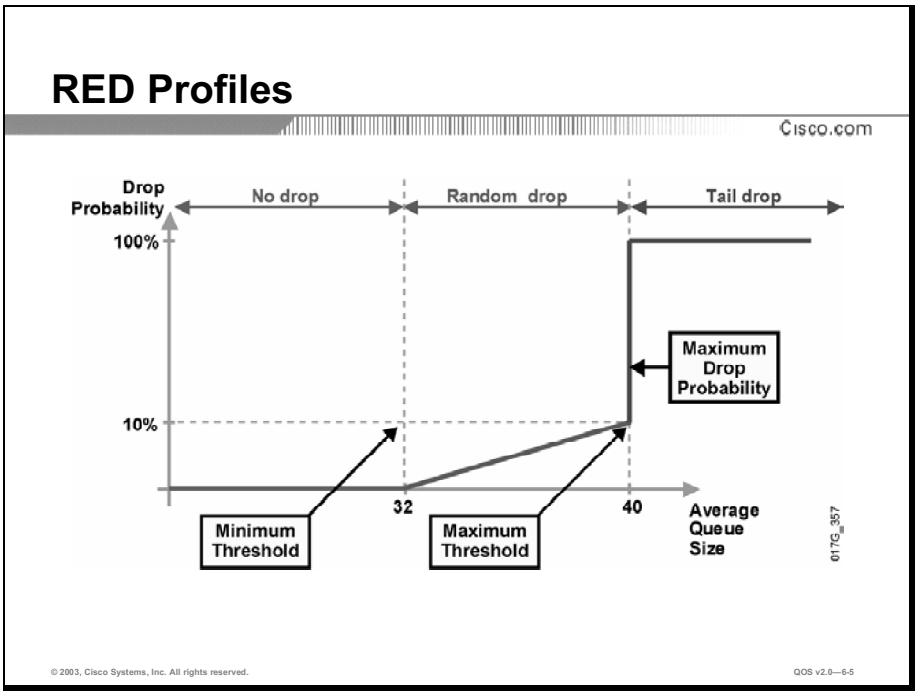

A RED traffic profile is used to determine the packet dropping strategy and is based on the average queue length. The probability of a packet being dropped is based on three configurable parameters contained within the RED profile:

- **Minimum threshold:** When the average queue length 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.
- **Maximum threshold:** When the average queue size is above the maximum threshold, all packets are dropped.
- **Mark probability denominator:** This is the fraction of packets that are 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. The linear increase of packet drops from the minimum threshold (0 drops) to the maximum threshold is based on this parameter and the queue size between the minimum and maximum thresholds.

The minimum threshold value should be set high enough to maximize the link utilization. If the minimum threshold is too low, packets may be dropped unnecessarily, and the transmission link will not be fully used.

The difference between the maximum threshold and the minimum threshold should be large enough to avoid global synchronization. If the difference is too small, many packets may be dropped at once, resulting in global synchronization.

The mark probability has the effect of controlling the number of packets that are dropped when the average queue length reaches the maximum threshold. If the value is set too low, the behavior of tail drop will rapidly be approached, resulting in too many dropped packets. If the value is set too large, RED dropping can be rendered ineffective.

Example: RED Traffic Profile

The following is an example of a RED traffic profile.

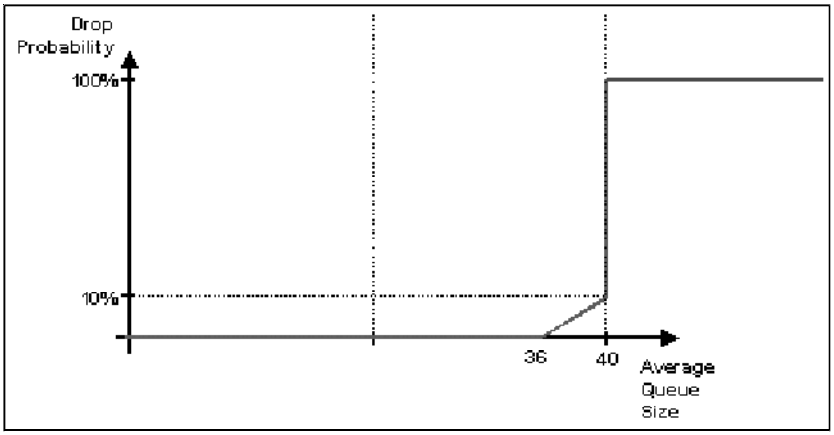

Sample RED Traffic Profile

In the RED traffic profile shown in the example, the minimum threshold is 36. RED will not drop any packets until the average queue size is 36 or greater.

If the average queue size is greater than or equal to 36, but less than 40, RED will randomly discard packets from the more aggressive traffic flows. The rate of packet discard will linearly increase as the length of the average queue increases.

The maximum threshold for this RED profile is 40. When the average queue length is 40 or greater, all packets will be discarded until the average queue size is below 40. At the moment in time that the average queue size reaches 40, the router will be dropping 1 out of every 10 packets (mark probability denominator = 10).

RED Modes

This topic describes the different packet drop modes of RED.

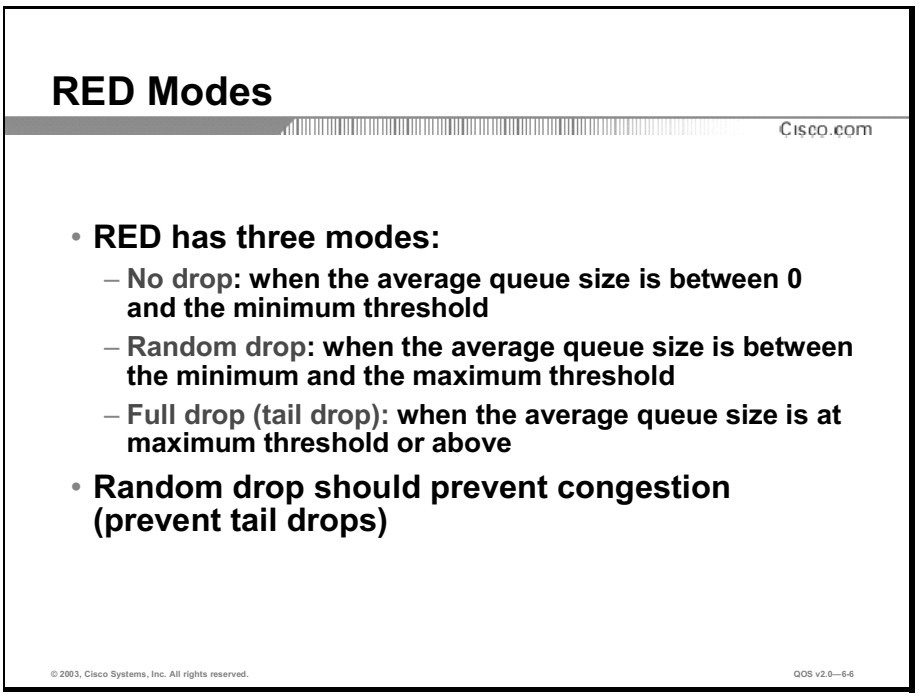

Based on the average queue size, RED has three dropping modes:

- When the average queue size is between 0 and the configured minimum threshold, no drops occur and all packets are queued.
- When the average queue size is between the configured minimum threshold, and the configured maximum threshold, random drop occurs, which is linearly proportional to the mark probability denominator and the average queue length.
- When the average queue size is at or higher than the maximum threshold, RED performs full (tail) drop in the queue. This event is unlikely, as RED should slow down TCP traffic ahead of congestion. If a lot of non-TCP traffic is present, RED cannot effectively drop traffic to reduce congestion, and tail drops are likely to occur.

TCP Traffic Before and After RED

This topic describes the effects of using RED on TCP traffic by comparing TCP traffic flows both before and after the application of RED.

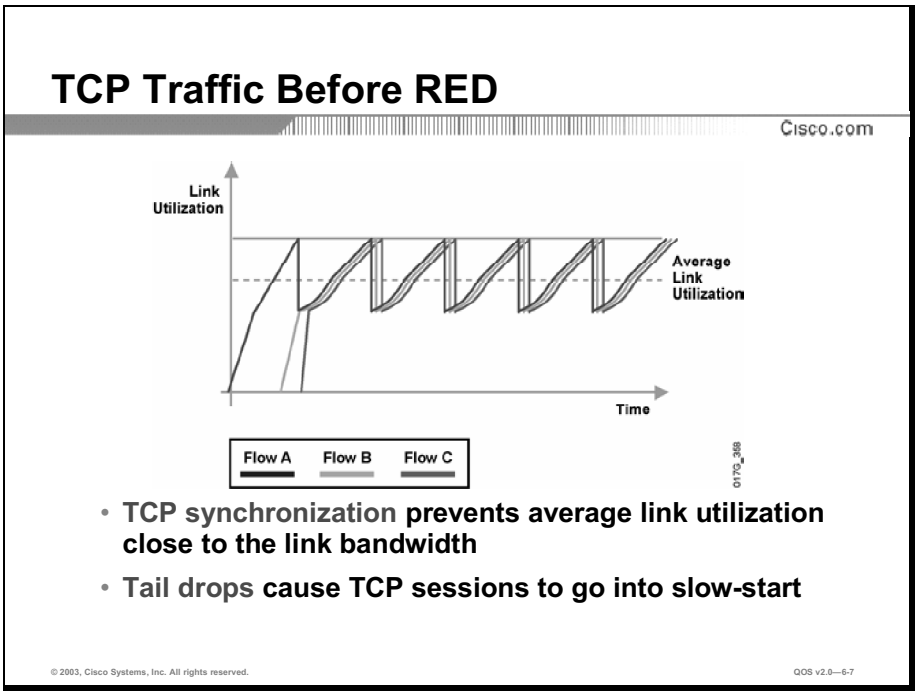

The figure shows TCP throughput behavior compared to link bandwidth in a congested network scenario where the tail-drop mechanism is in use on a router. The global synchronization phenomenon causes all sessions to slow down when congestion occurs, as all sessions are penalized when tail drop is used because it drops packets with no discrimination between individual flows.

When all sessions slow down, congestion on the router interface is removed and all TCP sessions restart their transmission at roughly the same time. Again, the router interface quickly becomes congested, causing tail drop. As a result, all TCP sessions back off again. This behavior cycles constantly, resulting in a link that is always underutilized on the average.

TCP Traffic After RED

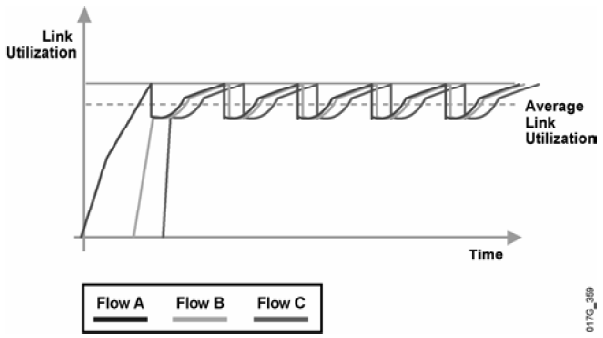

- **Average link utilization is much closer to link bandwidth**
- **Random drops cause TCP sessions to reduce window sizes**

The figure shows TCP throughput behavior compared to link bandwidth in a congested network scenario where RED has been configured on a router. RED randomly drops packets, influencing a small number of sessions at a time, before the interface reaches congestion. Overall throughput of sessions is increased, as well as average link utilization. Global synchronization is very unlikely to occur, as there is selective, but random dropping of adaptive traffic.

Applying Congestion Avoidance

This topic describes where congestion avoidance mechanisms are commonly deployed in enterprise and service provider networks.

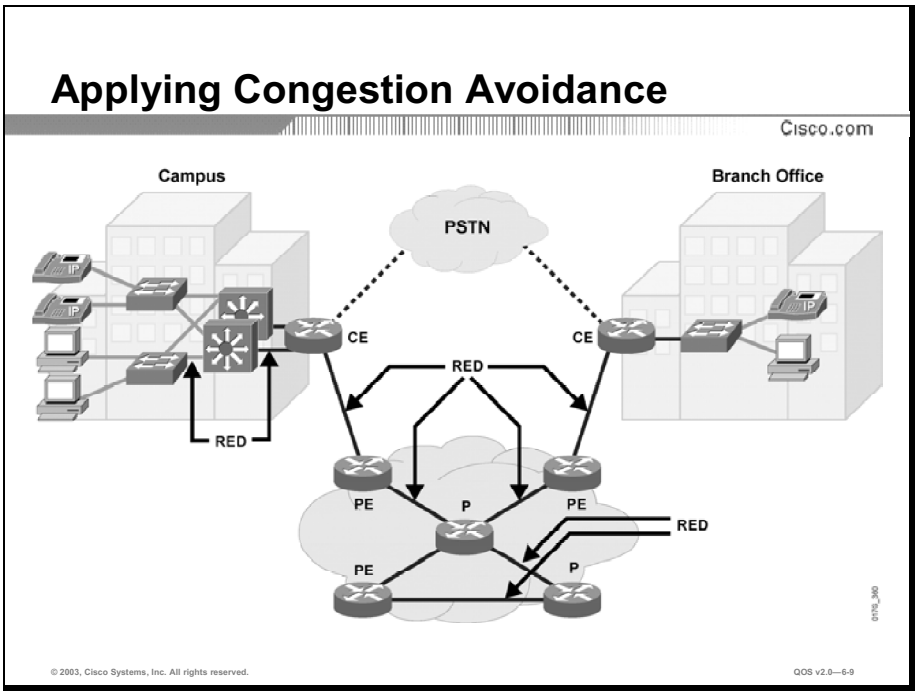

RED is most useful in enterprise and service provider networks on output interfaces where congestion is expected to occur. This typically relegates the use of RED to the core routers in a network rather than the routers at the network edge. Edge routers or switches typically classify and mark packets as they enter the network. Congestion avoidance mechanisms can use these packet markings to indicate a set of drop criteria for a traffic stream.

Congestion avoidance mechanisms are also applicable to the campus or LAN environment. In these networks, congestion avoidance is best used on interfaces that connect to WAN gateways, as these interfaces are typically sites for congestion to occur.

Summary

This topic summarizes the key points discussed in this lesson.

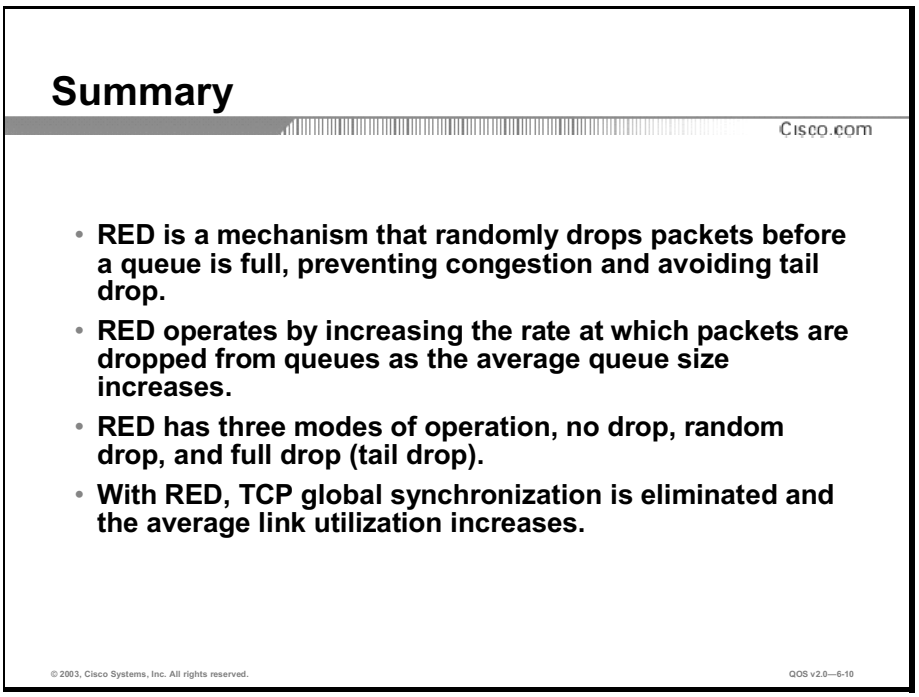

References

For additional information, refer to these resources:

- For a detailed discussion of TCP congestion behavior see Geoff Huston, Telstra, “The Future for TCP,” *Internet Protocol Journal*, Vol. 3, No. 3, September 2000, at the following URL: http://www.cisco.com/warp/public/759/ipj_3-3/ipj_3-3_futureTCP.html
- For more information on Random Early Detection, refer to, “Congestion Avoidance Overview” at the following URL: http://www.cisco.com/univercd/cc/td/doc/product/software/ios121/121cgcr/qos_c/qcprt3/qcdconav.htm

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) What are three ways in which RED avoids congestion? (Choose three.)

A) RED uses random dropping of packets to slow aggressive TCP flows.

B) RED provides a flow control mechanism directing senders to throttle the rate at which traffic is sent.

C) RED distributes losses over time and normally maintains a low queue depth.

D) RED increases the rate at which packets are dropped as congestion increases.

Q2) What are the three modes of RED? (Choose three.)

A) no drop

B) random drop

C) immediate drop

D) full drop

Q3) What is the purpose of the mark probability denominator?

A) It indicates at what average queue depth RED should begin tail drop.

B) It indicates at what average queue depth RED should begin random dropping.

C) It indicates the number of packets to drop when the average queue length is above the minimum threshold.

D) It is the fraction of packets dropped when the average queue depth is at the maximum threshold.

Q4) In a RED profile, what can result from setting the difference between the minimum and maximum threshold too small?

A) packets might never be dropped

B) global synchronization can occur

C) link utilization will be maximized

D) TCP congestive management will not operate properly

Q5) Where is RED typically applied in enterprise and service provider networks?

A) on the edge, to reduce the effects of congestion towards hosts and servers

B) at the distribution, on output interfaces pointing to network edge devices

C) on core devices where congestion is most likely to occur

D) on network gateways and interfaces connecting these devices to the network

Q6) What are two ways in which the application of RED helps to increase link utilization? (Choose two.)

A) It prioritizes traffic from fragile flows ahead of more aggressive flows.

B) RED eliminates global synchronization and its effects.

C) RED buffers packets in memory when interface queues are congested.

D) Random drops forces TCP session to reduce window sizes preventing congestion.

Q7) Match the terms below with their location on the figure.

A) Minimum Threshold

B) Maximum Threshold

C) No Drop

D) Random Drop

E) Full Drop

F) Mark probability denominator

G) Maximum Drop Probability

H) Average Queue Length

I) Drop Probability

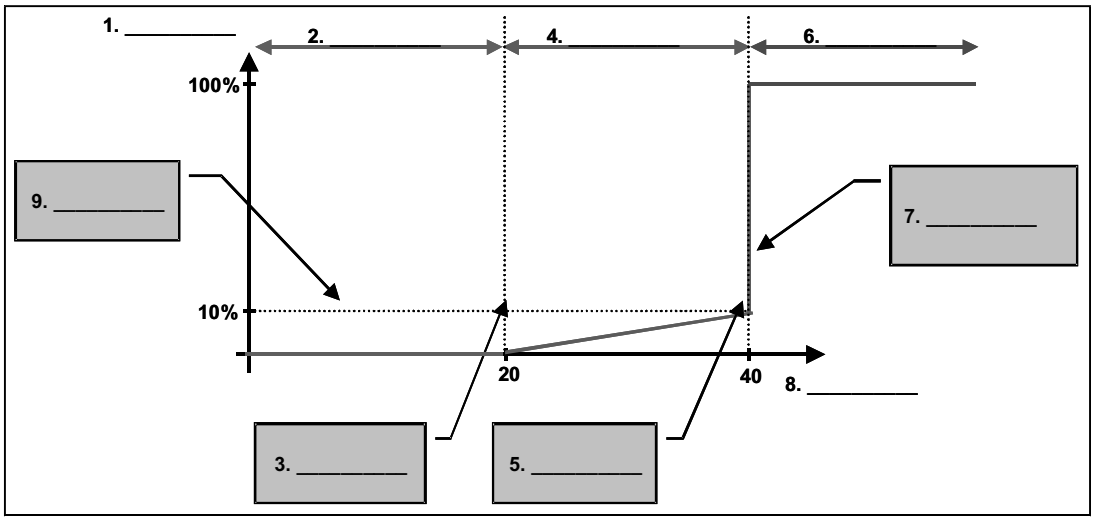

Quiz Answer Key

Q1) A, C, D

Relates to: Random Early Detection

Q2) A, B, D

Relates to: RED Modes

Q3) D

Relates to: RED Profiles

Q4) B

Relates to: RED Profiles

Q5) C

Relates to: Applying Congestion Avoidance

Q6) B, D

Relates to: TCP Traffic Before and After RED

Q7) 1-I

2-C

3-A

4-D

5-B

6-E

7-G

8-H

9-F

Relates to: RED Profiles

Configuring Class-Based Weighted RED

Overview

Congestion avoidance techniques monitor network traffic loads in an effort to anticipate and avoid congestion at common network bottleneck points using advanced packet discard techniques. This lesson introduces the congestion avoidance technique weighted random early detection (WRED), which is the Cisco implementation of RED.

Relevance

Congestion avoidance techniques offer a viable alternative to the default router congestion response, tail drop. WRED combines the capabilities of the RED algorithm with IP precedence for preferential traffic handling of higher-priority packets. When congestion begins on an interface, WRED can selectively discard lower-priority traffic, providing differentiated performance characteristics for different service classes.

Objectives

Upon completing this lesson, you will be able to configure CB-WRED to avoid congestion. This includes being able to meet these objectives:

- Describe WRED and how it can be used to prevent congestion
- Describe the traffic profiles used in WRED implementations
- Identify the Cisco IOS commands required to configure CB-WRED
- Identify the Cisco IOS commands required to configure DSCP-based CB-WRED
- Identify the Cisco IOS commands used to monitor CB-WRED

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Basic knowledge of internetworking with TCP/IP concepts
- Fundamentals of congestion avoidance with RED

Outline

The outline lists the topics included in this lesson.

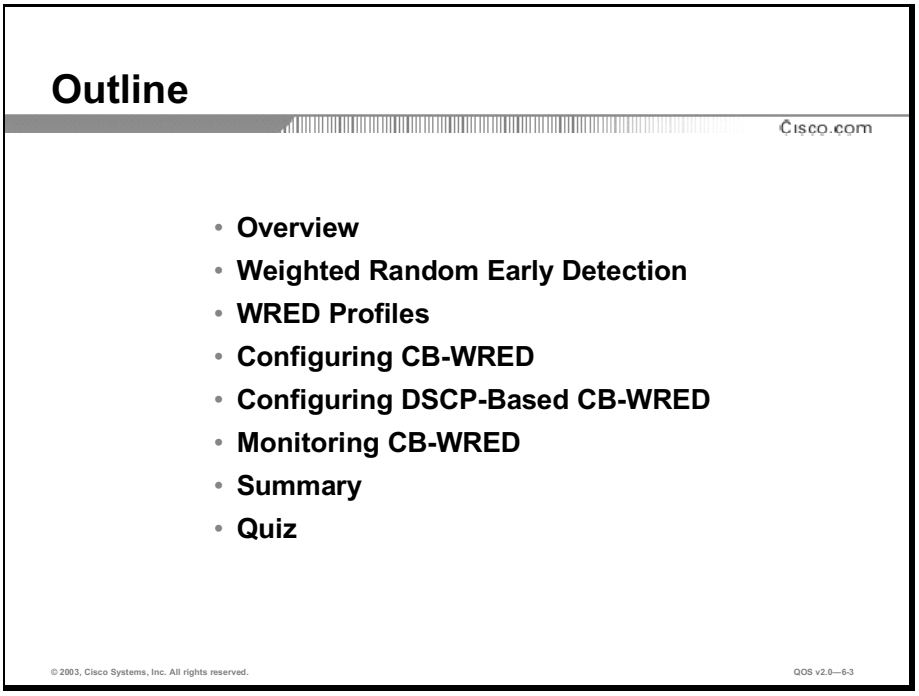

Weighted Random Early Detection

This topic describes WRED and how it can be used to prevent congestion.

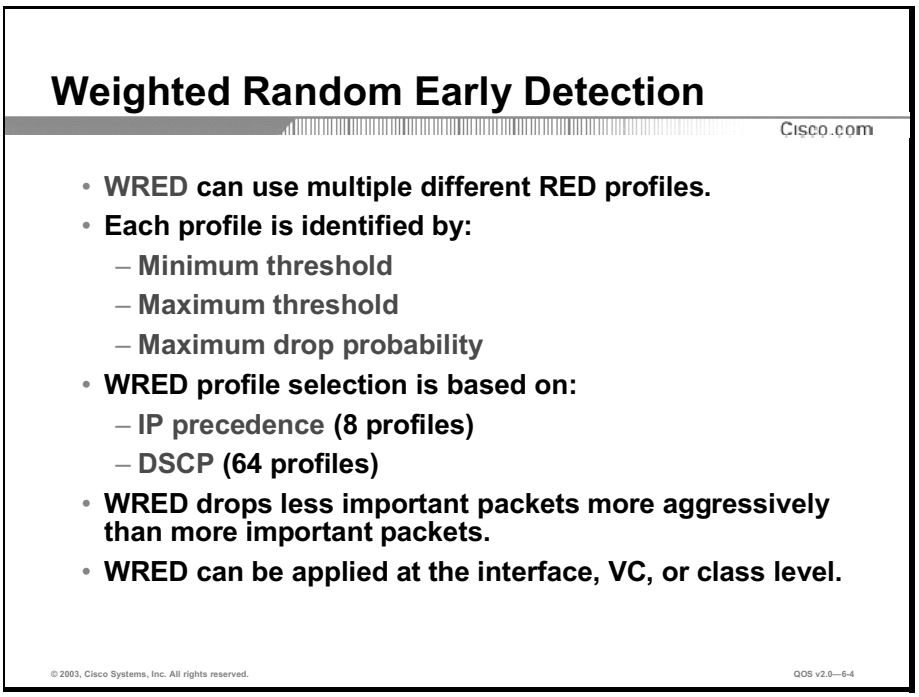

WRED combines RED with IP precedence or DSCP and performs packet dropping based on IP precedence or DSCP markings.

As with RED, WRED monitors the average queue length in the router and determines when to begin discarding packets based on the length of the interface queue. When the average queue length is greater than the user-specified "minimum threshold," WRED begins to randomly drop packets (both TCP and User Datagram Protocol [UDP]) with a certain probability. If the average length of the queue continues to increase such that it becomes larger than the user-specified "maximum threshold," WRED reverts to a "tail drop" packet discard strategy, where all incoming packets might be dropped.

The idea behind using WRED is to maintain the queue length at a level somewhere between the minimum and maximum thresholds, and to implement different drop policies for different classes of traffic. WRED can selectively discard lower-priority traffic when the interface becomes congested and can provide differentiated performance characteristics for different classes of service. WRED can also be configured so that non-weighted RED behavior is achieved.

For interfaces configured to use the Resource Reservation Protocol (RSVP), WRED chooses packets from other flows to drop rather than the RSVP flows. Also, IP precedence or DSCP governs which packets are dropped because traffic that is at a lower priority has a higher drop rate than traffic at a higher priority (and, therefore, lower priority is more likely to be throttled back). In addition, WRED statistically drops more packets from large users than small users. Therefore, traffic sources that generate the most traffic are more likely to be slowed down than traffic sources that generate little traffic.

WRED reduces the chances of tail drop by selectively dropping packets when the output interface begins to show signs of congestion. By dropping some packets early rather than waiting until the queue is full, WRED avoids dropping large numbers of packets at once and minimizes the chances of global synchronization. As a result, WRED maximizes the utilization of transmission lines.

WRED is only useful when the bulk of the traffic is TCP traffic. With TCP, dropped packets indicate congestion, so the packet source reduces its transmission rate. With other protocols, packet sources might not respond or might re-send dropped packets at the same rate, and so dropping packets might not decrease congestion.

WRED treats non-IP traffic as precedence 0, the lowest precedence. Therefore, non-IP traffic, in general, is more likely to be dropped than IP traffic.

WRED should be used wherever there is a potential bottleneck (congested link), which could very well be an access/edge link. However, WRED is normally used in the core routers of a network rather than at the network edge. Edge routers assign IP precedence or DSCP to packets as they enter the network. WRED uses these assigned values to determine how to treat different types of traffic.

Note that WRED is not recommended for any voice queue, although it may be enabled on an interface carrying voice traffic. WRED will not throttle back voice traffic because it is UDP-based. (The network itself should not be designed to lose voice packets because lost voice packets result in reduced voice quality.) WRED controls congestion by impacting other prioritized traffic and avoiding congestion helps to ensure voice quality.

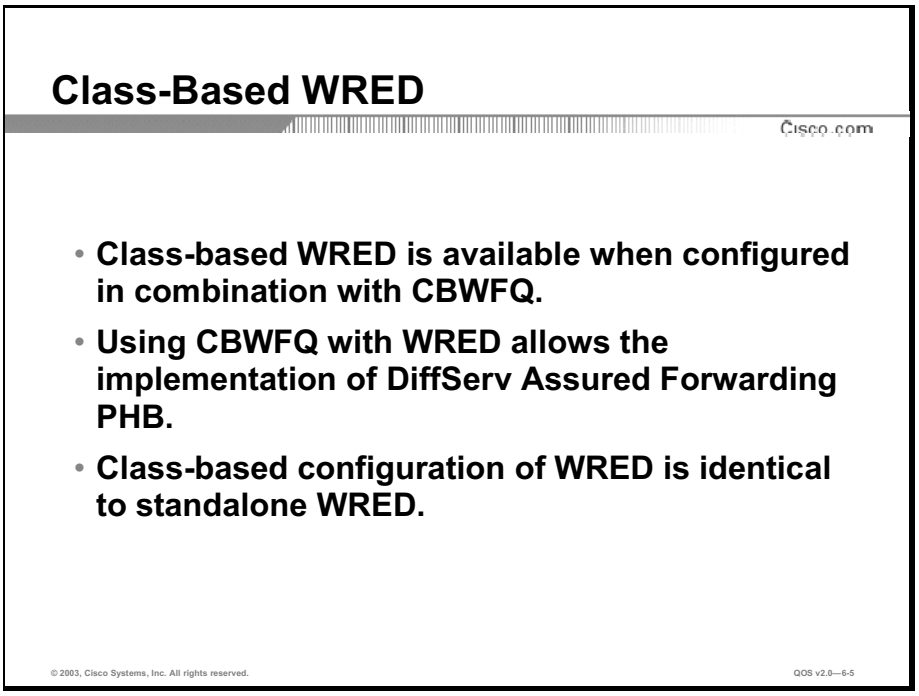

Congestion avoidance techniques monitor the network interface load in an effort to anticipate and avoid congestion at common network bottlenecks. Congestion avoidance is achieved through intelligent packet dropping techniques. Traditionally, Cisco IOS software used standalone RED and WRED mechanisms to avoid congestion on an interface. Those mechanisms can perform a differentiated drop based on the IP precedence or DSCP value.

The class-based weighted fair queuing (CBWFQ) system supports the use of WRED inside the queuing system, therefore implementing class-based weighted random early detection (CB-WRED). Each class is queued in its separate queue, and has a queue limit, performing tail drop by default. WRED can be configured as the preferred dropping method in a queue, implementing a differentiated drop based on traffic class and further on the IP precedence or DSCP value.

Note: The combination of CB-WFQ with WRED on a single device is currently the only way to implement the DiffServ Assured Forwarding per-hop behavior (AF PFB) using Cisco IOS software.

The class-based configuration of WRED is analogous to standalone WRED configuration.

Flow-based WRED is a variant of WRED that enforces more “fairness” in the way packets are dropped from different traffic flows. Flow-based WRED is not available within the CBWFQ queuing system and the Cisco IOS modular QoS command-line interface (MQC).

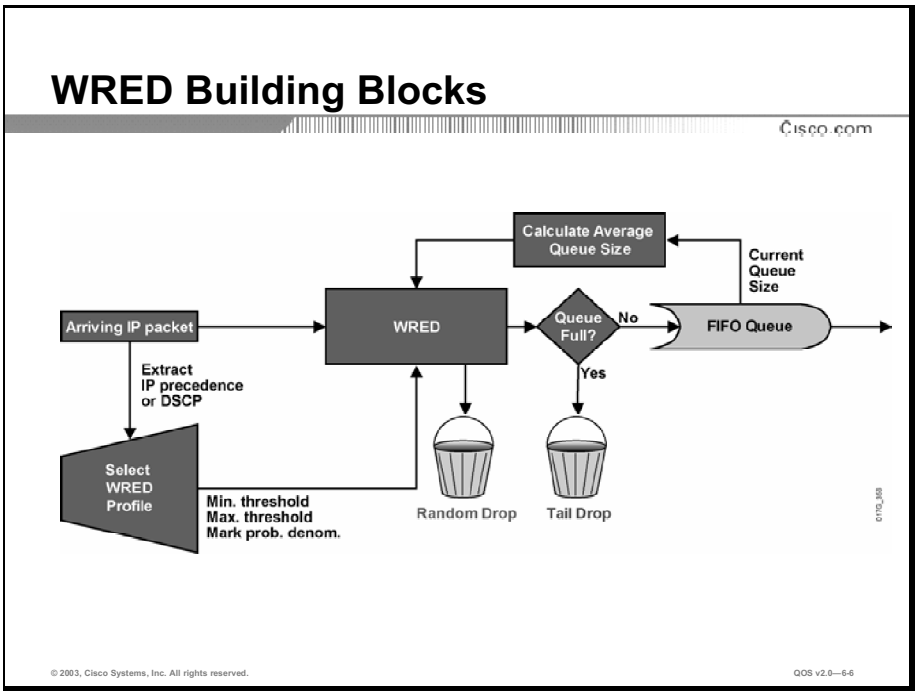

The figure shows how WRED is implemented, and the parameters that are used by WRED to influence packet drop decisions.

The router constantly updates the WRED algorithm with the calculated average queue length, which is based on the recent history of queue lengths.

Configured in the traffic profile are the parameters that define the drop characteristics used by WRED (minimum threshold, maximum threshold, and mark probability denominator). It is these parameters that define the WRED probability slopes.

When a packet arrives at the output queue, the IP precedence or DSCP value is used to select the correct WRED profile for the packet. The packet is then passed to WRED for processing. Based on the selected traffic profile and the average queue length, WRED calculates the probability for dropping the current packet and either drops it or passes it to the output queue.

If the queue is already full, the packet is tail-dropped. Otherwise, the packet will eventually be transmitted out onto the interface. If the average queue length is greater than the minimum threshold but less than the maximum threshold, based on the drop probability, WRED will either queue the packet or perform a random drop.

WRED Profiles

This topic describes the different traffic profiles that are used in WRED implementations.

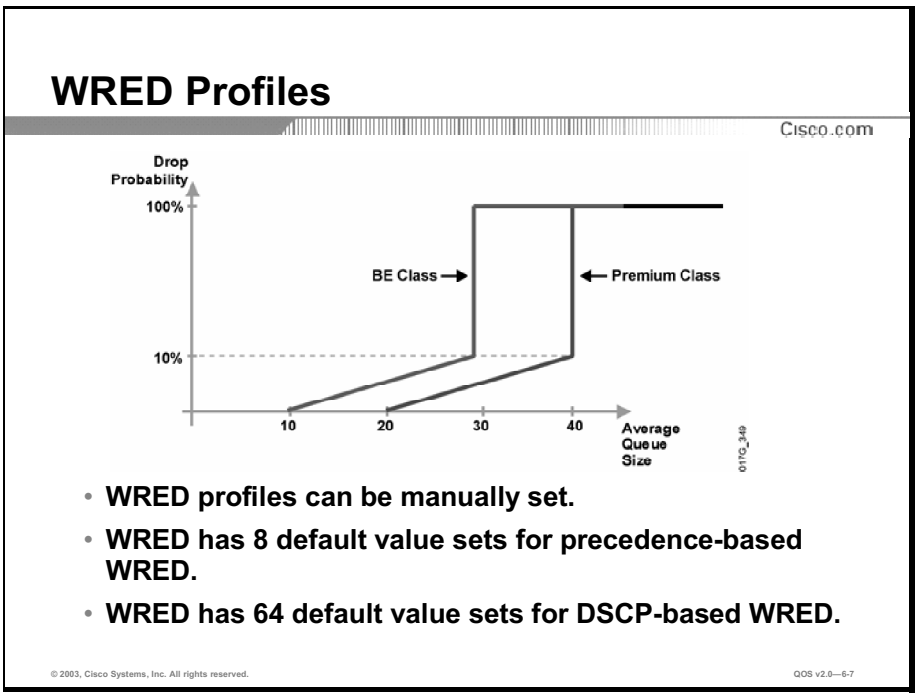

The figure shows two different WRED profiles that are used for traffic of two different QoS classes (“BE” class and “Premium” class).

The BE traffic class has a much lower minimum (10) and maximum threshold (30). As a result, traffic belonging to the BE class will be dropped much earlier and more aggressively than traffic from the Premium class. When heavy congestion occurs, traffic belonging to the BE class will ultimately be tail-dropped.

The Premium traffic class has been configured with higher minimum (20) and maximum thresholds (40). Therefore packet drop as a result of congestion will occur later (longer average queue size) and is less likely, as compared to the BE class. The differences in these traffic profiles, as defined in the figure, maintain differentiated levels of service in the event of congestion.

To avoid the need of setting all WRED parameters in a router, 8 default values are already defined for precedence-based WRED, and 64 DiffServ aligned values are defined for DSCP-based WRED. Therefore, the default settings should suffice in the vast majority of deployments.

By default, the maximum threshold for all DSCP values is 40. The default mark probability denominator for all DSCP values is 10.

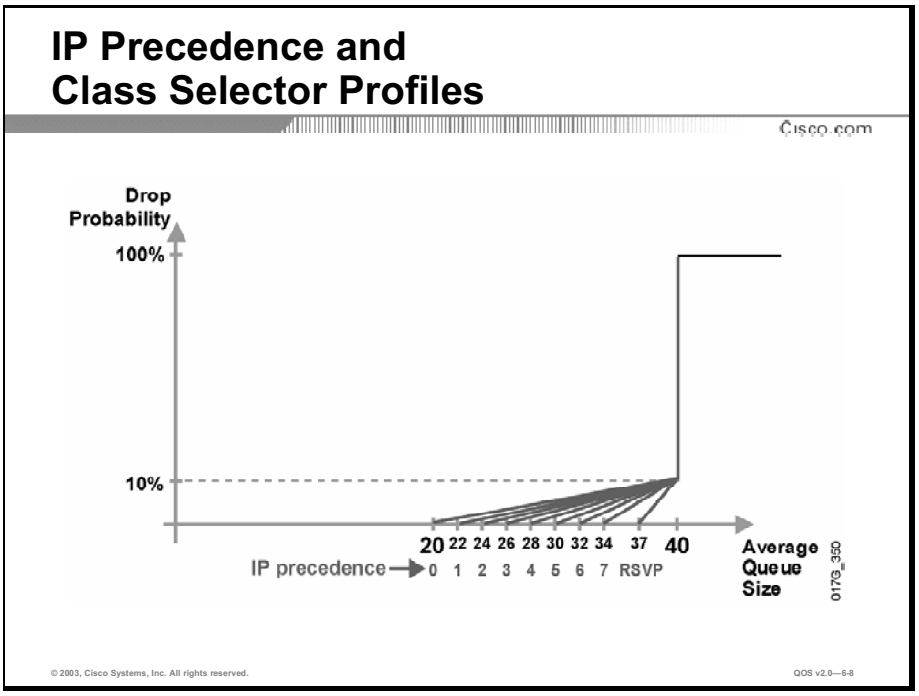

A PHB is the externally observable forwarding behavior applied at a DiffServ-compliant node to a DiffServ Behavior Aggregate (BA). With the ability of the system to mark packets according to DSCP setting, collections of packets (each with the same DSCP setting and sent in a particular direction) can be grouped into a DiffServ BA. Packets from multiple sources or applications can belong to the same DiffServ BA.

The class selector BA is used for backward compatibility with non-DiffServ-compliant devices (RFC 1812 compliant devices and, optionally, RFC 791 compliant devices). Therefore, the class selector range of DSCP values is used for backward compatibility with IP precedence. The same WRED profiles are applied to equal IP precedence and class selector values:

IP Precedence and Class Selector Profiles

IP Precedence	DSCP (Class Selector)	Default Minimum Threshold
0 (000)	Default (0)	20
1 (001)	cs1 (8) (001000)	22
2 (010)	cs2 (16) (010000)	24
3 (011)	cs3 (24) (011000)	26
4 (100)	cs4 (32) (100000)	28
5 (101)	cs5 (40) (101000)	30
6 (110)	cs6 (48) (110000)	32
7 (111)	cs7 (56) (111000)	34
RSVP	RSVP	37

DSCP-Based WRED (Expedited Forwarding)

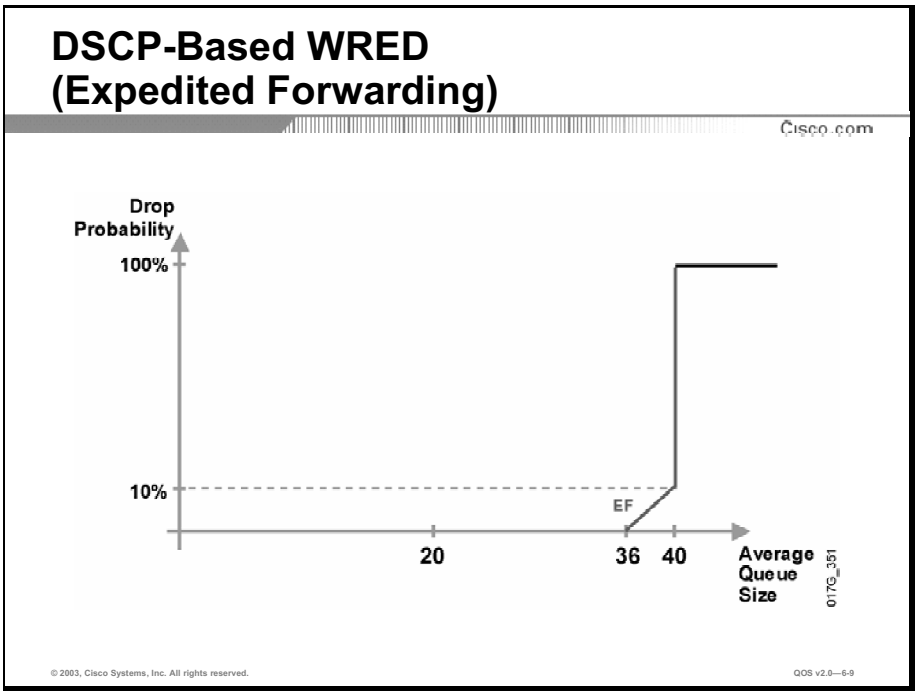

In DSCP, the Expedited Forwarding (EF) PHB is identified based on the following parameters:

- Ensures a minimum departure rate to provide the lowest possible delay to delay-sensitive applications
- Guarantees bandwidth to prevent starvation of the application if there are multiple applications using EF PHB
- Polices bandwidth to prevent starvation of other applications or classes that are not using this PHB
- Packets requiring EF should be marked with DSCP binary value “101110” (46 or 0x2E)

For the EF DiffServ traffic class, WRED configures itself by default so that the minimum threshold is very high, thus increasing the probability of no drops being applied to that traffic class. It is expected then, that EF traffic should be dropped very late, as compared to other traffic classes, and is therefore prioritized in the event of congestion.

Expedited Forwarding Profile

DSCP (Six Bits)	Default Minimum Threshold
EF (101110)	36

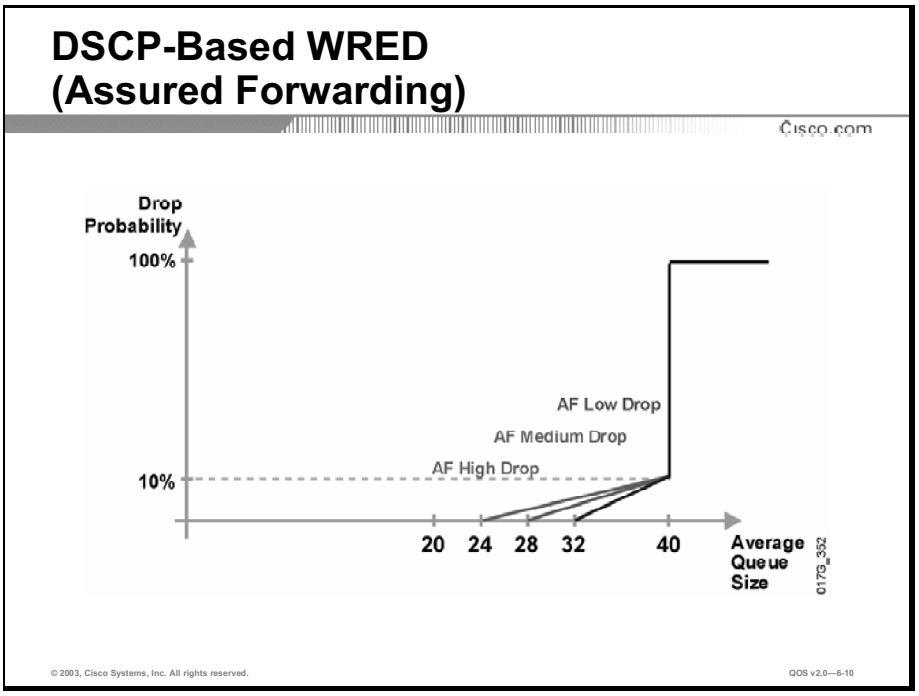

In DSCP, the Assured Forwarding (AF) PHB is identified based on the following parameters:

- Guarantees a certain amount of bandwidth to an AF class
- Allows access to extra bandwidth, if available
- Packets requiring AF PHB should be marked with DSCP value “aaadd0” where “aaa” is the number of the class and “dd” is the drop probability or drop preference of the traffic class.

There are four standard-defined AF classes. Each class should be treated independently and have bandwidth allocated that is based on the QoS policy. For the AF DiffServ traffic class, WRED configures itself by default for three different profiles, depending on the drop preference DSCP marking bits. Therefore, AF traffic should be classified into the three possible classes based on the sensitivity of the application or applications represented by the class to packet drops.

Assured Forwarding Profiles

Assured Forwarding Class	Drop Probability	(AF Class) DSCP	Default Minimum Threshold
AF class 1	Low Drop Prob	(AF11) 001010	32
	Medium Drop Prob	(AF12) 001100	28
	High Drop Prob	(AF13) 001110	24
AF class 2	Low Drop Prob	(AF21) 010010	32
	Medium Drop Prob	(AF22) 010100	28
	High Drop Prob	(AF23) 010110	24
Assured Forwarding Class	**Drop Probability**	**(AF Class) DSCP**	**Default Minimum Threshold**
AF class 3	Low Drop Prob	(AF31) 011010	32
	Medium Drop Prob	(AF32) 011100	28
	High Drop Prob	(AF33) 011110	24
AF class 4	Low Drop Prob	(AF41) 100010	32
	Medium Drop Prob	(AF42) 100100	28
	High Drop Prob	(AF43) 100110	24

Configuring CB-WRED

This topic describes the Cisco IOS commands that are required to configure CB-WRED.

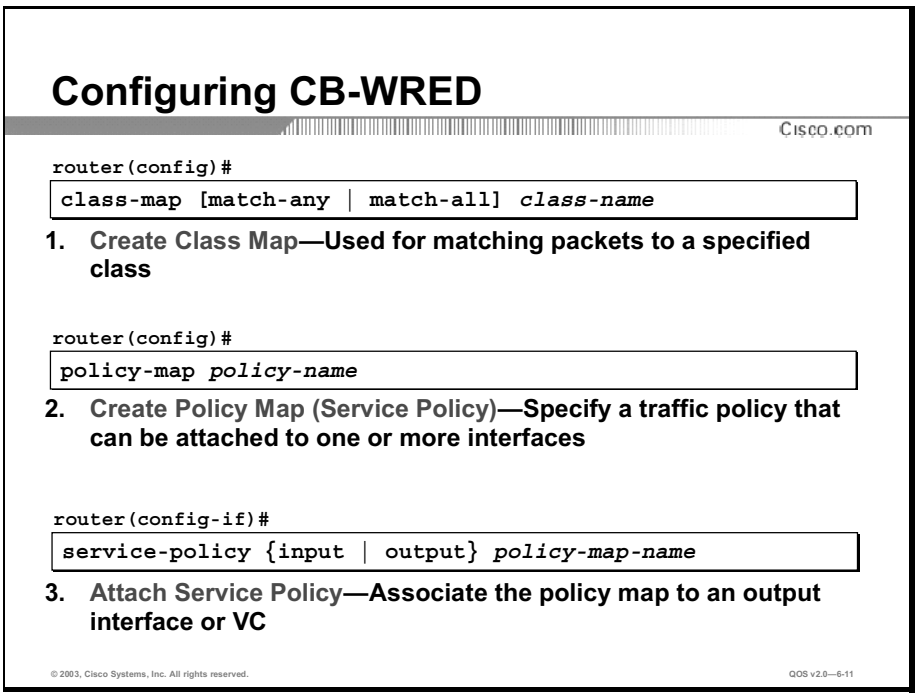

To configure CB-WRED (WRED at the class level with CB-WFQ), the dscp-based and prec-based arguments are configured within MQC. Specific CB-WRED configuration arguments are applied within a policy map. The policy map configuration can then be applied wherever policy maps are attached (for example, at the interface level, the per-virtual circuit [VC] level, or the shaper level).

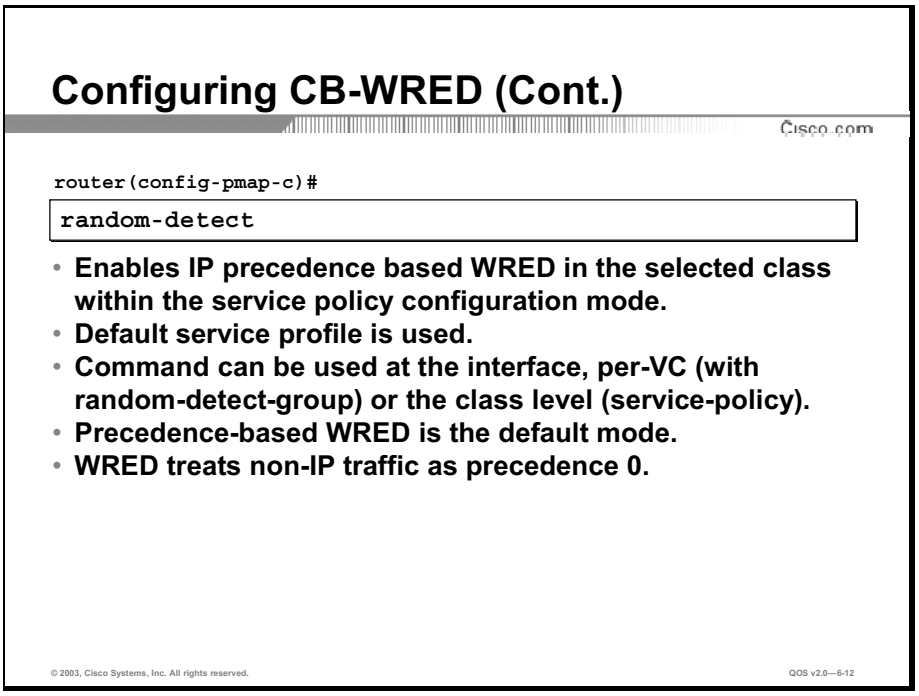

The **random-detect** command is used to enable WRED on an interface. By default, WRED is precedence-based and uses eight default WRED profiles, one for each value of IP precedence.

Within the CBWFQ system, WRED is used to perform per-queue dropping within the class queues. Therefore, each class queue has its own WRED method, which can be further weighed based on the IP precedence or DSCP value. Each queue can therefore be configured with a separate dropping policy to implement different drop policies for every class of traffic.

WRED will treat all non-IP traffic as precedence 0. As a result, non-IP traffic is more likely to be dropped than IP traffic.

If the random-detect command is used on virtual IP (VIP)-based interfaces, distributed WRED (DWRED) is enabled and the VIP CPU is responsible for WRED dropping. This can significantly increase router performance when used in the context of distributed Cisco Express Forwarding (CEF) switching, which is a prerequisite for DWRED functionality. Also, DWRED can be combined with distributed weighted fair queuing (DWFQ), enabling truly distributed queuing and congestion avoidance techniques, running independently from the central CPU.

WRED cannot be configured on the same interface as custom queuing (CQ), priority queuing (PQ), or WFQ. However, both DWRED and DWFQ can be configured on the same interface. In addition, CB-WRED can be configured in conjunction with CBWFQ. Restricting non-distributed, non-class-based WRED to only FIFO queuing on an interface is typically not a major issue because WRED is usually applied in the network core, where advanced queuing mechanisms are not typically used. WRED is suited for the network core as it has a relatively low performance impact on routers. Further, DWRED or CB-WRED can be used to overcome this limitation by combining WRED with WFQ.

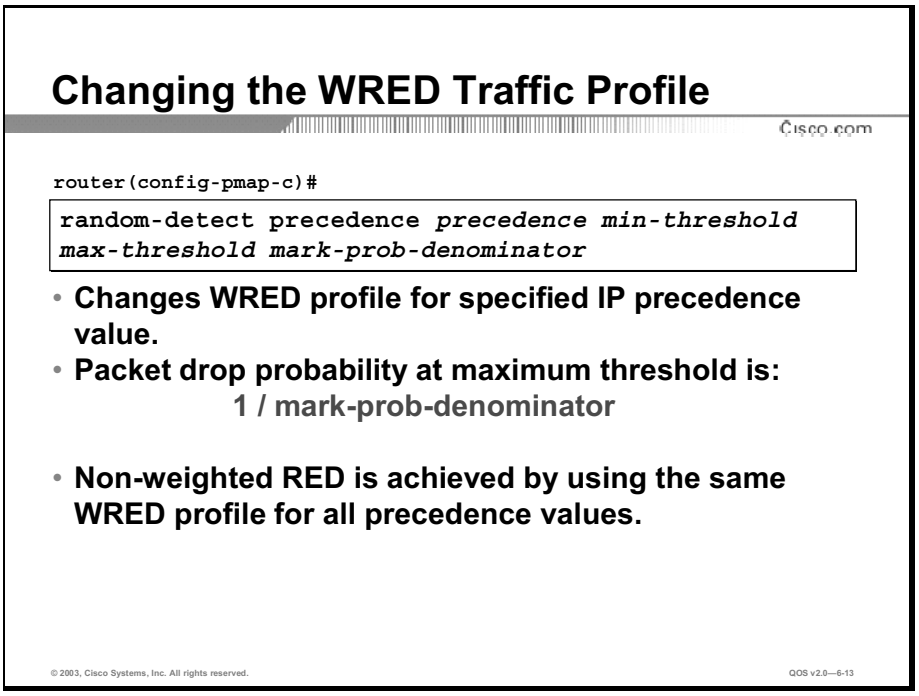

When WRED is enabled, default values are selected for each traffic profile based on the weight used (IP precedence or DSCP). Network administrators can then modify these default values to match their specific administrative QoS policy goals. When modifying the default WRED profile for IP precedence, the following values are configurable:

- **Minimum threshold:** When the average queue depth is above the minimum threshold, WRED 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.

Note: The default maximum threshold is equal to the default hold queue size (40) on an interface. The size of the hold queue is equivalent to the number of packets that can be held within a queue. The hold queue length ranges from 0 to 4096, and therefore, the minimum/maximum threshold range is 1 to 4096. The default maximum threshold will reflect the defined hold queue size. Thus, if the hold queue is changed, the maximum threshold will change.

- **Maximum threshold:** When the average queue size is above the maximum threshold, all packets are dropped.

Note: If the difference between the maximum threshold and the minimum threshold is too small, many packets might be dropped at once, resulting in global synchronization.

- **Mark probability denominator:** This is the fraction of packets dropped when the average queue depth is at the maximum threshold. For example, if the denominator is 10, one out of every 10 packets is dropped when the average queue is at the maximum threshold.

Note: The maximum probability of drop at the maximum threshold can be expressed as 1/mark-prob-denominator. The maximum drop probability is 10 percent if default settings are used that have a mark probability denominator value of 10. The value of the mark probability can range from 1 to 65536.

If required, RED can be configured as a special case of WRED, by assigning the same profile to all eight IP precedence values.

Note: The default WRED parameter values are based on the best available data. Cisco recommends that these parameters should not be changed from their default values unless you have determined that your applications will benefit from the changed values.

Changing WRED Sensitivity to Bursts

Cisco.com

```
router(config-pmap-c)#
random-detect exponential-weighting-constant n
```

- **WRED takes the average queue size to determine the current WRED mode (no drop, random drop, full drop)**

$$Q_{avg}(t+1) = Q_{avg}(t) \cdot (1 - 2^{-n}) + Q_t \cdot 2^{-n}$$

New average queue size | Previous average queue size | Current queue size

0179_355

- **High values of N allow short bursts**
- **Low values of N make WRED more burst-sensitive**
- **Default value (9) should be used in most scenarios**
- **Average output queue size with N=9 is**

$$Q_{ave}(t+1) = Q_{ave}(t) * 0.998 + Q_t * 0.002$$

WRED does not calculate the drop probability using the current queue length, but rather uses the average queue length. The average queue length is constantly recalculated using two terms: the previously calculated average queue size and the current queue size. An exponential weighting constant N influences the calculation by weighing the two terms, therefore influencing how the average queue size follows the current queue size, in the following way:

- For high values of N, the previous average becomes more important. A large factor will smooth out the peaks and lows in queue length. The average queue size is unlikely to change very quickly, avoiding drastic swings in size. The WRED process will be slow to start dropping packets, but it may continue dropping packets for a time after the actual queue size has fallen below the minimum threshold. The slow-moving average will accommodate temporary bursts in traffic.

Note: If the value of N gets too high, WRED will not react to congestion. Packets will be transmitted or dropped as if WRED were not in effect.

- For low values of N, the average queue size closely tracks the current queue size. The resulting average may fluctuate with changes in the traffic levels. In this case, the WRED process responds quickly to long queues. When the queue falls below the minimum threshold, the process will stop dropping packets.

Note: If the value of N gets too low, WRED will overreact to temporary traffic bursts and drop traffic unnecessarily.

The default value of N is 9. This value should suffice for most scenarios except perhaps those involving extremely high-speed interfaces (like OC12), where it can be increased slightly (to about 12) to allow more bursts.

Example: CBWFQ Using IP Precedence with CB-WRED

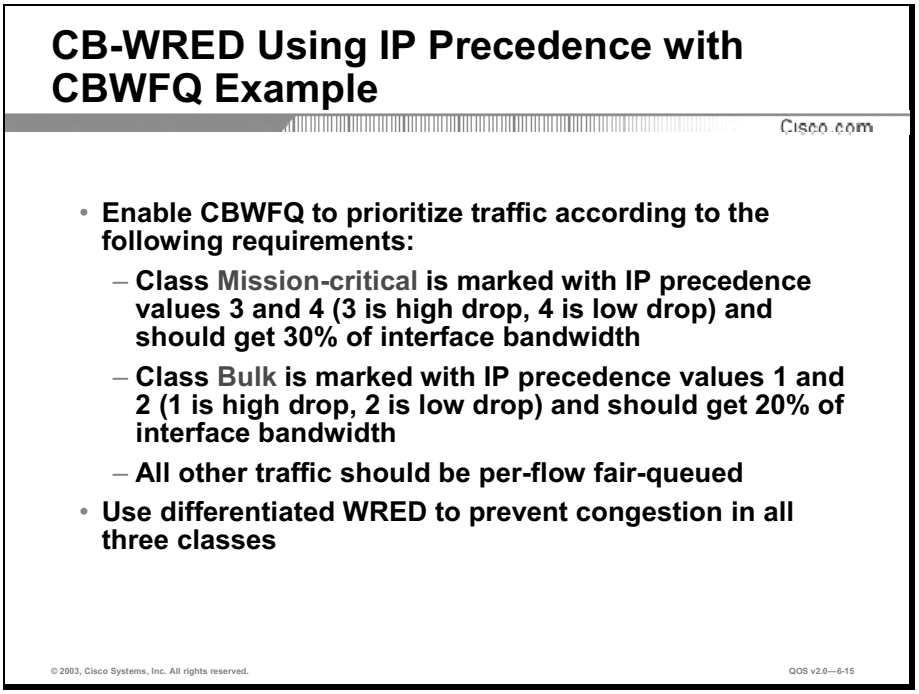

The following example of CBWFQ with WRED focuses on a network, which provides three different service levels for three traffic classes:

- **Mission-critical class:** Marked with IP precedence values of 3 and 4 (3 is used for high drop, and 4 is used for low drop within the service class) should get 30 percent of an interface bandwidth
- **Bulk class:** Marked with IP precedence values of 1 and 2 (1 being high-drop, and 2 being low-drop service) should get 20 percent of the interface bandwidth
- **Best-effort class:** Should get the remaining bandwidth share, and should be fair-queued

To enforce this service policy, a router will use CBWFQ to perform bandwidth sharing and WRED within service classes to perform differentiated drop.

CB-WRED Using IP Precedence with CBWFQ Example (Cont.)

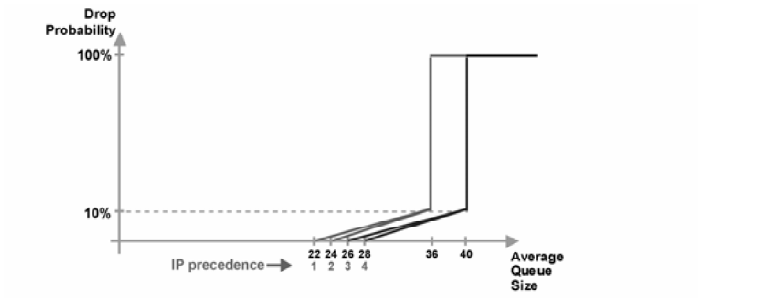


```
class-map Mission-critical
 match ip precedence 3 4
!
class-map Bulk
 match ip precedence 1 2
!
policy-map Policy1
   class Mission-critical
    bandwidth percent 30
    random-detect
    random-detect precedence 3 26 40 10
    random-detect precedence 4 28 40 10
```

```
class Bulk
    bandwidth percent 20
    random-detect
    random-detect precedence 1 22 36 10
    random-detect precedence 2 24 36 10
  class class-default
   fair-queue
   random-detect
!
```


The figure shows the WRED traffic profile representing the QoS service policy and the configuration that is used to implement the example service policy. The traffic is classified based on the precedence bits, and all non-contract traffic is classified into the default class.

- The Mission-critical class is guaranteed at least 30 percent of bandwidth with a custom WRED profile that establishes a low-drop and a high-drop per-hop behavior.
- The bulk class is guaranteed at least 20 percent of bandwidth, is configured with somewhat lower WRED drop thresholds, and is therefore more likely to be dropped than the Mission-critical class in the event of interface congestion.
- All other traffic is part of the default class and is fair-queued with default WRED parameters.

Configuring DSCP-Based CB-WRED

This topic describes the Cisco IOS commands that are required to configure DSCP-based CB-WRED.

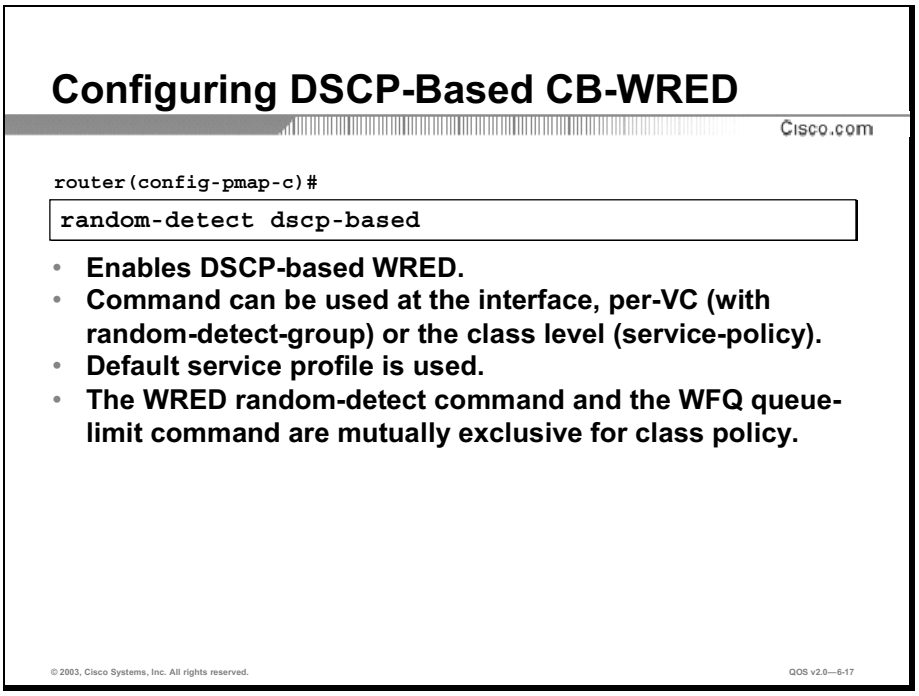

The **random-detect dscp-based** command is used to enable DSCP-based WRED on an interface. By default, WRED is precedence-based, and uses eight default WRED profiles, one for each value of IP precedence. Changing WRED weighting to values based on DSCP increases the number of WRED traffic profiles to 64.

You can configure WRED as part of the policy for a standard class or the default class. The WRED **random-detect** command and the WFQ **queue-limit** command are mutually exclusive for class policy. If you configure WRED, its packet drop capability is used to manage the queue when packets exceeding the configured maximum count are enqueued. If you configure the WFQ **queue-limit** command for class policy, tail drop is used.

WRED cannot be configured on the same interface as CQ, PQ, or WFQ. However, both DWRED and DWFQ can be configured on the same interface. In addition, CB-WRED can be configured in conjunction with CBWFQ. Restricting non-distributed, non-class-based WRED only to FIFO queuing on an interface is not a major issue because WRED is usually applied in the network core, where advanced queuing mechanisms are not typically deployed. WRED is suited for the network core as it has a relatively low performance impact on routers. Further, DWRED or CB-WRED can be used to overcome this limitation by combining WRED with WFQ.

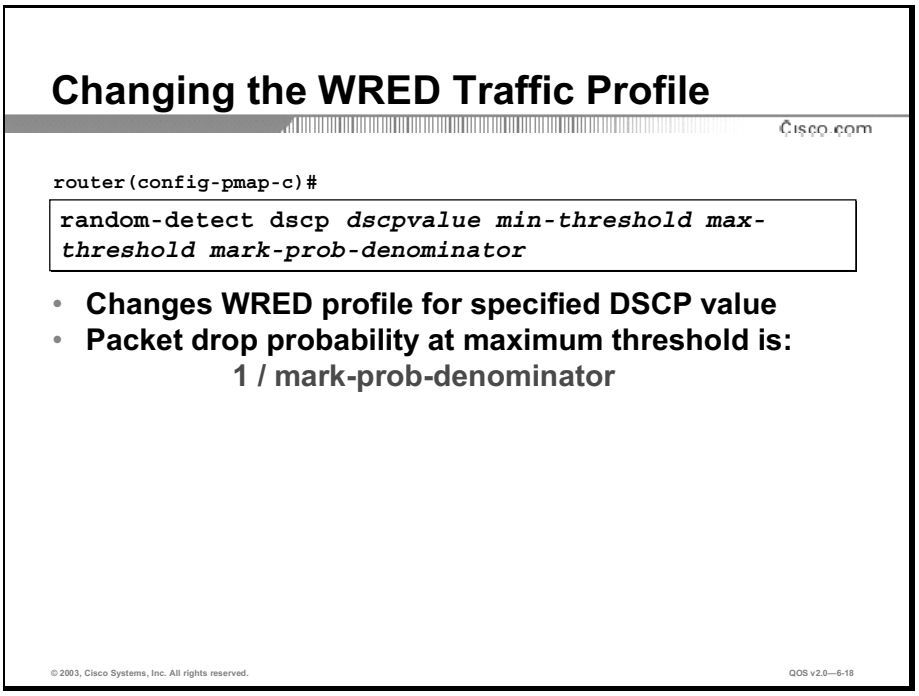

When DSCP-based WRED is enabled, default values are selected for each traffic profile based on DSCP. Network administrators can then modify these default values to match their specific administrative QoS policy goals. When modifying the default WRED profile for DSCP, the following values are configurable:

- **Minimum threshold:** When the average queue depth is above the minimum threshold, WRED 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.

Note: The default maximum threshold is equal to the default hold queue size (40) on an interface. The size of the hold queue is equivalent to the number of packets that can be held within a queue. The hold queue length ranges from 0 to 4096, and therefore, the minimum/maximum threshold range is 1 to 4096. The default maximum threshold will reflect the defined hold queue size. Thus, if the hold queue is changed, the maximum threshold will change.

- **Maximum threshold:** When the average queue size is above the maximum threshold, all packets are dropped.

Note: If the difference between the maximum threshold and the minimum threshold is too small, many packets might be dropped at once, resulting in global synchronization.

- **Mark probability denominator:** This is the fraction of packets dropped when the average queue depth is at the maximum threshold. For example, if the denominator is 10, one out of every 10 packets is dropped when the average queue is at the maximum threshold.

Note: The maximum probability of drop at the maximum threshold can be expressed as 1/mark-prob-denominator. The maximum drop probability is 10 percent if default settings are used that have a mark probability denominator value of 10. The value of the mark probability can range from 1 to 65536.

Note: The default WRED parameter values are based on the best available data. Cisco recommends that these parameters should not be changed from their default values.

Example: CB-WRED Using DSCP with CBWFQ

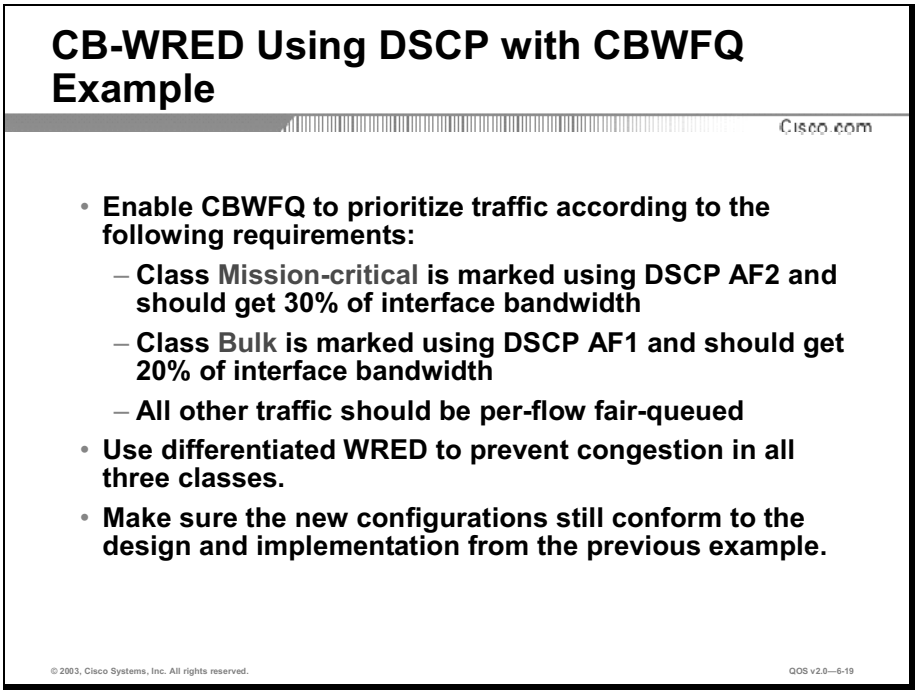

In the previous example of CBWFQ with WRED, the service policy was implemented using classes of service based on IP precedence. In this example, the same service policy will be configured. In this case, however, DSCP-based classes of service are used. Remember that the DiffServ model itself provides defined traffic classes and their associated PHB. DiffServ-based classification is used in this example:

- **Mission-critical class:** Marked using DSCP AF class 2 should get 30 percent of an interface bandwidth
- **Bulk class:** Marked using DSCP AF class 1 should get 20 percent of the interface bandwidth
- **Best-effort class:** Traffic should get the remaining bandwidth share, and should be fair-queued

To enforce this service policy, a router will use CBWFQ to perform bandwidth sharing, and WRED within service classes to perform differentiated drop.

CB-WRED Using DSCP with CBWFQ Example (Cont.)

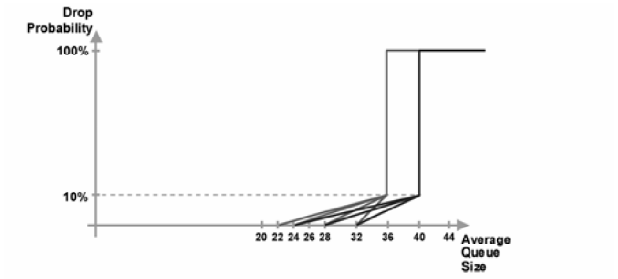


```
class-map Mission-critical
 match ip dscp af21 af22 af23 cs2
!
class-map Bulk
 match ip dscp af11 af12 af13 cs1
!
policy-map Policy1
   class Mission-critical
    bandwidth percent 30
    random-detect dscp-based
    random-detect dscp af21 32 40 10
    random-detect dscp af22 28 40 10
    random-detect dscp af23 24 40 10
    random-detect dscp cs2 24 40 10
```

```
class Bulk
   bandwidth percent 20
   random-detect dscp-based
   random-detect dscp af11 32 36 10
   random-detect dscp af12 28 36 10
   random-detect dscp af13 24 36 10
   random-detect dscp cs1 22 36 10
class class-default
   fair-queue
   random-detect dscp-based
!
```


The configuration example shows how traffic classification is performed using DSCP-based classes, representing the Mission-critical class as the AF1 class, and using the AF2 class as the Bulk class. WRED DSCP-based parameters are sent reflecting the class-dependent drop strategy.

- The Mission-critical class is guaranteed at least 30 percent of bandwidth, with a custom WRED profile, which establishes three different drop probabilities for AF class 2.
- The Bulk class is guaranteed at least 20 percent of bandwidth, is configured with three different drop probabilities for AF class 1, and has a somewhat lower WRED maximum threshold. As a result, Bulk class traffic is more likely to be dropped than the Mission-critical class in the event of interface congestion.

All other traffic is part of the default class, is fair-queued, with default WRED parameters.

Monitoring CB-WRED

This topic describes the Cisco IOS commands that are required to monitor CB-WRED.

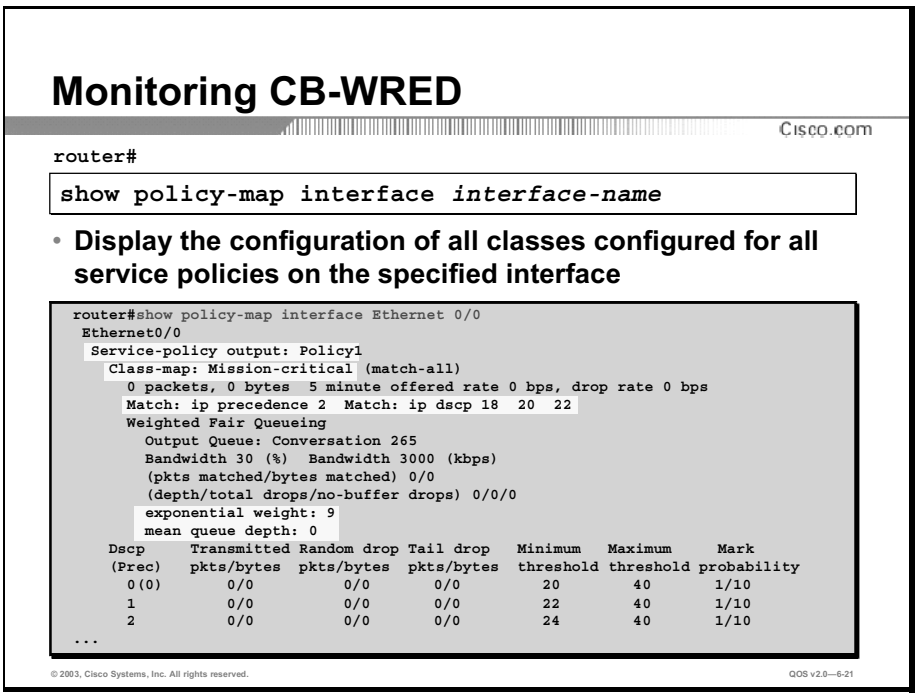

The **show policy-map interface** command displays the configuration of all classes configured for all service policies on the specified interface. This includes all WRED parameters implementing the dropping policy on the specified interface.

The following table explains some of the key fields of the **show policy-map interface** command.

show policy-map interface Parameters

Parameter	Description
`Service-policy output`	Name of the output service policy applied to the specified interface or VC.
`Class-map`	Class of traffic being displayed. Output is displayed for each configured class in the policy. The choice for implementing class matches (for example, match-all or match-any) can also appear next to the traffic class.
`Match`	Match criteria specified for the class of traffic. Choices include criteria such as IP precedence, IP DSCP value, Multiprotocol Label Switching (MPLS) experimental value, access groups, and QoS groups.
`exponential weight`	Exponent used in the average queue size calculation for a WRED parameter group.
`mean queue depth`	Average queue depth based on the actual queue depth on the interface and the exponential weighting constant. It is a fluctuating average. The minimum and maximum thresholds are compared against this value to determine drop decisions.

Summary

This topic summarizes the key points discussed in this lesson.

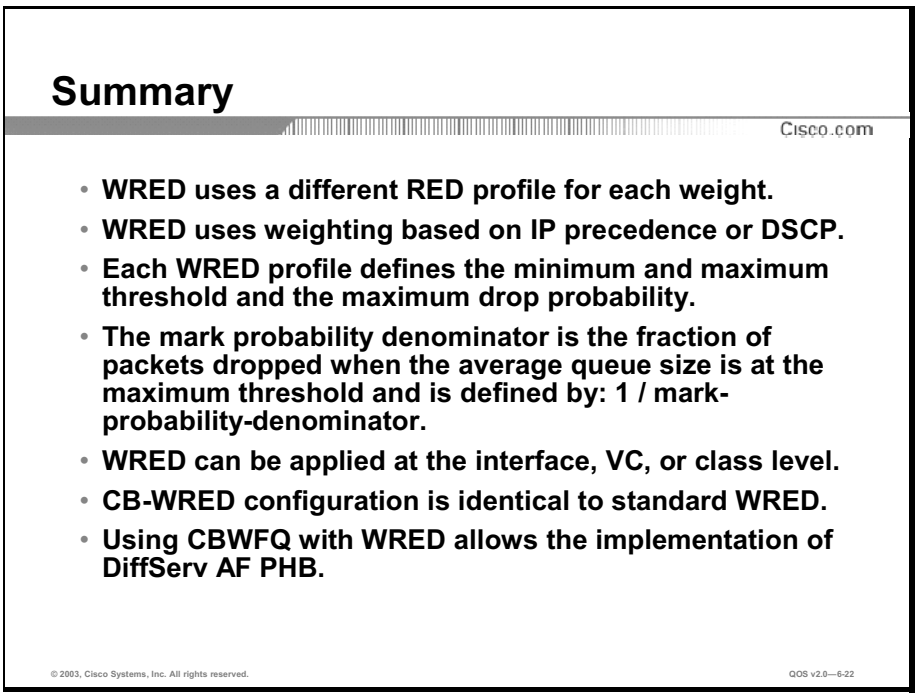

References

For additional information, refer to these resources:

- For more information on WRED and configuring WRED, refer to, “Configuring Weighted Random Early Detection” at the following URL: http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/12cgcr/qos_c/qcpart3/qcwred.htm
- For more information on DSCP-Based WRED, refer to, “DiffServ Compliant Weighted Random Early Detection” at the following URL: http://www.cisco.com/univercd/cc/td/doc/product/software/ios121/121newft/121t/121t5/dtdswred.htm
- For information regarding WRED on Cisco GSR 12000 routers, refer to, “Weighted Random Early Detection on the Cisco 12000 Series Router” at the following URL: http://www.cisco.com/univercd/cc/td/doc/product/software/ios112/ios112p/gsr/wred_gs.htm

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) Which two profiles are supported by WRED? (Choose two.)

A) one traffic profile for all traffic

B) eight profiles for IP precedence-based WRED

C) 64 profiles for DSCP-based WRED

D) up to 256 user-defined profiles

Q2) What two factors is the “weight” based on in WRED? (Choose two.)

A) a user-definable metric ranging from 0 to 32768

B) IP precedence

C) DiffServ Code Point

D) QoS-Group

Q3) A maximum drop probability of 20 percent is required for a WRED profile. Which value should be configured as the mark probability denominator?

A) 2

B) 10

C) 20

D) depends upon the length of the queue limit

Q4) What are two ramifications of modifying the exponential weighting constant in WRED? (Choose two)

A) Setting the constant too low can cause WRED to drop traffic unnecessarily.

B) Setting the constant too low can cause WRED to not drop traffic until the max threshold has been reached.

C) Setting the constant too high can make WRED unresponsive to congestion.

D) Setting the constant too high can cause WRED to begin dropping high levels of traffic before it should.

Q5) Which of the following commands provides service level for IP precedence 3 with min-threshold of 20, max-threshold of 40 and mark-prob-denominator value for 5 percent drop probability?

A) **random-detect precedence 3 40 20 5**

B) **random-detect precedence 3 20 40 5**

C) **random-detect precedence 3 20 40 20**

D) **random-detect precedence 5 2 4 5**

Q6) What three actions does the, **(config-pmap-c)# random-detect dscp 8 24 40** command complete? (Choose three.)

A) enables DSCP-based WRED

B) configures the minimum threshold to 8

C) configures the minimum threshold to 24

D) configures the maximum threshold to 24

E) configures the maximum threshold to 40

F) configures the mark probability denominator to 40

G) configures the mark probability to 10

Q7) Match the DSCP-based WRED profile with the DSCP PHB it implements as shown in the figures below.

A) AF

B) EF

C) CS

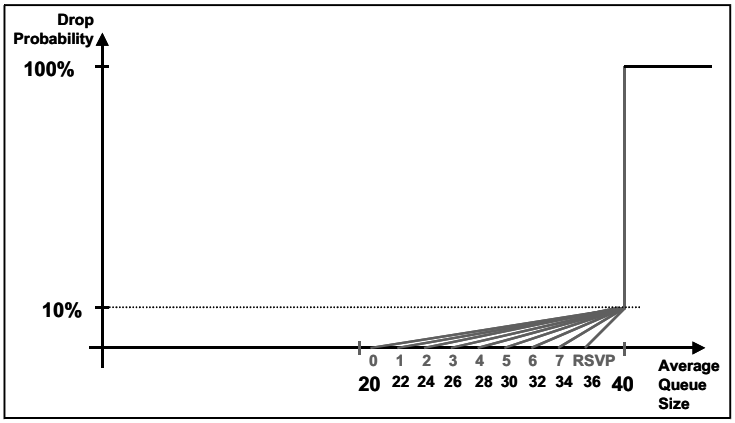

_____ 1.

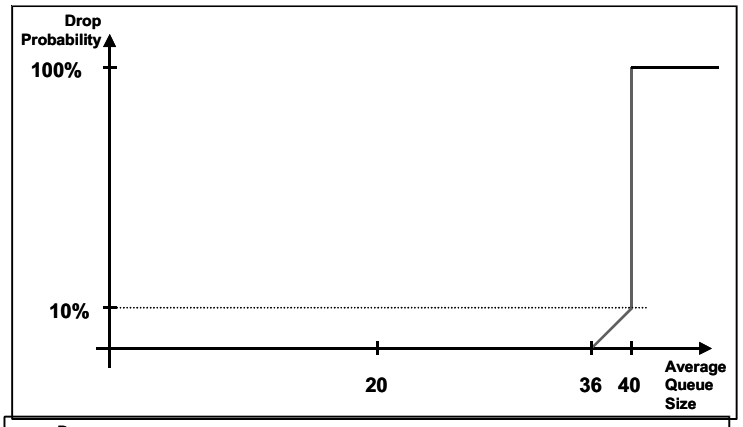

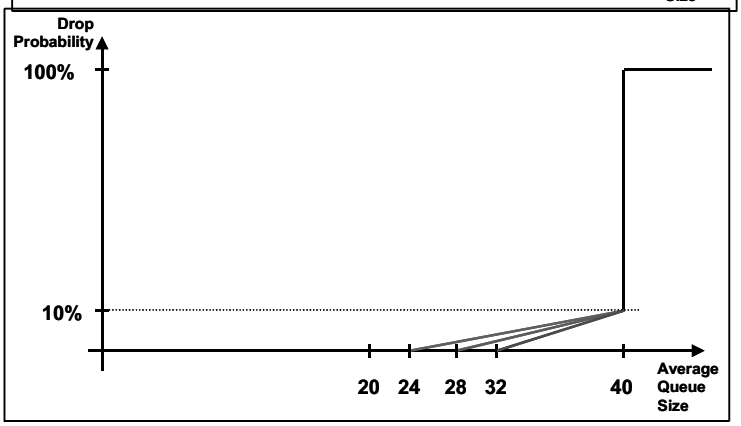

_____ 2.

_____ 3.

Q8) Given the following Cisco router console output, fill in the blanks for each question.

```
router# show policy-map interface serial 4/0

   Serial4/0
    Service-policy output: AVVID (2022)
      Class-map: silver (match-all) (2023/2)
       251162 packets, 375236028 bytes
       1 minute offered rate 612000 bps, drop rate 0 bps
       Match: ip dscp 18  20  22  (2025)
      Weighted Fair Queueing
        Output Queue: Conversation 265
          Bandwidth 25 (%)
          (pkts matched/bytes matched) 3/4482
          (depth/total drops/no-buffer drops) 0/0/0
          mean queue depth: 0

     Dscp      Random drop    Tail drop     Minimum   Maximum
Mark
     (Prec)    pkts/bytes    pkts/bytes   threshold threshold
probability
       18          0/0           0/0           20        40
1/20
       20          0/0           0/0           20        40
1/15
       22          0/0           0/0           20        40
1/10
```

A) What QoS mechanism has been configured? _____

B) What traffic types belong to the traffic class shown? _____

C) What DSCP PHB is being used in this service class? _____

D) If the interface becomes congested, at what average queue length will the interface resort to tail drop? _____

E) At the time the average queue length reaches the maximum threshold, what percentage of traffic will be dropped by WRED for all traffic types in this traffic class? _____, _____, _____

Quiz Answer Key

Q1) B, C

Relates to: Weighted Random Early Detection

Q2) B, C

Relates to: Weighted Random Early Detection

Q3) C

Relates to: WRED Profiles

Q4) A, C

Relates to: Configuring CB-WRED

Q5) B

Relates to: Configuring CB-WRED

Q6) C, E, G

Relates to: Configuring DSCP-Based CB-WRED

Q7) 1-C

2-B

3-A

Relates to: WRED Profiles

Q8) A) CB-WRED with CB-WFQ

B) DSCP 18, DSCP 20, DSCP 22

C) Assured Forwarding (AF)

D) 40

E) DSCP 18 - 5%, DSCP 20 - 6.67%, DSCP 22 - 10%

Relates to: Monitoring CB-WRED

Case Study: WRED Traffic Profiles

Overview

This case study activity provides information regarding the QoS administrative policy requirements of a small to mid-sized network. Your task is to work with a partner to evaluate the QoS requirements, and based on these requirements, create WRED traffic profiles that you can use to implement the required QoS administrative policy. You will discuss your traffic profile with the instructor and other classmates, and the instructor will present a solution for the case study to the class.

Relevance

The creation of traffic profiles is an important step in correctly implementing an active queue management strategy using congestion avoidance mechanisms such as WRED.

Objectives

In this activity, you will create the appropriate WRED traffic profile to properly implement a customer QoS administrative policy. Upon completing this case study, you will be able to meet these objectives:

- Review customer QoS requirements
- Identify the service classes required to implement the policy
- Create WRED traffic profiles that can be used to implement the policy
- Present a solution to the case study

Learner Skills and Knowledge

To benefit fully from this activity, you must have these prerequisite skills and knowledge:

- Basic knowledge of internetworking with TCP/IP concepts
- Fundamentals of congestion avoidance with WRED

Outline

The outline lists the topics included in this activity.

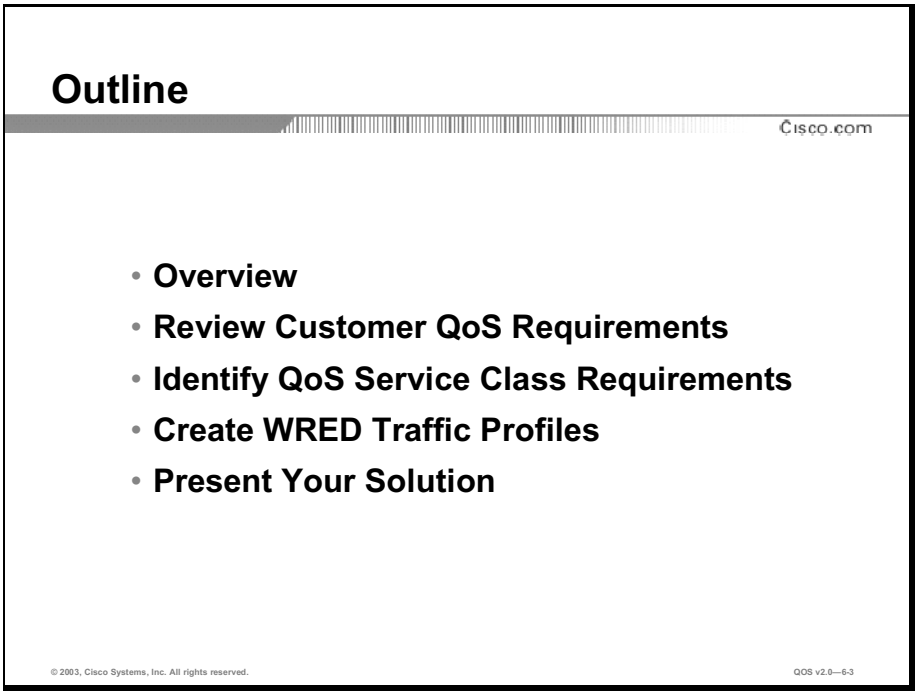

Required Resources

These are the resources required to complete this exercise:

- Case Study Activity: WRED Traffic Profiles
- A workgroup consisting of two learners

Job Aids

No job aids are required to complete this case study.

Case Study Tasks

The activity includes these tasks:

Step 1 Review customer QoS requirements. Completely read the customer requirements provided.

Step 2 Identify QoS service class requirements. With the aid of your partner, identify the service classes required to implement the administrative QoS policy based on customer requirements.

Step 3 Create WRED traffic profiles. Create the WRED traffic profiles required to properly implement the administrative QoS policy.

Step 4 Present your solution. After the instructor presents a solution to the case study, present your solution to the class with your partner.

Case Study Verification

You have completed this activity when your case study solution has been presented to the class and you have justified any major deviations from the case study solution supplied by the instructor.

Review Customer QoS Requirements

This case study involves analyzing an administrative QoS policy of LCR Incorporated, a fictitious manufacturer of recumbent bicycles. The company has provided you with a short description of their requirements. It is your task to provide the network engineers from LCR with a QoS solution to meet their requirements.

Read the customer requirements and discuss them with your partner. Identify the different classes of service required and hence the number of WRED traffic profiles required to solve this customer problem.

Company Background

LCR Incorporated began making recumbent bicycles in the garage of its owner Patrick Cagney, in 1984. Since that time, the company has grown to be a global provider of recumbent bicycles. Headquartered in St. Petersburg, Florida, LCR has two manufacturing facilities and five sales offices in the United States.

Each site utilizes dedicated 100 Mb switching to the desktop and contains a distributed server farm. Each site connects over a private WAN connection to the corporate headquarters using an IP-enabled Frame Relay service from a global service provider. WAN link speeds are all T1 (1.544 Mbps).

Customer Situation

LCR Incorporated is currently experiencing application performance problems and has an urgent need to resolve them. Internet usage at LCR is extremely high because most of the sales and customer contacts of the company use the Internet. The company currently has redundant, 3-Mbps Internet connections at its headquarters. Much of the use of the Internet, however, is for non-business-critical applications. Therefore, Internet browsing and non-critical applications should be treated as the lowest priority.

Many of the applications at LCR are distributed between sites because they require collaboration between members of the LCR staff. Examples are Oracle and Citrix. Manufacturing and Finance use Oracle databases to manage inventory, shipping, order entry, and customer billing. These systems are integrated across the company and reside in the main data center at the headquarter location. Citrix is heavily used for quality assurance monitoring of manufacturing and its automated systems. LCR has indicated that the Oracle application and Citrix transactions are critical to the company. Internet traffic should not be allowed to interfere with Oracle or Citrix transactions.

Working with the network engineering staff at LCR and the service provider, you have been enlisted to assist LCR by defining QoS requirements for their network. Their first priority is to deploy active congestion management mechanisms across the provider backbone to ease the congestion issues they are experiencing.

Identify QoS Service Class Requirements

Identify the different traffic classes required to implement the customer administrative QoS policy. Use the table below to help you with your answer choices. Write your answers on the lines below:

Customer Traffic: _____ PHB: _____ DSCP: _____

Customer Traffic: _____ PHB: _____ DSCP: _____

Customer Traffic: _____ PHB: _____ DSCP: _____

QoS Service Classes

PHB	DSCP	DSCP Value	Intended Protocols and Applications	Service Class	Service Class and Configuration
EF	EF	101110	Interactive Voice	Voice Bearer	Admission Control = RSVP Queuing = Priority
AF1	AF11 AF12 AF13	001010 001100 001110	General Data Service, FTP, Backups	Bulk Data	Queuing = Rate Based Active Queue Mgt = WRED minth AF13 < maxth AF13 <= minth AF12 < maxth AF12 <= minth AF11 < maxth AF11
AF2	AF21 AF22 AF23	010010 010100 010110	Database access, transaction services, interactive traffic, preferred data service	Transactional	Queuing = Rate Based Active Queue Mgt = WRED minth AF23 < maxth AF23 <= minth AF22 < maxth AF22 <= minth AF21 < maxth AF21
AF3	AF31 AF32 AF33	011010 011100 011110	Locally defined mission-critical applications	Mission-critical	Queuing = Rate Based Active Queue Mgt = WRED minth AF33 < maxth AF33 <= minth AF32 < maxth AF32 <= minth AF31 < maxth AF31
AF4	AF41 AF42 AF43	100010 100100 100110	Interactive video and associated voice	Interactive Video	Admission Control = RSVP Queuing = Rate Based Active Queue Mgt = WRED minth AF43 < maxth AF43 <= minth AF42 < maxth AF42 <= minth AF41 < maxth AF41
CS6	Class 6	110000	Border Gateway Control (BGP), Open Shortest Path First (OSPF), etc.	Routing (Reserved)	Queuing = Rate Based Small guaranteed minimum rate Active Queue Mgt = RED minth < maxth, but minth is deep to minimize loss
CS4	Class 4	100000	Often proprietary	Streaming Video	Admission Control = RSVP Queuing = Rate Based Active Queue Mgt = RED minth < maxth
CS3	Class 3	011000	Session Initiation Protocol (SIP), H.323, etc.	Call Signaling	Queuing = Rate Based Small guaranteed minimum rate Active Queue Mgt = RED minth < maxth, but minth is deep to minimize loss
CS1	Class 1	001000	User-selected service, PPP Applications	Less-than-Best-Effort Data (Scavenger)	Queuing = Rate Based No bandwidth guarantee Active Queue Mgt = RED minth < maxth
Default	Default (Best Effort) Class 0	000000	Unspecified traffic, E-mail, Internet	Best-Effort	Queuing = Rate Based Minimal bandwidth guarantee Active Queue Mgt or Per-flow fair queuing Active Queue Mgt = RED minth < maxth

Create WRED Traffic Profiles

Create a WRED traffic profile for each of the service classes identified in the previous section. Use the following table to assist you in creating your profile. When completing each profile, be sure to draw the traffic profile and include all information on the blank profile graphic provided.

Cisco IOS Default WRED Profile Values

PHB	Minimum Threshold	Maximum Threshold	Mark Probability
af11	32	40	1/10
af12	28	40	1/10
af13	24	40	1/10
af21	32	40	1/10
af22	28	40	1/10
af23	24	40	1/10
af31	32	40	1/10
af32	28	40	1/10
af33	24	40	1/10
af41	32	40	1/10
af42	28	40	1/10
af43	24	40	1/10
cs1	22	40	1/10
cs2	24	40	1/10
cs3	26	40	1/10
cs4	28	40	1/10
cs5	30	40	1/10
cs6	32	40	1/10
cs7	34	40	1/10
EF	36	40	1/10
RSVP	36	40	1/10
Default (BE)	20	40	1/10

Traffic Profile 1:

Traffic Class: _____ PHB: _____

WRED Traffic Profile Parameters:

Minimum Threshold: _____ Maximum Threshold: _____

Mark Probability Denominator: _____

Traffic Profile 2:

Traffic Class: _____ PHB: _____

WRED Traffic Profile Parameters:

Minimum Threshold: _____ Maximum Threshold: _____

Mark Probability Denominator: _____

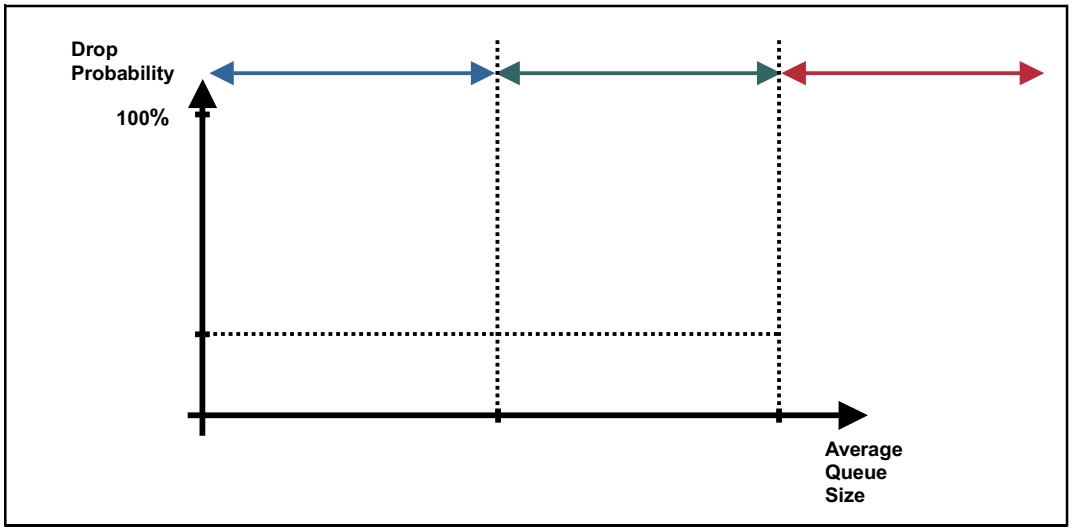

Traffic Profile 3:

Traffic Class: _____ PHB: _____

WRED Traffic Profile Parameters:

Minimum Threshold: _____ Maximum Threshold: _____

Mark Probability Denominator: _____

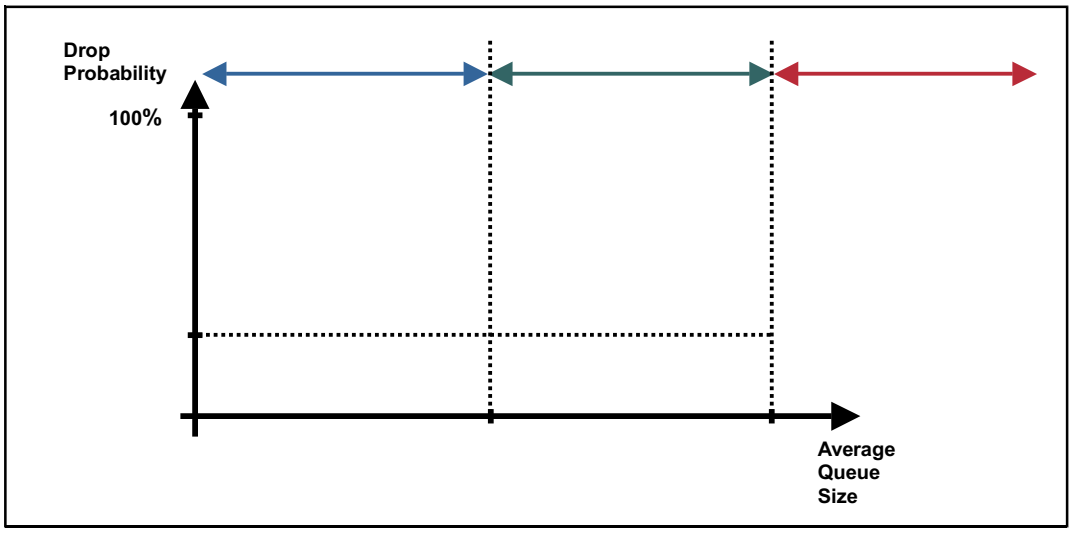

Present Your Solution

Together with your partner, present your solution to the class. Include the following information:

- Customer service class requirements
- WRED traffic profiles
- Justification for differences from the solution presented by the instructor

Case Study Answer Key

Identify Customer QoS Requirements

Customer Traffic: Oracle	PHB: AF2	DSCP: AF21
Customer Traffic: Citrix	PHB: AF2	DSCP: AF22
Customer Traffic: Internet	PHB: Default (BE)	DSCP: 0

Create WRED Traffic Profiles

Traffic Profile 1:

Traffic Class: Transactional PHB: AF21

Minimum Threshold: 32 Maximum Threshold: 40 Mark Probability Denominator: 10

Traffic Profile 2:

Traffic Class: Transactional PHB: AF22

Minimum Threshold: 28 Maximum Threshold: 40 Mark Probability Denominator: 10

Traffic Profile 3:

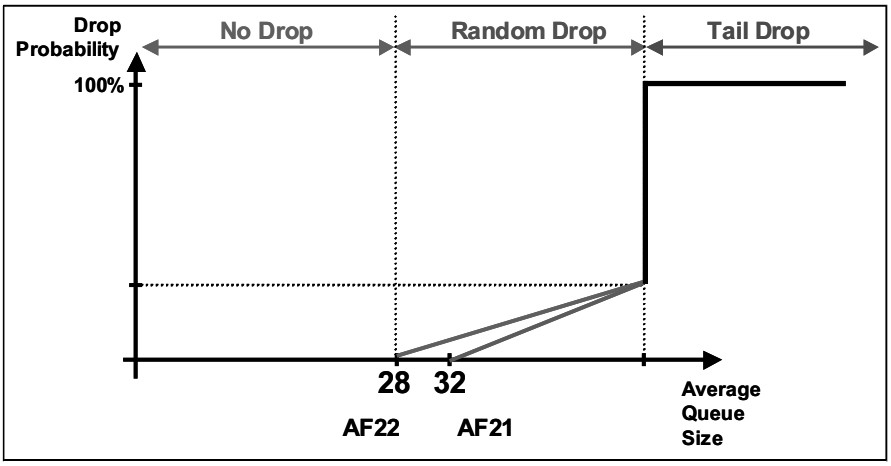

Traffic Class: Default PHB: 0

Minimum Threshold: 20 Maximum Threshold: 40 Mark Probability Denominator: 10

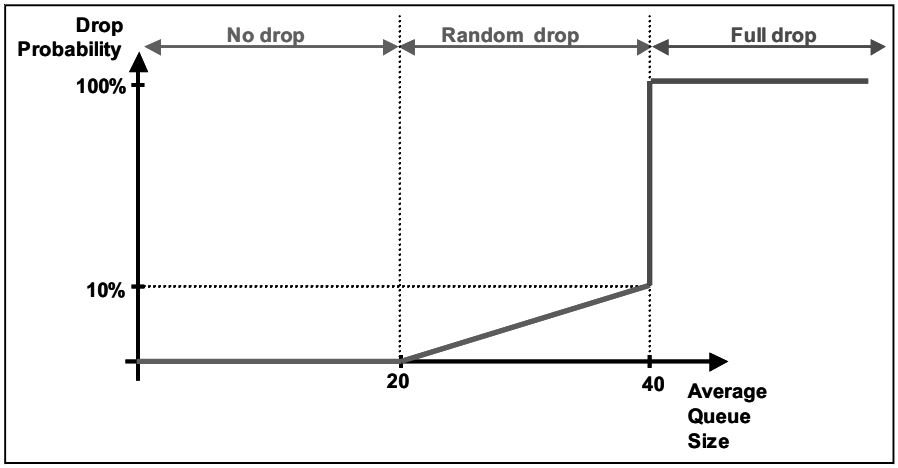

Configuring Explicit Congestion Notification

Overview

Congestion avoidance techniques monitor network traffic loads in an effort to anticipate and avoid congestion at common network bottleneck points. Congestion avoidance is achieved through packet dropping by using more complex techniques than simple tail drop. With the addition of ECN extensions to IP, routers now have an alternative method of indicating congestion to peers. This lesson introduces the concept of ECN and the Cisco IOS commands that are required to configure and monitor ECN.

Relevance

Congestion avoidance techniques offer a viable alternative to the default router congestion response, tail drop. ECN extends the number of available options for avoiding congestion by allowing routers to signal peers about congestive states without dropping packets.

Objectives

Upon completing this lesson, you will be able to configure ECN to enhance the congestion avoidance features of WRED. This includes being able to meet these objectives:

- Describe the ECN extensions to IP
- Identify key characteristics of the ECN field in IP
- Explain how ECN interacts with WRED
- Identify the Cisco IOS commands required to configure ECN
- Identify the Cisco IOS commands used to monitor ECN

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Basic knowledge of internetworking with TCP/IP concepts
- Congestion avoidance with WRED

Outline

The outline lists the topics included in this lesson.

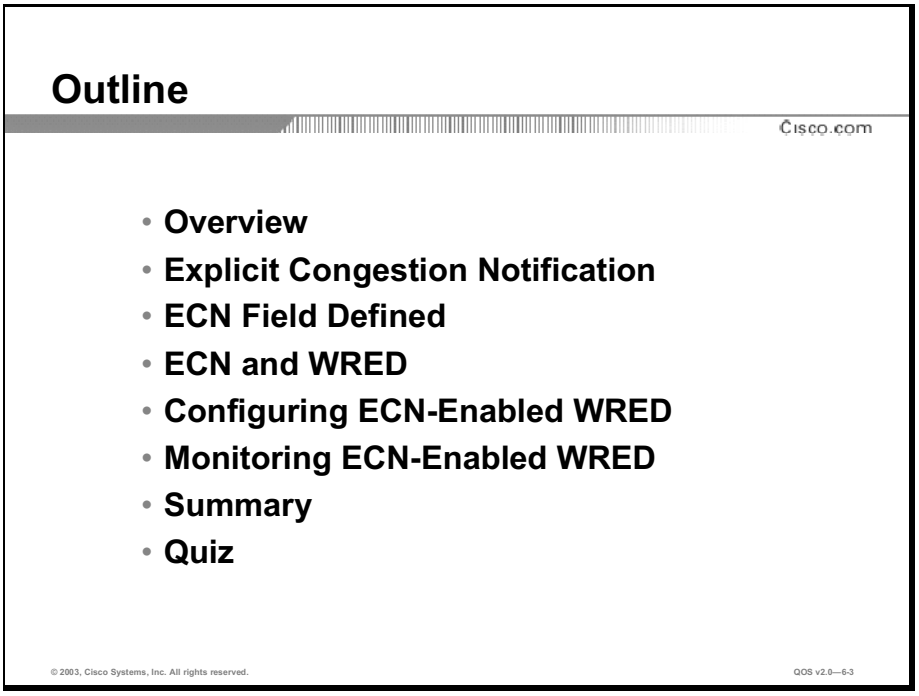

Explicit Congestion Notification

This topic describes ECN as an extension to IP.

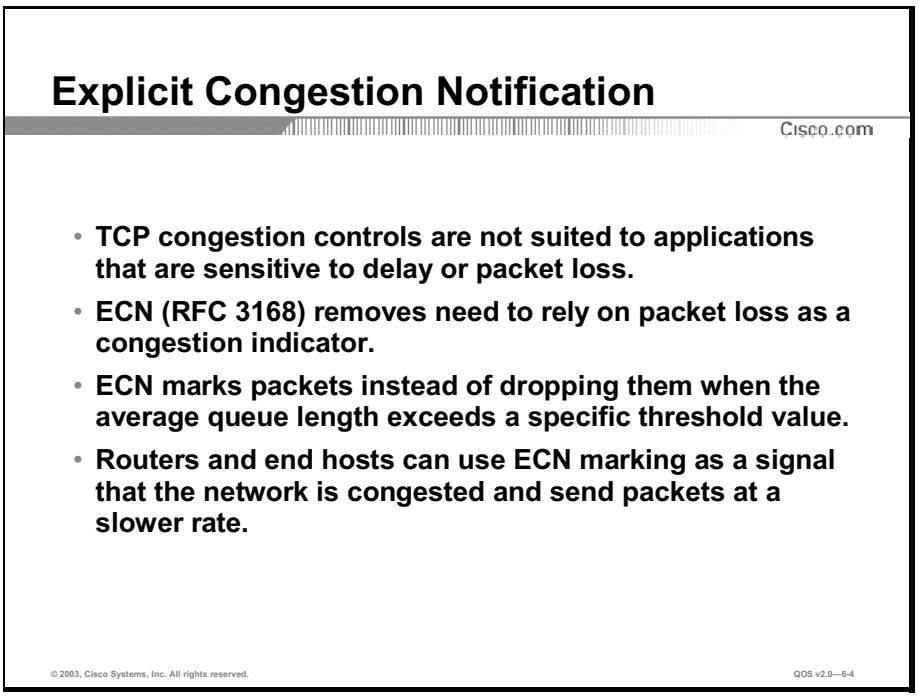

TCP determines how many unacknowledged packets it can send (window size) by gradually increasing the number of packets it sends until it experiences a dropped packet. As a result, TCP tends to cause router queues to build up at network bottleneck points. When queues become full, tail drop begins dropping all incoming packets until there is room in the queue. Tail drop does not provide differential treatment and therefore some of the fragile flow packets, sensitive to latency, may be dropped. In addition, tail drop can lead to global synchronization of packet loss across multiple flows.

Active queue management mechanisms such a RED or WRED, detect congestion before queues fill and overflow. Through the use of selective packet discard, these mechanisms provide congestion indication to end nodes. Therefore, active queue management (congestion avoidance) mechanisms can reduce queuing delays for all traffic sharing a specific queue. In addition, active queue management means that it is no longer necessary to rely on buffer overflow as the only means of indicating congestion.

Traditional active queue management mechanisms, such as RED, rely on the use of packet drops to indicate congestion. Packet dropping in these mechanisms is based on the average queue length exceeding a predefined threshold, rather than only when queues overflow. However, because packets are dropped prior to queues actually overflowing, the router dropping the packet is not always constrained by memory limitations and needs to actually drop the packet. With the “Addition of Explicit Congestion Notification to IP” (RFC 3168), active queue management allows routers to signal that congestion has been experienced by the router, instead of relying on the use of packet drops. Through the use of signaling congestion, aggressive flows can be slowed, thus reducing the impact of congestion and packet loss on latency-sensitive flows.

ECN Field Defined

This topic describes the characteristics of the ECN field.

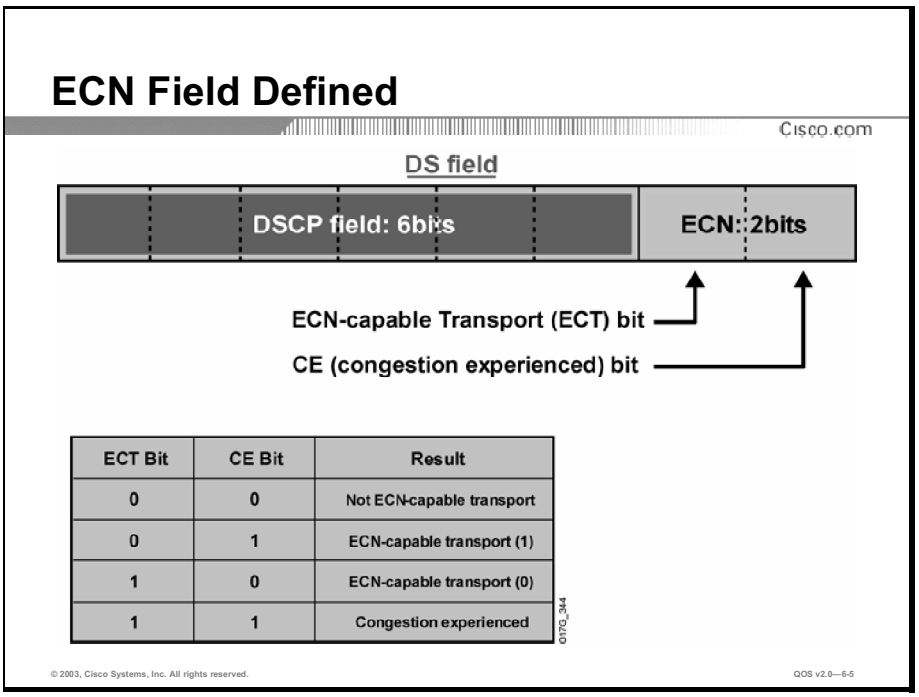

"The Addition of Explicit Congestion Notification to IP" (RFC 3168), redefines the Differentiated Services (DiffServ) field (former type of service [ToS] byte) to contain an ECN-specific field. The ECN field consists of the last two low-order bits of the DiffServ field and is comprised of the ECN-capable Transport (ECT) bit and the congestion experienced (CE) bit.

The ECT bit and the CE bit can be used to make four ECN field combinations of 00, 01, 10, and 11. The different ECT and CE bit combinations in the ECN field have the following meaning:

- 00: The ECN field combination indicates that a packet is not using ECN.
- 01 and 10: The ECN field combinations, called ECT(1) and ECT(0), respectively, are set by the data sender to indicate that the endpoints of the transport protocol are ECN-capable. Routers will treat these two field combinations identically. Data senders can use either one or both of these two combinations.
- 11: The ECN field combination indicates to the endpoints that congestion has been experienced. Packets arriving at a full queue of a router will be dropped.

ECN and WRED

This topic describes how ECN interacts with WRED.

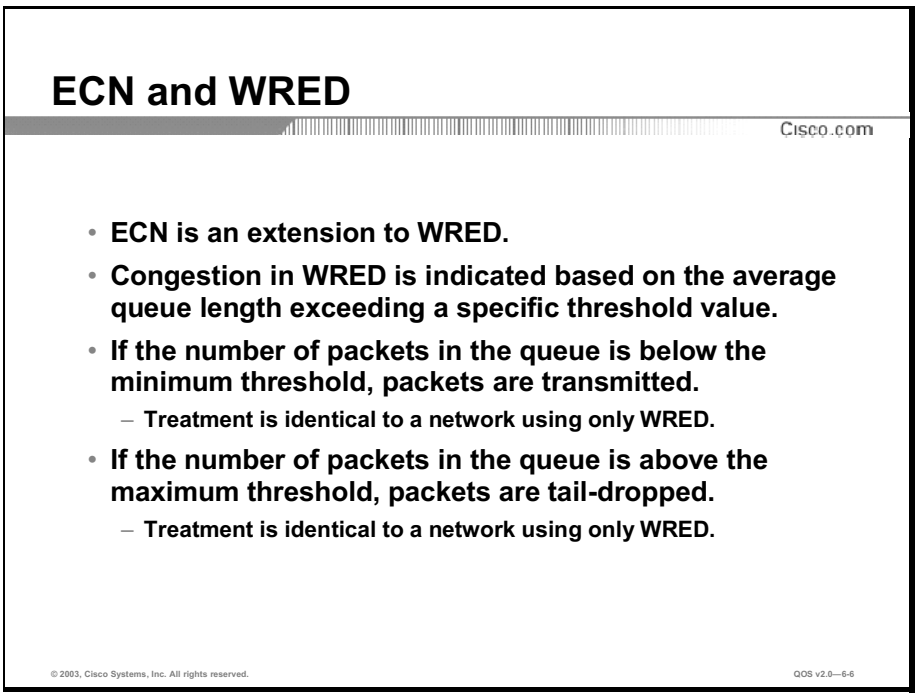

On Cisco IOS routers, ECN is an extension to WRED functionality. WRED is an active queue management mechanism that utilizes packet drops as a congestion indicator to endpoints. Packets are dropped by WRED based on the average queue length exceeding a specific set of predefined threshold values (minimum and maximum threshold). ECN is an extension to WRED in that ECN marks packets instead of dropping them when the average queue length exceeds a specific threshold value. When ECN is configured with WRED, routers and end hosts would use this marking as a signal that the network is congested and would slow down the rate at which packets are sent.

One important aspect of ECN is that it must be interoperable with non-ECN-compliant devices. Because ECN is configured as an extension to WRED, packets are treated differently by WRED when ECN has been enabled.

If the average queue length is below the defined WRED minimum threshold, all packets are queued and transmitted normally. This behavior is identical to devices that are configured to use non-ECN-enabled WRED.

If the average queue length is greater than the maximum threshold, packets are tail-dropped. This behavior is identical to devices configured to use non-ECN-enabled WRED.

ECN and WRED (Cont.)

Cisco.com

If the number of packets in the queue is between the minimum and maximum threshold, one of three scenarios can occur:

- **ECN-capable endpoints and WRED determine that the packet should be dropped based on the drop probability:**
  - **ECT and CE bits for the packet are changed to 1 and the packet is transmitted.**
- **Non ECN-capable endpoints:**
  - **The packet may be dropped based on the WRED drop probability.**
- **The network is experiencing congestion:**
  - **The packet is transmitted and no further marking is required.**

Where non-ECN-enabled WRED and ECN-enabled WRED routers differ is in how packets are treated where the average queue length is between the minimum and maximum thresholds.

If the average queue length is greater than the defined WRED minimum threshold but less than the defined WRED maximum threshold, one of three actions is possible:

- If endpoints support ECN (ECN-capable) and the WRED algorithm determines that the packet should be dropped based on the drop probability, both the ECT and CE bits are marked (set to 1), to indicate that congestion has been experienced. The packet is then transmitted towards its destination. When the ECN-capable endpoint receives the packet, it will slow its transmission rate.
- If endpoints do not support ECN (non-ECN-capable), it is up to the WRED algorithm to determine wheter to drop the packet or not based on the drop probability. This behavior is identical to devices that are configured to use non-ECN-enabled WRED.
- If the incomming packet already has the ECT and CE bits marked (set to 1), indicating that congestion has been experienced, the packet is transmitted towards its destination without modifying its ECN marking. Transmision of the packet towards the endpoint without modifying its ECN markings (or dropping the packet) are important because the ECN bits are the signal that congestion is occurring and the endpoint should slow its packet transmission rate. However, if the incomming packet arrives at a full output queue, the packet will be tail-dropped.

Configuring ECN-Enabled WRED

This topic describes the Cisco IOS commands that are required to configure ECN extensions to WRED.

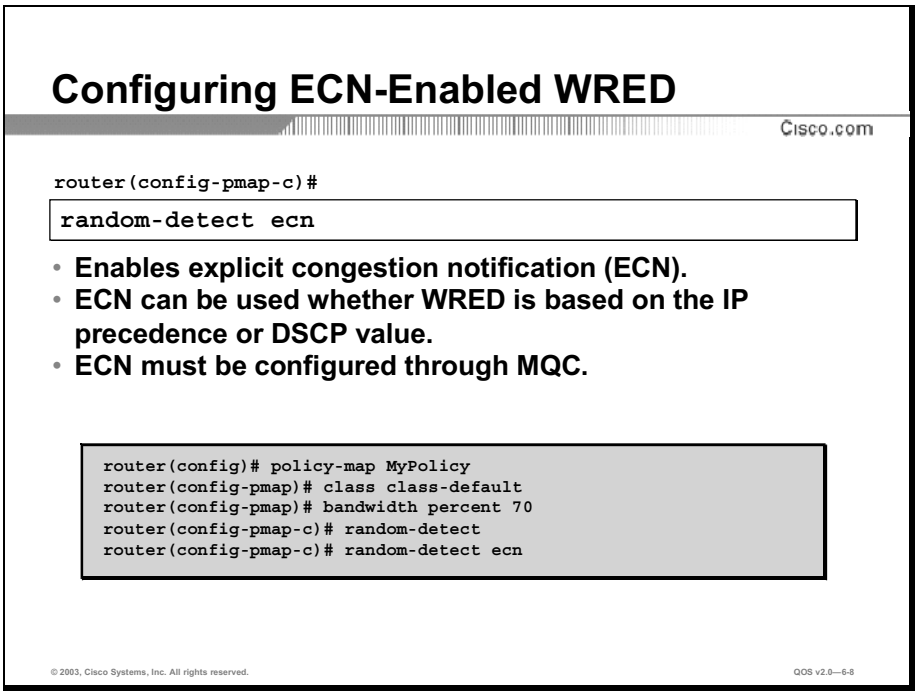

The ECN feature must be configured through the MQC. ECN is configured as part of a policy map after CB-WRED has been enabled. ECN can be used whether the CB-WRED configuration is based on IP precedence or DSCP.

Note: The ECN feature was introduced in Cisco IOS release 12.2(8)T.

Monitoring ECN-Enabled WRED

This topic describes the Cisco IOS commands that are required to monitor ECN-enabled WRED.

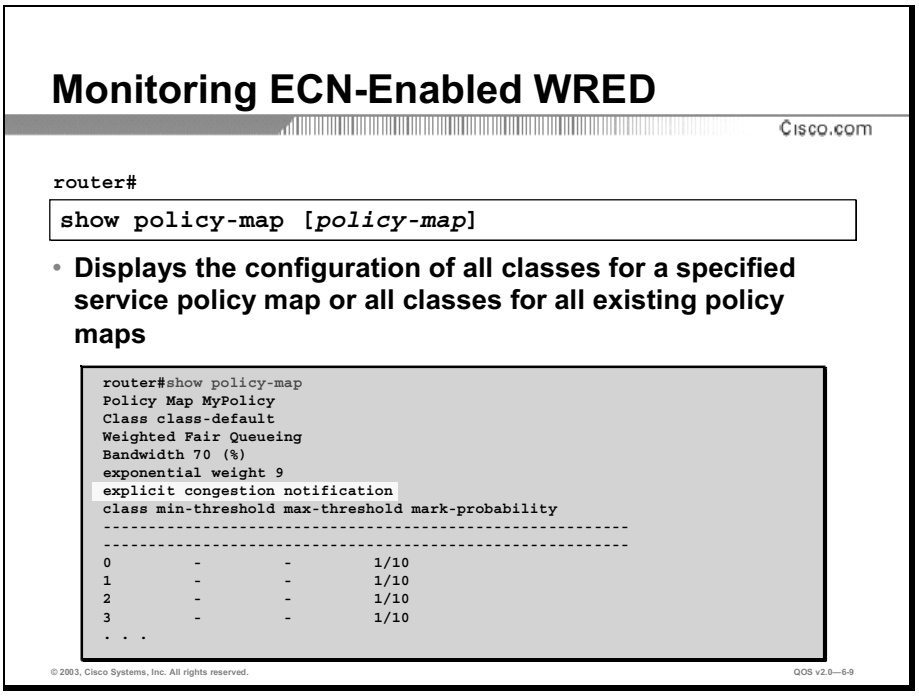

The **show policy-map** command displays the configuration of a service policy map created using the **show policy-map** command. The **show policy-map** command will display ECN marking information only if ECN is enabled on the interface. The following table explains some of the key fields of the **show policy-map** command.

show policy-map Parameters

Parameter	Description
`explicit congestion notification`	Indication that ECN is enabled.
`class`	IP precedence value.
`min-threshold`	Minimum threshold. Minimum WRED threshold in number of packets.
`max-threshold`	Maximum threshold. Maximum WRED threshold in number of packets.
`mark-probability`	Fraction of packets dropped when the average queue depth is at the maximum threshold.

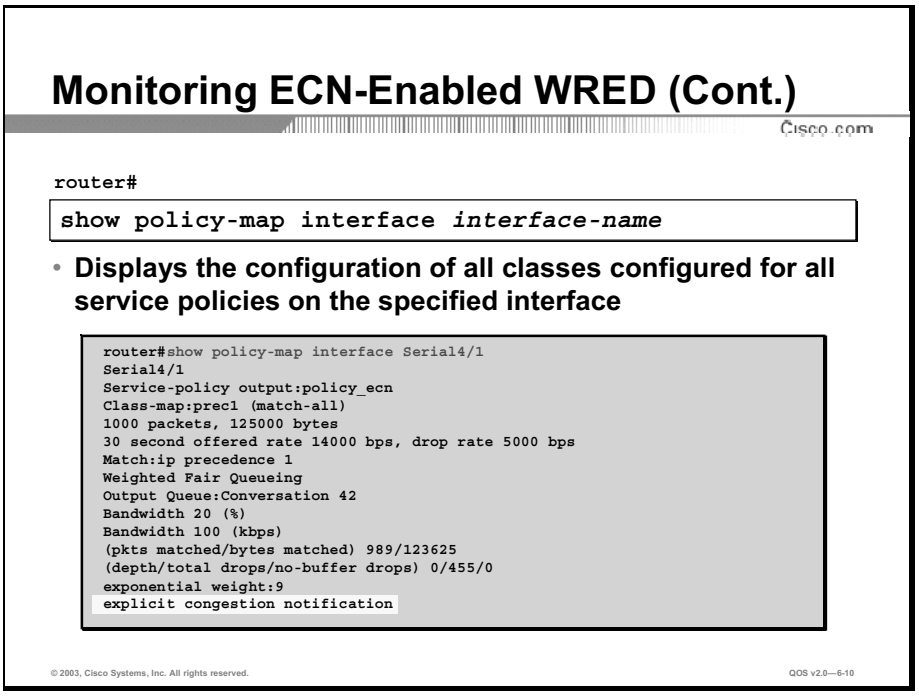

The **show policy-map interface** command displays the configuration of all classes configured for all service policies on the specified interface. The counters displayed after the **show policy-map interface** command is entered are updated only if congestion is present on the interface. The **show policy-map interface** command displays ECN marking information only if ECN is enabled on the interface. The following table explains some of the key fields of the **show policy-map interface** command.

show policy-map interface Parameters

Parameter	Description
`explicit congestion notification`	Indication that ECN is enabled.

Monitoring ECN-Enabled WRED (Cont.)

Cisco.com

```
mean queue depth:0
class Transmitted Random drop Tail drop  Minimum   Maximum   Mark
      pkts/bytes  pkts/bytes  pkts/bytes threshold threshold probability
0         0/0       0/0        0/0          20        40        1/10
1       545/68125   0/0        0/0          22        40        1/10
2         0/0       0/0        0/0          24        40        1/10
3         0/0       0/0        0/0          26        40        1/10
4         0/0       0/0        0/0          28        40        1/10
5         0/0       0/0        0/0          30        40        1/10
6         0/0       0/0        0/0          32        40        1/10
7         0/0       0/0        0/0          34        40        1/10
rsvp      0/0       0/0        0/0          36        40        1/10

     class        ECN Mark
                  pkts/bytes
       0             0/0
       1             43/5375
       2             0/0
       3             0/0
       4             0/0
       5             0/0
       6             0/0
       7             0/0
     rsvp            0/0
```


show policy-map interface Parameters (Cont.)

Parameter	Description
mean queue depth	Average queue depth based on the actual queue depth on the interface and the exponential weighting constant. It is a moving average. The minimum and maximum thresholds are compared against this value to determine drop decisions.
class	IP precedence value.
Transmitted pkts/bytes	Number of packets (also shown in bytes) transmitted.
Random drop pkts/bytes	Number of packets (also shown in bytes) randomly dropped when the mean queue depth is between the minimum threshold value and the maximum threshold value for the specified IP precedence value.
Tail drop pkts/bytes	Number of packets dropped when the mean queue depth is greater than the maximum threshold value for the specified IP precedence value.
Minimum threshold	Minimum WRED threshold in number of packets.
Maximum threshold	Maximum WRED threshold in number of packets.
Mark probability	Fraction of packets dropped when the average queue depth is at the maximum threshold.
ECN Mark pkts/bytes	Number of packets (also shown in bytes) marked by ECN.

Summary

This section summarizes the key points discussed in this lesson.

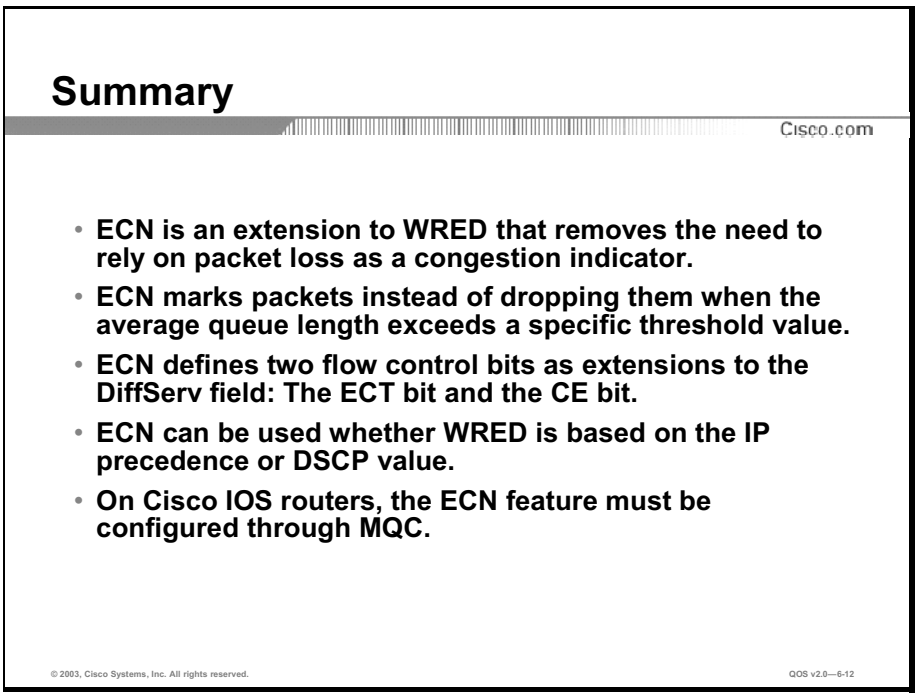

References

For additional information, refer to these resources:

- For more information on using WRED with ECN, refer to, “WRED—Explicit Congestion Notification,” at the following URL: http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122t/122t8/ftwrdecn.htm
- For more information on ECN, refer to, “RFC 3168: The Addition of Explicit Congestion Notification to IP,” at the following URL: http://www.ietf.org/rfc/rfc3168.txt

Next Steps

For the associated lab exercise, refer to the following section of the course Lab Guide:

- Lab Exercise 6-1: Configuring DSCP-Based WRED

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) What is a key advantage of active queue management using Explicit Congestion Notification?

A) Packets from aggressive flows are randomly dropped as a means of avoiding global synchronization, congested queues, and tail drop.

B) Congestion is signaled without drops, causing aggressive flows to be slowed, reducing the impact of congestion and packet loss on latency-sensitive flows.

C) Packet drops are used as an indication that senders should reduce transmission rates, removing the congested state.

D) Packets are not dropped unless the queue exhausts its packet buffer memory, allowing endpoints to request a transmission rate for sending flows into the network controls congestion. ECN maintains the flow transmission process.

Q2) What two bits are defined as part of the ECN extensions to IP? (Choose two.)

A) transmission control bit

B) ECN-capable transport bit

C) congestion-experienced bit

D) Forward Explicit Congestion Notification bit

Q3) What is indicated by the ECN field bit combination of 00?

A) Congestion has been experienced in the network.

B) The packet is not from an ECN- capable endpoint.

C) The endpoint recognizes ECN, but no congestion has been encountered.

D) Packets have been dropped because of experienced congestion.

Q4) How does ECN extend the capabilities of WRED-configured Cisco routers?

A) ECN allows differential packet dropping based on IP precedence or DSCP.

B) ECN allows routers to dynamically allocate packet buffers based on congestion signaling.

C) ECN removes the need to drop packets when the average queue length is between the two defined WRED thresholds.

D) ECN uses signaling to manage previously defined packet transmission rates of network endpoints.

Q5) In an ECN-enabled WRED configuration, what happens if the average queue length grows larger than the minimum threshold but the endpoints are not ECN-capable?

A) The router will automatically allocate more packet buffers to prevent packet loss.

B) The router will request ECN options from the endpoints.

C) The packet will be dropped or transmitted according to normal WRED drop probability calculations.

D) The packet will be dropped as a means of slowing the packet transmission rate.

Q6) What two requirements must be met prior to configuring ECN on Cisco IOS routers? (Choose two.)

A) CB-WRED must first be configured using the **random-detect** command.

B) CBWFQ must first be configured using MQC.

C) WRED must be configured to use DSCP as its weight.

D) The ECN bits must be reset (marked 0,0) using a policy map.

Q7) What two Cisco IOS commands can be used to verify that ECN has been enabled? (Choose two.)

A) **show interface**

B) **show policy-map**

C) **show policy-map interface**

D) **show interface random-detect**

Quiz Answer Key

Q1) B

Relates to: Explicit Congestion Notification

Q2) B, C

Relates to: ECN Field Defined

Q3) B

Relates to: ECN Field Defined

Q4) C

Relates to: ECN and WRED

Q5) C

Relates to: ECN and WRED

Q6) A, B

Relates to: Configuring ECN-Enabled WRED

Q7) B, C

Relates to: Monitoring ECN-Enabled WRED

Module Assessment

Overview

Use this assessment to test what you learned in this module. The correct answers and solutions are found in the Module Assessment Answer Key.

Quiz: Congestion Avoidance

Complete the Quiz to assess what you have learned in the module.

Objectives

This activity tests your knowledge on how to meet these objectives:

- Explain the problems that may result from the limitations of TCP congestion management mechanisms on a converged network
- Explain how RED can be used to avoid congestion
- Configure CB-WRED to avoid congestion
- Configure ECN to enhance the congestion avoidance features of WRED

Instructions

Complete these steps:

Step 1 Answer all questions in this quiz by selecting the best answer(s) to each question.

Step 2 Verify your results against the answer key located at the end of this section.

Step 3 Review the topics in this module that relate to the questions that you answered incorrectly.

Q1) What are two ways in which TCP manages congestion? (Chose two.)

A) TCP uses tail drop on queues that have reached their queue limit.

B) TCP uses dropped packets as an indication that congestion has occurred.

C) TCP uses variable window sizes to reduce and increase the rates at which packets are sent.

D) TCP measures the average size of device queues and drops packets, linearly increasing the amount of dropped packets with the size of the queue.

Q2) What are two active congestion management mechanisms available on Cisco IOS routers? (Choose two.)

A) tail drop

B) weighted round robin

C) explicit congestion notification

D) weighted random early detection

Q3) Two stations (A and B) are communicating using TCP. Station A has negotiated a TCP window size of 5 and as a result sends 5 packets to station B.

Station A receives 3 ACK messages from station B indicating ACK 3.

Which two of the following descriptions best describe the status of the communication between A and B? (Choose two)

A) Station B is acknowledging receipt of packets 1, 2, and 3, but has lost packets 4 and 5.

B) Station A initiates a fast-retransmit and immediately sends packet 3 to B.

C) Station B has not received packet 3.

D) Station B has received packets 1, 2, and 3, but not packet 4. It cannot be determined where packet 5 was received at B until packet 4 has been sent.

E) Station A will send packets 4 and 5 to station B upon receipt of the station B ACK.

Q4) What are three important limitations of using a tail-drop mechanism to manage queue congestion? (Choose three.)

A) Tail drop can cause many flows to synchronize, lowering overall link utilization.

B) Tail drop can cause starvation of fragile flows.

C) Tail drop increases the amount of packet buffer memory required, as queues must be full before congestion management becomes active.

D) Tail drop results in variable delays, which can interfere with delay-sensitive traffic flows.

Q5) What are three advantages of active congestion management using RED? (Choose three.)

A) RED uses selective packet discard to eliminate global synchronization of TCP flows.

B) RED avoids congestion by ensuring that interface queues never become full.

C) RED increases the overall utilization of links.

D) RED uses selective packet discard to penalize aggressive flows.

Q6) A specific RED profile has been configured with a mark probability denominator of 1. What is the effect of this configuration on packet loss as the average queue length reaches the maximum threshold?

A) Given this configuration, no packets will be dropped until the average queue length is greater than the maximum threshold.

B) For every active traffic flow, one packet will be discarded.

C) When the average queue length is at the maximum threshold, all packets are dropped.

D) This is an invalid configuration.

Q7) Refer to the following RED traffic profile. How will the RED traffic profile in the figure affect the traffic flows to which it is applied? (Choose two.)

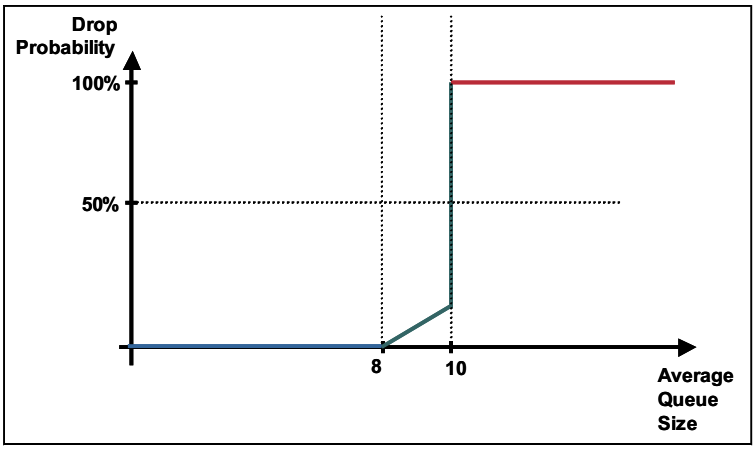

Figure: RED Traffic Profile

A) Packets may be dropped unnecessarily as the minimum threshold is too low.

B) This profile can result in global synchronization as the difference between the minimum and maximum thresholds is too small.

C) RED will not be effective as the mark probability denominator is 50 when the average queue length reaches the maximum threshold.

D) The reduced size of the maximum threshold will prevent tail drop and maximize link utilization.

Q8) What are the three traffic drop modes in Random Early Detection? (Choose three.)

A) no drop

B) full drop

C) random drop

D) deferred drop

Q9) What two QoS markers can you base the “weight” in WRED on when configuring CB-WRED? (Choose two.)

A) CoS

B) DSCP

C) QoS group

D) IP precedence

Q10) What are two requirements for configuring CB-WRED? (Choose two.)

A) An MQC configuration that includes a policy map must be configured.

B) CEF must be enabled for IP.

C) Random detect must be enabled for DSCP-based CB-WRED.

D) A previous configuration of CB-WFQ must be present.

Q11) Given the following CB-WRED configuration, what command should be entered in the Bronze traffic class to properly enable CB-WRED using a minimum threshold of 22, a maximum threshold of 36, and a drop probability of 10 percent?

```
class-map Bronze
 match ip dscp cs1
!
policy-map Policy1
  class Bronze
    bandwidth percent 15
    random-detect dscp-based
[ <- _____ -> ]
!
  class class-default
    fair-queue
    random-detect dscp-based
```

A) **random-detect dscp-based 22 36 10**

B) **random-detect dscp-based cs1 22 36 10**

C) **random-detect dscp cs1 22 36**

D) **random-detect dscp cs1 10 22 36**

Q12) What will a router do with a newly arriving packet if its output queue is full and ECN fields are both set to a 1?

A) drop the last packet on the queue and enqueue the newly arriving packet

B) perform a tail drop and drop the new packet

C) move the packet to the head of the queue to ensure that the receiver is signaled about the network congestion condition

D) allocate additional interface buffers to store the packet since it contains congestion notification information

Q13) Given the following configuration, what Cisco IOS configuration command must be added to the default class to enable Explicit Congestion Notification (ECN)?

```
policy-map MyPolicy
  class class-default
  bandwidth percent 70
  random-detect
[ <- _____ -> ]
```

A) **wred ecn**
B) **ecn enable**
C) **random-detect ecn**
D) **random-detect ecn enable**

Scoring

You have successfully completed the quiz for this lesson when you earn a score of 80 percent or better.

Module Assessment Answer Key

Q1) B, C

Relates to: Introduction to Congestion Avoidance

Q2) C, D

Relates to: Weighted Random Early Detection (WRED), Explicit Congestion Notification (ECN)

Q3) B, C

Relates to: Introduction to Congestion Avoidance

Q4) A, B, D

Relates to: Introduction to Congestion Avoidance

Q5) A, C, D

Relates to: Introduction to RED

Q6) C

Relates to: Introduction to RED

Q7) A, B

Relates to: Introduction to RED

Q8) A, B, C

Relates to: Introduction to RED

Q9) B, D

Relates to: Configuring Class-Based Weighted RED

Q10) A, D

Relates to: Configuring Class-Based Weighted RED

Q11) C

Relates to: Configuring Class-Based Weighted RED

Q12) B

Relates to: Configuring Explicit Congestion Notification

Q13) C

Relates to: Configuring Explicit Congestion Notification

Module Summary

This topic summarizes the key points discussed in this module.

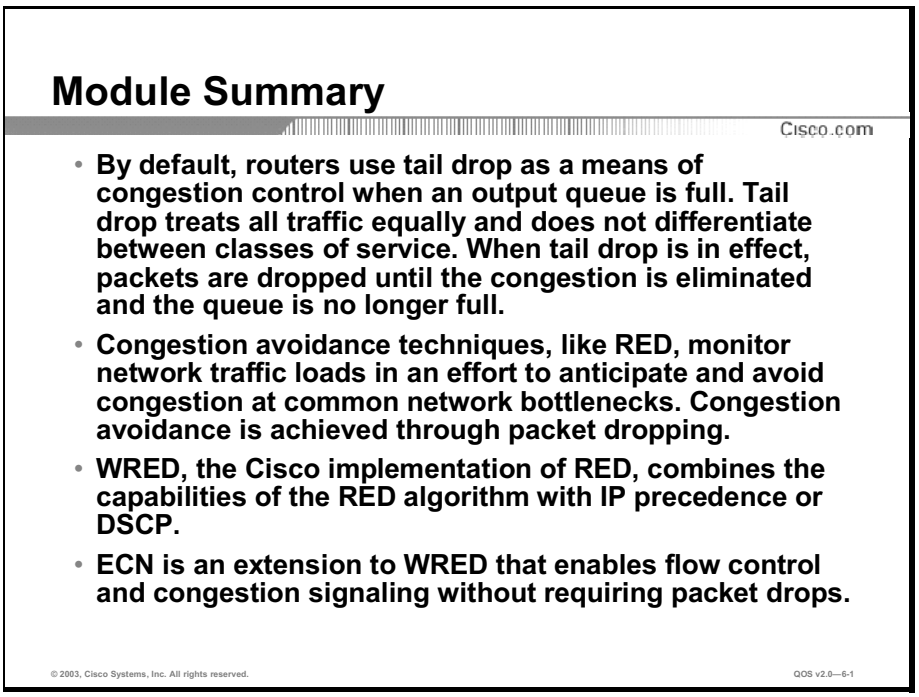

Congestion management is an area of concern for all networks that require a differentiated treatment of packet flows. Active queue management mechanisms address the limitations of relying solely on TCP congestion management techniques, which simply wait for queues to overflow and then drop packets to signal that congestion has occurred. Congestion avoidance mechanisms such as RED and WRED allow for specific packet flows to be selectively penalized and slowed by applying a traffic profile. Traffic flows are matched against this profile and transmitted or dropped depending upon the average length of the interface output queue. In addition, RED and WRED are extremely effective tools at preventing global synchronization of many TCP traffic flows. Another active queue management technique is ECN. ECN is an extension to WRED that allows for signaling to be sent to ECN-enabled endpoints, instructing them to reduce their packet transmission rates. ECN also provides the benefit of not requiring packet drops when the WRED drop probability indicates otherwise.

Module 7

Traffic Policing and Shaping

Overview

Within a network, different forms of connectivity can have significantly different costs to an organization. For instance, a LAN connection will cost considerably less than a WAN connection for an equal amount of bandwidth. Because WAN bandwidth is relatively expensive, many organizations would like to limit the amount of traffic that specific applications can send. This is especially true when enterprise networks use Internet connections for remote site and extranet connectivity. Downloading non-business-critical images, music, and movie files can greatly reduce the amount of bandwidth available to other mission-critical applications. Traffic policing and traffic shaping are two quality of service (QoS) techniques that can be used to limit the amount of bandwidth a specific application can use on a link.

From a services perspective, many service providers would like to install a larger bandwidth connection to customers but provision a smaller circuit so that incremental bandwidth upgrades do not require provisioning new circuits or installing new equipment. Called sub-rate access, traffic policing and traffic shaping techniques can also assist in this regard.

In this module, the operation of traffic policing and traffic shaping, and how these techniques can be used to rate-limit traffic is discussed. As Frame Relay WANs have specific requirements, class-based traffic shaping on Frame Relay networks is also covered in this module.

Module Objectives

Upon completing this module, you will be able to use Cisco QoS traffic policing and traffic-shaping mechanisms to effectively limit the rate of network traffic.

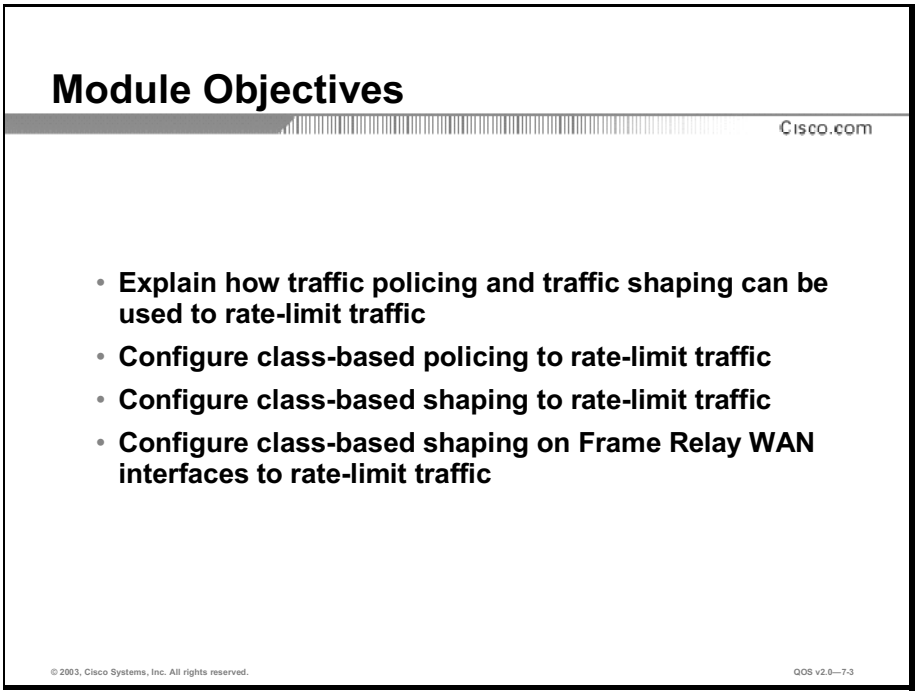

Module Outline

The outline lists the components of this module.

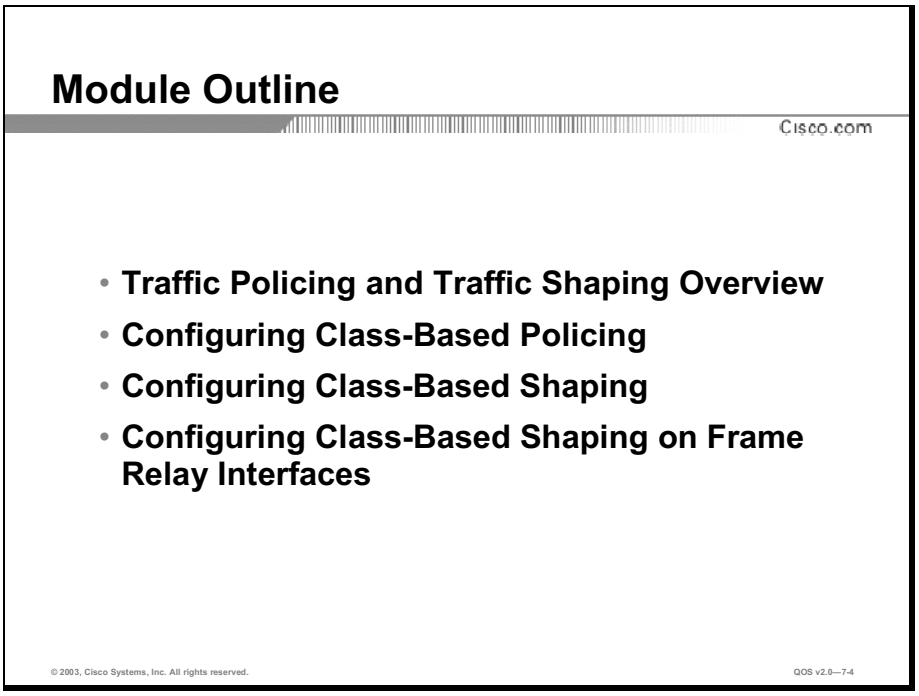

Traffic Policing and Traffic Shaping Overview

Overview

Traffic policing can be used to control the maximum rate of traffic sent or received on an interface. Traffic policing is often configured on interfaces at the edge of a network to limit traffic into or out of the network. Traffic that falls within the rate parameters is sent, whereas traffic that exceeds the parameters is dropped or sent with a different priority.

Traffic shaping can be used to control the traffic going out an interface in order to match its flow to the speed of the remote target interface and to ensure that the traffic conforms to policies contracted for it. Thus, traffic adhering to a particular profile can be shaped to meet downstream requirements, thereby eliminating bottlenecks in topologies with data-rate mismatches.

Traffic policing and traffic shaping differ in the way they respond to traffic violations. Policing typically drops traffic, while shaping typically queues excess traffic by using a shaping queue to hold packets and shape the flow when the data rate of the source is higher than expected.

This lesson describes the traffic-policing and traffic-shaping QoS mechanisms that are used to limit the available bandwidth to traffic classes. Because both traffic policing and traffic shaping use the token bucket metering mechanism, this lesson also explains how a token bucket works.

Relevance

Enterprise and service provider networks have a variety of requirements for traffic conditioning. Traffic policing and traffic shaping are important traffic conditioning tools to allow traffic rates to be controlled.

Objectives

Upon completing this lesson, you will be able to explain how traffic policing and traffic shaping can be used to condition traffic. This includes being able to meet these objectives:

- Describe the purpose of traffic conditioning using traffic policing and traffic shaping
- List key benefits of traffic conditioning using traffic policing and traffic shaping
- Differentiate between the features of traffic policing and traffic shaping
- Explain how network devices measure traffic rates
- Explain how traffic can be policed using a single token bucket scheme
- Explain how traffic can be policed using a dual token bucket scheme
- Explain how traffic can be policed using a dual-rate metering scheme
- Explain how traffic can be shaped using a single token bucket scheme
- Identify the key traffic policing and shaping mechanisms available in Cisco IOS software and differentiate among them
- Identify the points in a network where traffic conditioning can most effectively be employed

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- A good understanding of Frame Relay operation and configuration
- A good understanding of traffic classification

Outline

The outline lists the topics included in this lesson.

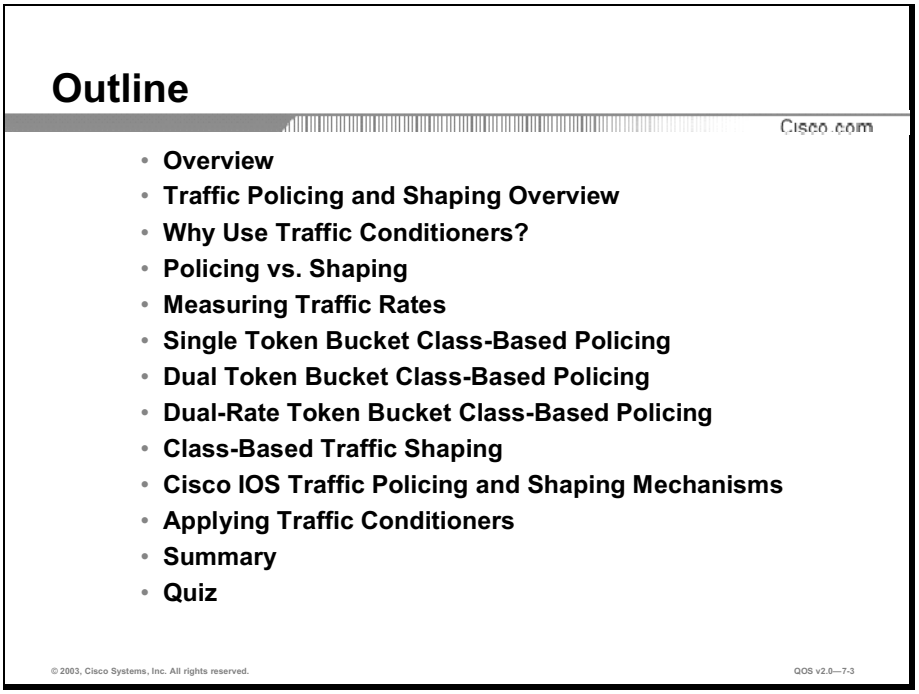

Traffic Policing and Shaping Overview

This topic describes the purpose of traffic conditioning using traffic policing and traffic shaping.

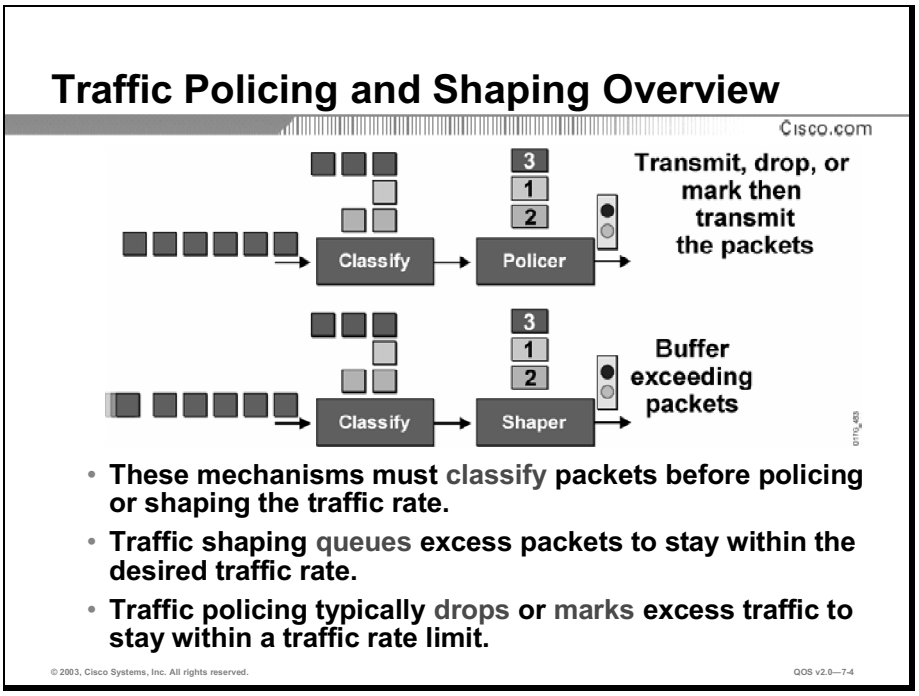

Both traffic shaping and policing mechanisms are traffic-conditioning mechanisms that are used in a network to control the traffic rate. Both mechanisms use classification so that they can differentiate traffic. They both measure the rate of traffic and compare it to the configured traffic-shaping or traffic-policing policy.

The difference between traffic shaping and policing can be described in terms of their implementation:

- Traffic shaping buffers excessive traffic so that it stays within the desired traffic rate. With traffic shaping, traffic bursts are smoothed out by queuing the excess traffic to produce a steadier flow of data. Reducing traffic bursts helps reduce congestion in the network.
- Traffic policing drops excess traffic in order to control traffic flow within specified rate limits. Traffic policing does not introduce any delay to traffic that conforms to traffic policies. It can, however, cause more TCP retransmissions, because traffic in excess of specified limits is dropped.

In Cisco IOS software, traffic-policing mechanisms such as class-based policing or committed access rate (CAR) also have marking capabilities in addition to rate-limiting capabilities. Instead of dropping the excess traffic, traffic policing can alternatively mark and then send the excess traffic. This allows the excess traffic to be re-marked with a lower priority before they are sent out.

Traffic shapers like class-based shaping, generic traffic shaping, Frame Relay traffic shaping (FRTS), or virtual IP (VIP)-based distributed traffic shaping in Cisco IOS software do not have the ability to mark traffic.

Why Use Traffic Conditioners?

This topic lists the key benefits of traffic conditioning using traffic policing and traffic shaping.

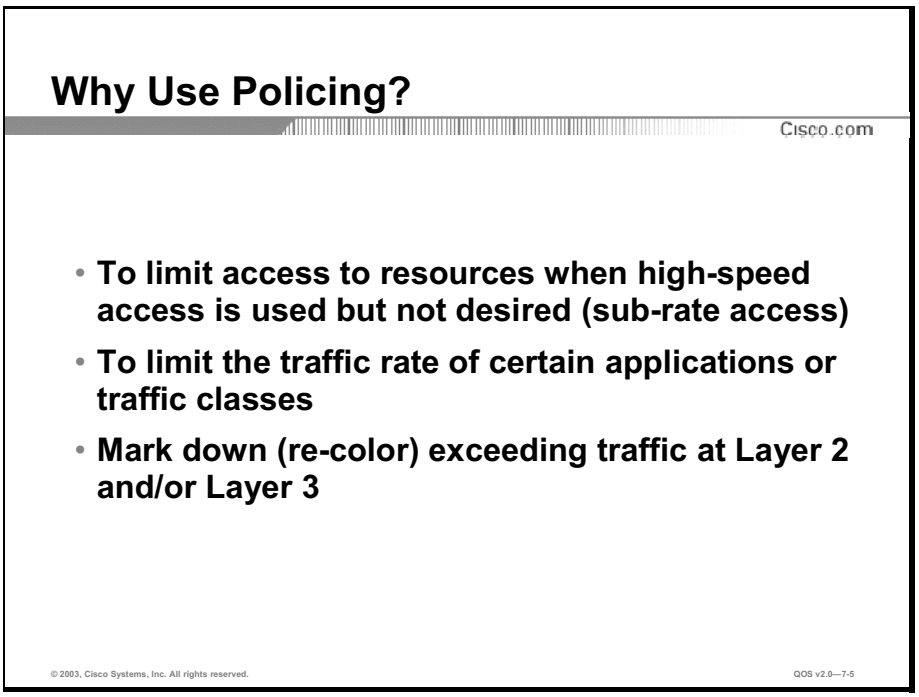

Traffic policing is typically used to satisfy one of the following requirements:

- Limiting the access rate on an interface when high-speed physical infrastructure is used in transport. Rate limiting is typically used by service providers to offer customers sub-rate access. For example, a customer may have an OC-3 connection to the service provider but pay only for a T1 access rate. The service provider can rate-limit the customer traffic to T1 speed.
- Engineering bandwidth so that traffic rates of certain applications or classes of traffic follow a specified traffic rate policy. For example, rate limiting traffic from file sharing applications to 64 kbps maximum.
- Re-marking excess traffic with a lower priority at Layer 2 and Layer 3 or both before sending them out. Cisco class-based traffic policing can be configured as a multiaction policer to mark packets at both Layer 2 and Layer 3. For example, excess traffic can be re-marked to a lower differentiated services code point (DSCP) value and also have the Frame Relay discard eligible (DE) bit set before the packet is sent out.

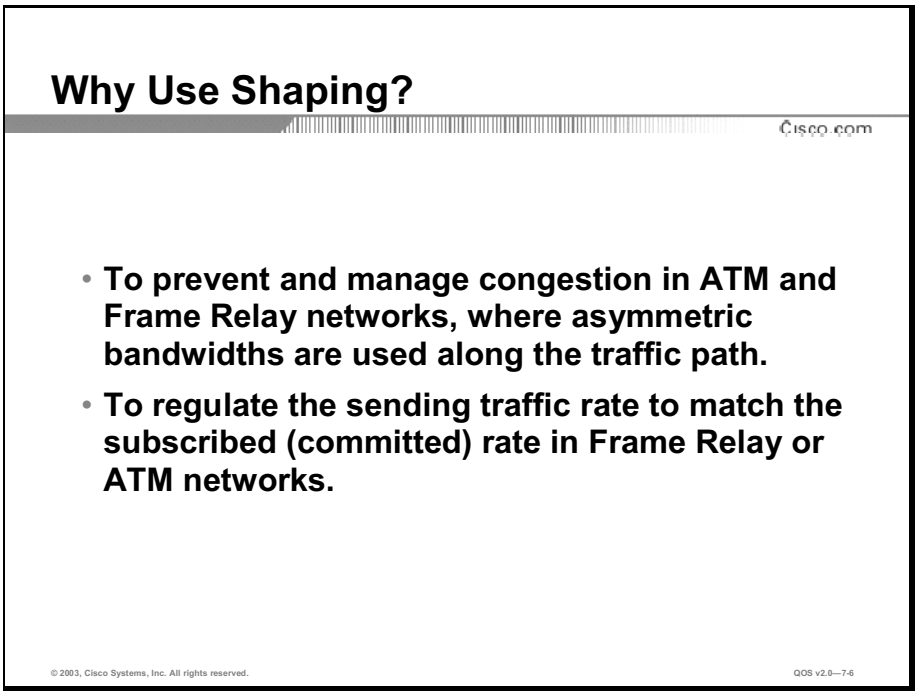

Traffic shaping is typically used to:

- Prevent and manage congestion in ATM and Frame Relay networks, where asymmetric bandwidths are used along the traffic path. If shaping is not used, then buffering can occur at the slow (usually the remote) end, which can lead to queuing causing delays and overflow causing drops.
- Prevent dropping of noncompliant traffic by the ATM or Frame Relay service provider by not allowing the traffic to burst above the subscribed (committed) rate. This allows the customer to keep local control of traffic regulation.

Example: Traffic Policing

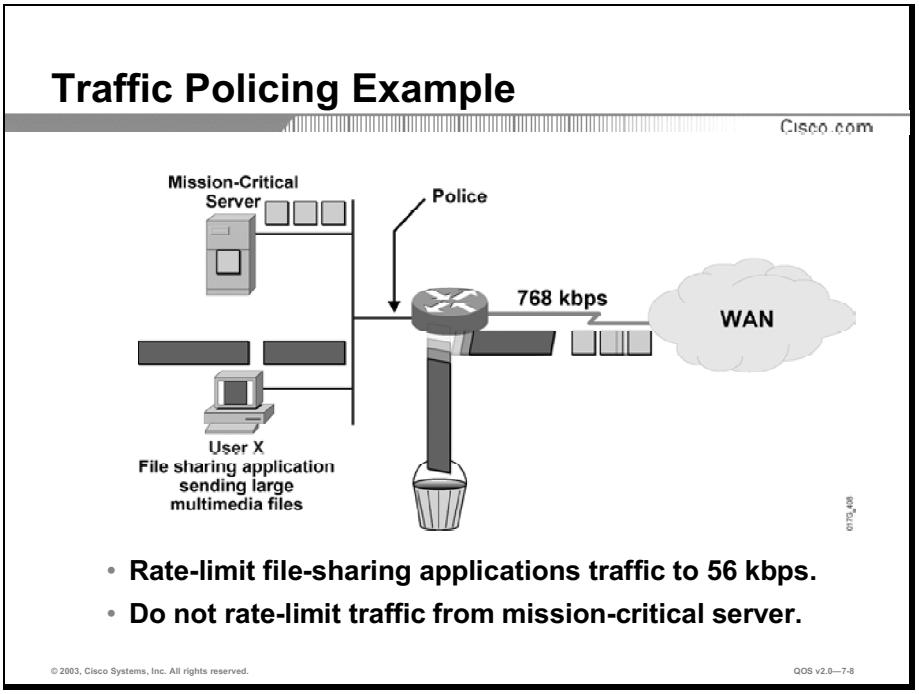

The figure shows an application for traffic policing. Traffic policing can be used to divide the shared resource (the upstream WAN link) between many flows. In this example, the router Fast Ethernet interface has an input traffic-policing policy applied to it, where the mission-critical server traffic rate is not rate-limited but the User X file-sharing application traffic is rate-limited to 56 kbps. All file-sharing application traffic from User X that exceeds the rate limit of 56 kbps will be dropped.

Example: Traffic Shaping

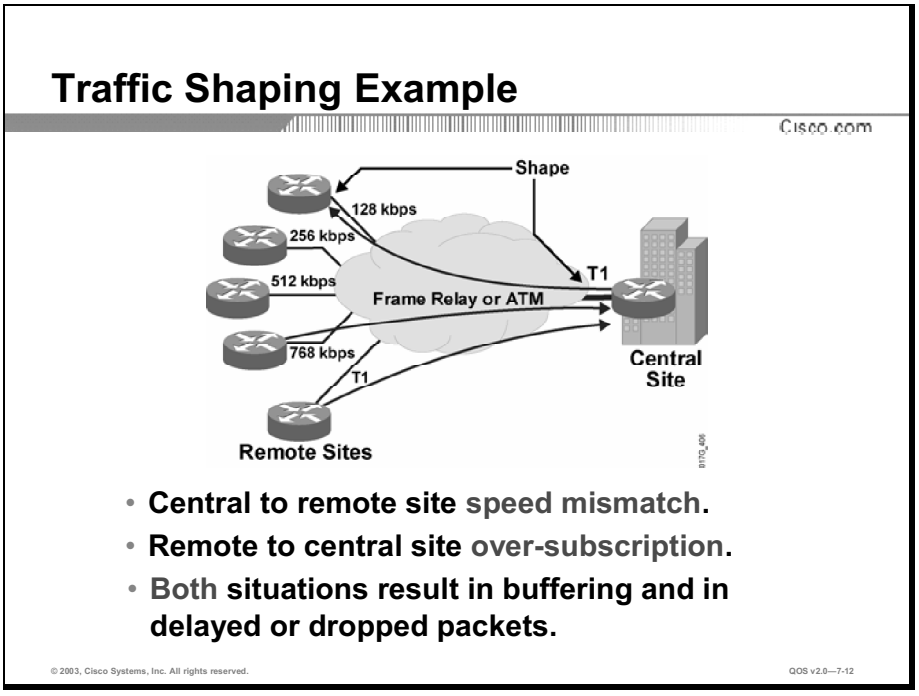

Traffic-shaping tools limit the transmit rate from a source by queuing the excess traffic. This limit is typically a value lower than the line rate of the transmitting interface. Traffic shaping can be used to account for speed mismatches that are common in nonbroadcast multiaccess (NBMA) networks, such as Frame Relay and ATM.

In the figure, two types of speed mismatches are shown:

- The central site can have a higher speed link than the remote site. Thus, traffic shaping can be deployed at the central site router to shape the traffic rate out of the central site router to match the link speed of the remote site. For example, the central router can shape the permanent virtual circuit (PVC) (going to the top remote-site router) outgoing traffic rate to 128 kbps to match that remote site link speed. At each remote site router, traffic shaping is also implemented to shape the remote site outgoing traffic rate to 128 kbps to match the committed information rate (CIR).
- The aggregate link speed of all the remote sites can be higher than the central site link speed (over-subscribing the central site link speed). In this case, the remote-site routers can be configured for traffic shaping to avoid oversubscription at the central site. For example, the bottom two remote-site routers can be configured to shape the PVC outgoing traffic rate to 256 kbps to avoid the central-site router from being over-subscribed.

Policing vs. Shaping

This topic describes the difference between the features of traffic policing and traffic shaping.

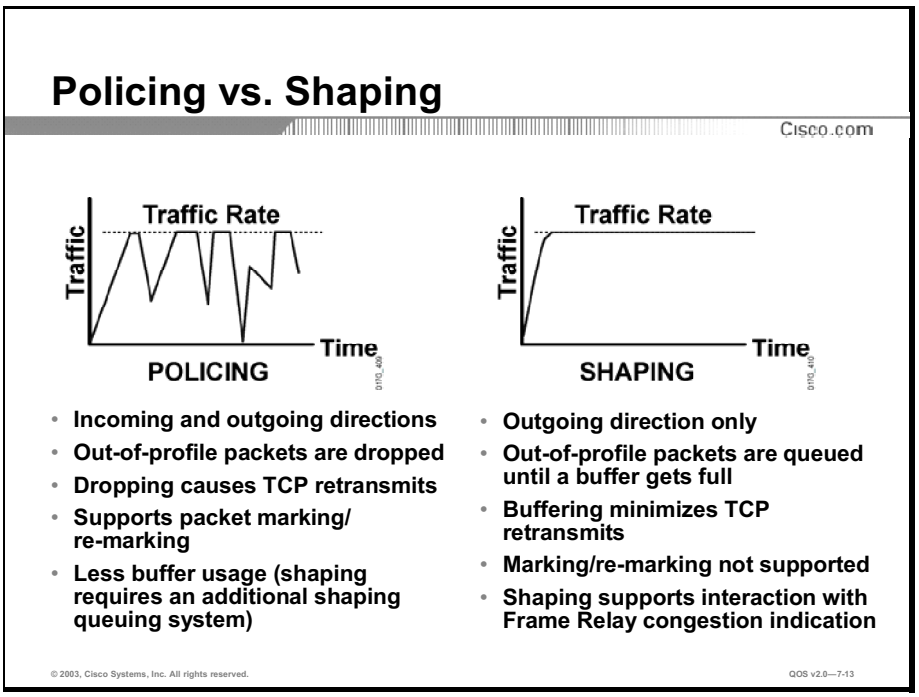

Shaping queues excess traffic by holding packets inside a shaping queue. Traffic shaping is used to shape the outbound traffic flow when the outbound traffic rate is higher than a configured shape rate. Traffic shaping smoothes traffic by storing traffic above the configured rate in a shaping queue. Therefore, shaping increases buffer utilization on a router and causes nondeterministic packet delays. Traffic shaping can also interact with a Frame Relay network, adapting to indications of Layer 2 congestion in the WAN. For example, if the backward explicit congestion notification (BECN) bit is received, the router can lower the rate limit to help reduce congestion in the Frame Relay network.

Policing can be applied to either the inbound or outbound direction while a shaper can only be applied in the outbound direction. Policing drops nonconforming traffic instead of queuing them like a shaper and also supports marking of traffic. Traffic policing is more efficient in terms of memory utilization than traffic shaping because no additional queuing of packets is needed.

Both traffic policing and shaping ensure that traffic does not exceed a bandwidth limit, but they have different impacts on the traffic:

- Policing drops packets more often, generally causing more retransmissions of connection-oriented protocols like TCP.
- Shaping adds variable delay to traffic, possibly causing jitter.

Measuring Traffic Rates

This topic explains how a token bucket can be used by network devices to measure traffic rates.

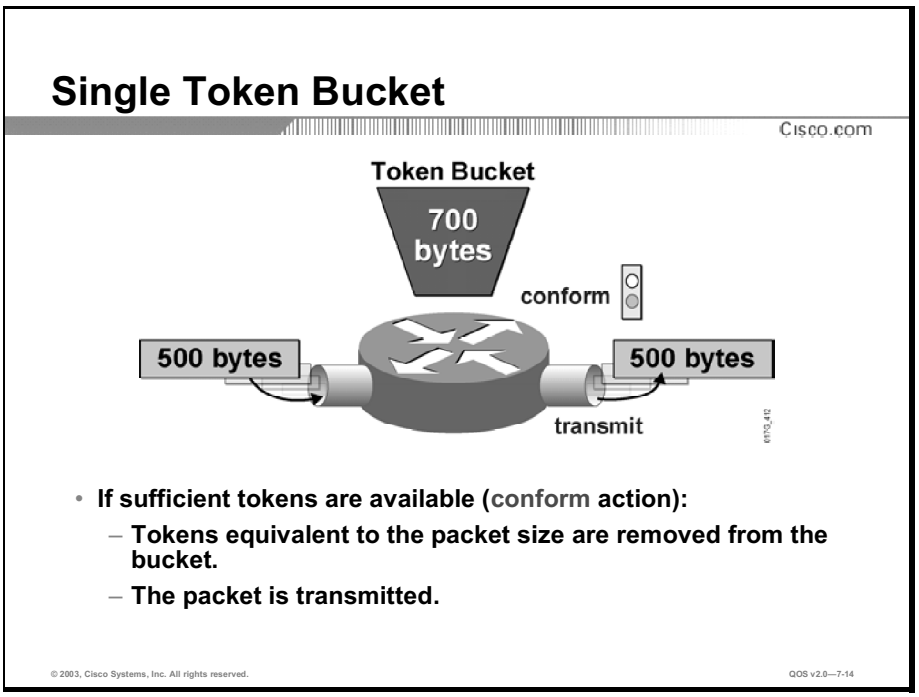

The token bucket is a mathematical model that is used by routers and switches to regulate traffic flow. The model has two basic components:

- **Tokens:** Where each token represents the permission to send a fixed number of bits into the network. Tokens are put into a token bucket at a certain rate by the Cisco IOS software.
- **Token bucket:** Has the capacity to hold a specified amount of tokens. Each incoming packet, if forwarded, takes tokens from the bucket, representing the packet size. If the bucket fills to capacity, newly arriving tokens are discarded. Discarded tokens are not available to future packets. If there are not enough tokens in the token bucket to send the packet, the traffic conditioning mechanisms may:
  - Wait for enough tokens to accumulate in the bucket: traffic shaping
  - Discard the packet: traffic policing

Using a single token bucket model, the measured traffic rate can be conforming or exceeding the specified traffic rate. The measured traffic rate is conforming if there are enough tokens in the single token bucket to transmit the traffic. The measured traffic rate is exceeding if there are not enough tokens in the single token bucket to transmit the traffic.

The figure shows a single token bucket traffic-policing implementation. Starting with a current capacity of 700 bytes worth of tokens accumulated in the token bucket, when a 500-byte packet arrives at the interface, its size is compared to the token bucket capacity (in bytes). The 500-byte packet conforms to the rate limit (500 bytes < 700 bytes), and the packet is forwarded: 500 bytes worth of tokens are taken out of the token bucket leaving 200 bytes worth of tokens for the next packet.

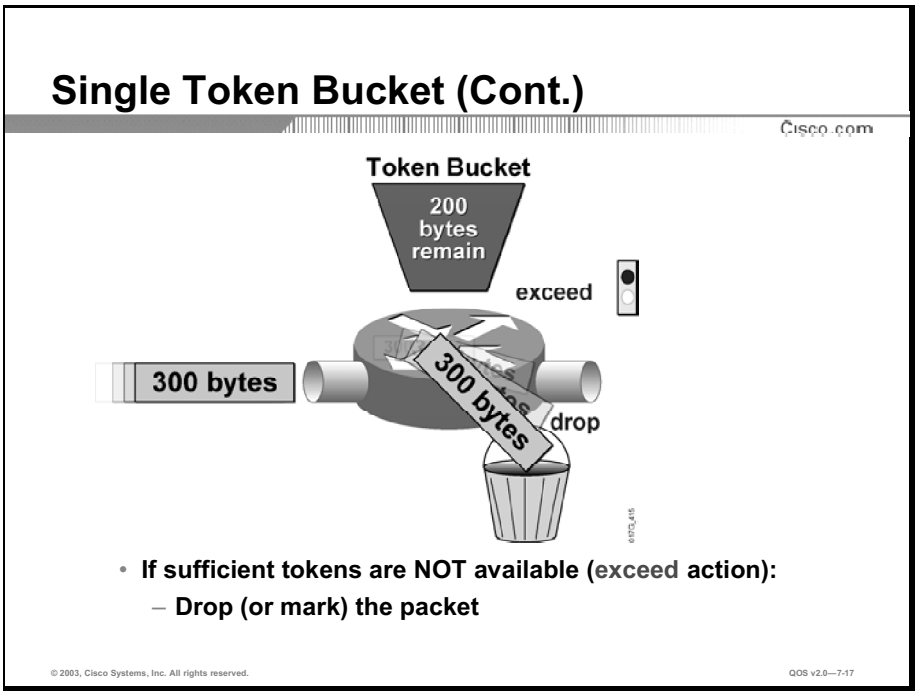

Continuing on with the single token bucket example from the previous slide, when the next 300-byte packet arrives immediately after the first packet, and no new tokens have been added to the bucket (which is done periodically), the packet exceeds the rate limit. The current packet size (300 bytes) is greater than the current capacity of the token bucket (200 bytes), and the exceed action is performed. The exceed action can be to drop or mark the packet when traffic policing.

Example: Token Bucket as a Piggy Bank

Think of a token bucket as a piggy bank. Every day you can insert a dollar into the piggy bank (the token bucket). At any given time, you can only spend what you have saved up in the piggy bank. On the average, if your saving rate is a dollar per day, your long-term average spending rate will be one dollar per day if you constantly spend what you saved. However, if you do not spend any money on a given day, then you can build up your savings in the piggy bank up to the maximum limit the piggy bank can hold. For example, if the size of the piggy bank is limited to hold five dollars and if you save and do not spend for five straight days, the piggy bank will contain five dollars. When the piggy bank fills to its capacity, you will not be able to put any more money in it. Then at any time, you can spend up to five dollars (bursting above the long-term average rate of one dollar per day).

Conforming rate using the piggy bank example means if you have two dollars in the piggy bank and you try to spend one dollar, then that is considered conforming because you are not spending more than what you have saved in the piggy bank.

Exceeding rate, using the piggy bank example, means that if you have two dollars in the piggy bank and you try to spend three dollars, it is considered exceeding because you are spending more than what you have saved in the piggy bank.

Single Token Bucket Class-Based Policing

This topic explains how traffic can be policed using a single token bucket scheme.

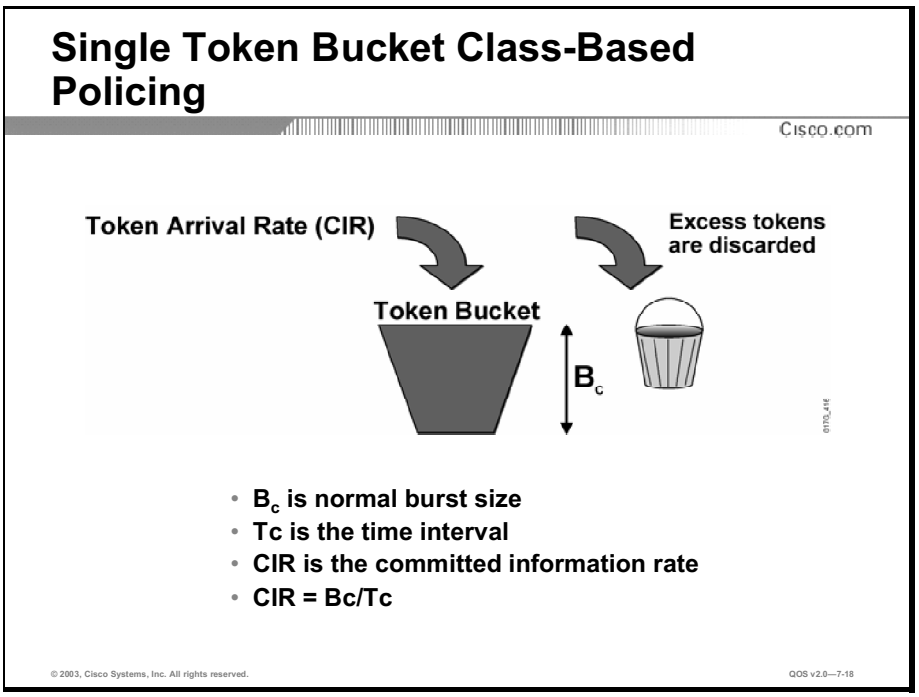

Token bucket operations rely on parameters such as: CIR, committed burst (Bc), and committed time window (Tc). CIR is the committed information rate. Bc is known as the normal burst size. Tc is an interval constant that represents time. The mathematical relationship between CIR, Bc, and Tc is:

- CIR (bps) = Bc (bits) / Tc (sec)

With traffic policing, new tokens are added into the token bucket based on the inter-packet arrival rate and the CIR. Every time a packet is policed, new tokens are added back into the token bucket. The amount of tokens added back into the token bucket is calculated as follows:

- (Current Packet Arrival Time – Previous Packet Arrival Time) * CIR

An amount (Bc) of tokens is forwarded without constraint in every time interval (Tc). For example, if 8000 bits (Bc) worth of tokens are placed in the bucket every 250 milliseconds (Tc), the router can steadily transmit 8000 bits every 250 milliseconds if traffic constantly arrives at the router.

- CIR (normal burst rate) = 8000 bits (Bc) / 0.25 seconds (Tc) = 32 kbps

When configuring Cisco IOS class-based traffic policing, it is recommended to allow the IOS software to automatically calculate the optimal Bc and Tc value based on the configured CIR.

Without any excess bursting capability, if the token bucket fills to capacity (Bc of tokens), the token bucket will overflow and newly arriving tokens are discarded. Using the example where the CIR is 32 kbps (Bc = 8000 bits and Tc = 0.25 seconds), the maximum traffic rate can never exceed a hard rate limit of 32 kbps.

Dual Token Bucket Class-Based Policing

This topic explains how traffic can be policed using a dual token bucket scheme.

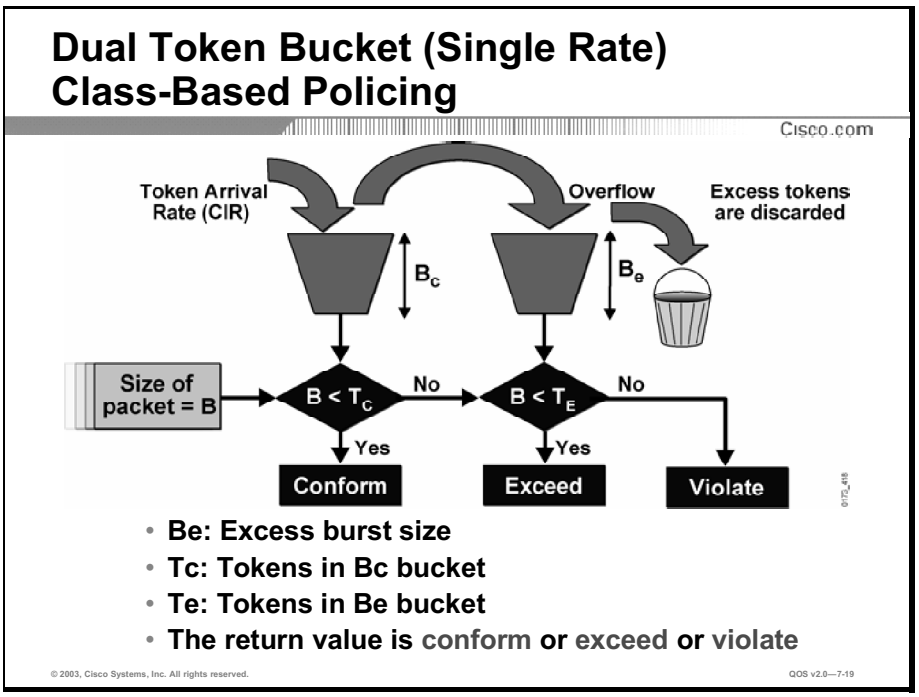

Class-based traffic policing can be configured to support excess bursting capability. With excess bursting, after the first token bucket is filled to Bc, extra (excess) tokens can be accumulated in a second token bucket. Excess burst (Be) is known as the excess burst size. Be is the maximum amount of excess traffic over and above Bc that can be sent during the time interval *after* a period of inactivity.

With a single rate metering mechanism, the second token bucket with a maximum size of Be fills at the same rate (CIR) as the first token bucket. If the second token bucket fills up to capacity, then no more tokens can be accumulated and the excess tokens are discarded.

When using a dual token bucket model, instead of a single token bucket, the measured traffic rate can be conforming, exceeding or violating:

- **Conforming:** There are enough tokens in the first token bucket with a maximum size of Bc.
- **Exceeding:** There are not enough tokens in the first token bucket but there are enough tokens in the second token bucket with a maximum size of Be.
- **Violating:** There are not enough tokens in the first or second token bucket.

With dual token bucket traffic policing, the typical actions performed can be:

- Send all conform traffic
- Re-mark (to a lower priority)
- Send all exceeding traffic
- Drop all violating traffic

The main benefit of using a dual token bucket method is the ability to distinguish between traffic that exceeds the Bc but not the Be. This enables a different policy to be applied to packets in the Be category.

Referring to the piggy bank example, think of the CIR as the savings rate (one dollar per day). Bc is how much you can save into the piggy bank per day (one dollar). Tc is the interval at which you put money into the piggy bank (one day). Be (five dollars) allows you to burst over the average spending rate of one dollar per day if you are not spending a dollar per day.

Dual Token Bucket (Single Rate) Class-Based Policing (Cont.)

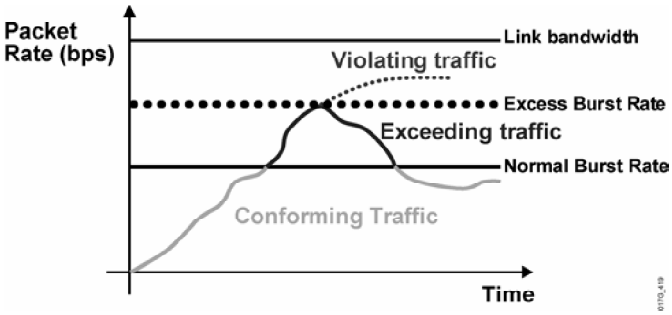

- **Traffic is conforming, exceeding, or violating**

Using a dual token bucket model allows traffic exceeding the normal burst rate (CIR) to be metered as exceeding and traffic that exceeds the excess burst rate to be metered as violating traffic. Different actions then can be applied to the conforming, exceeding, and violating traffic.

Dual-Rate Token Bucket Class-Based Policing

This topic explains how traffic can be policed using a dual-rate metering scheme.

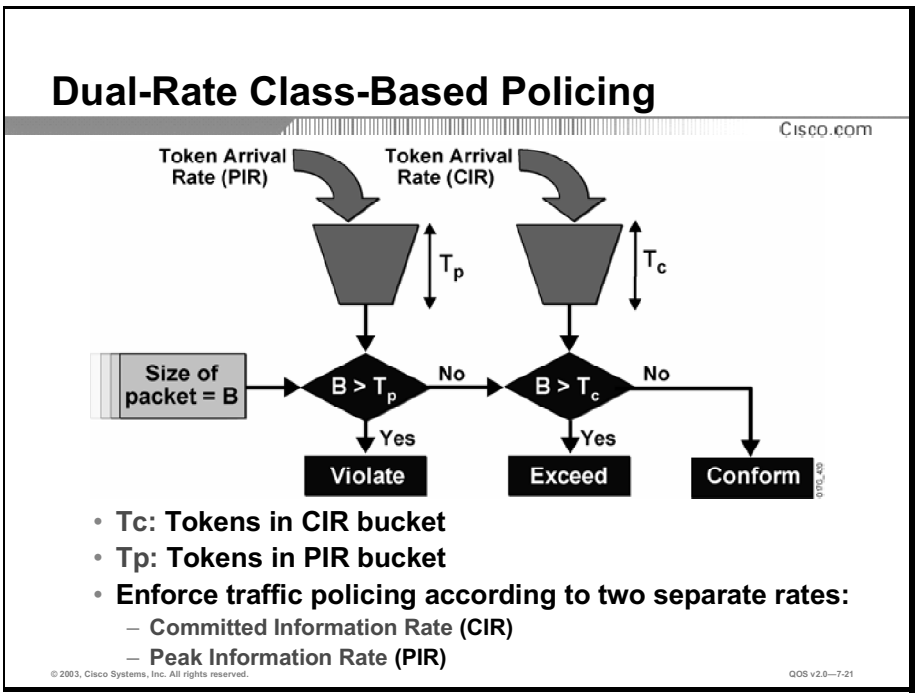

With dual-rate metering, traffic rate can be enforced according to two separate rates: CIR and peak information rate (PIR). Before this feature was available, you could meter traffic using a single rate based on the CIR with single or dual buckets. Dual-rate metering supports a higher level of bandwidth management and supports a sustained excess rate based on the PIR.

With dual-rate metering, the PIR token bucket fills at a rate based on the packet arrival rate, and the configured PIR and the CIR token bucket fills at a rate based on the packet arrival rate and the configured CIR.

When a packet arrives, the PIR token bucket is fist checked to see if there are enough tokens in the PIR token bucket to send the packet. The violating condition occurs if there are not enough tokens in the PIR token bucket to transmit the packet. If there are enough tokens in the PIR token bucket to send the packet, then the CIR token bucket is checked. The exceeding condition occurs if there are enough tokens in the PIR token bucket to transmit the packet but not enough tokens in the CIR token bucket to transmit the packet. The conforming condition occurs if there are enough tokens in the CIR bucket to transmit the packet.

Dual-rate metering is often configured on interfaces at the edge of a network to police the rate of traffic entering or leaving the network. In the most common configurations, traffic that conforms is sent and traffic that exceeds is sent with a decreased priority or is dropped. Users can change these configuration options to suit their network needs.

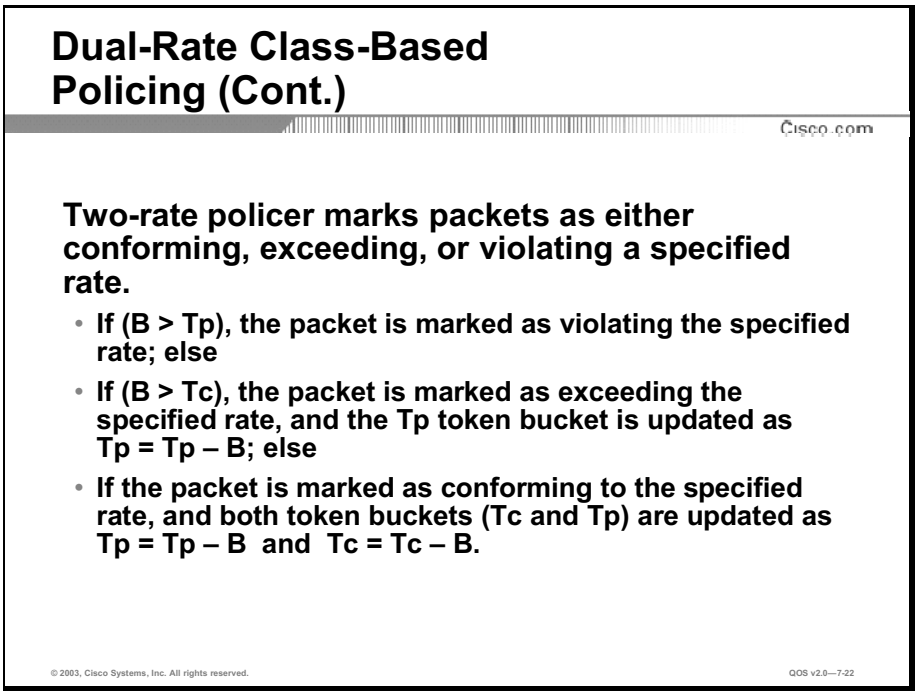

In addition to rate limiting, a traffic policing using dual-rate metering allows marking of traffic according to whether the packet conforms, exceeds, or violates a specified rate.

The token bucket algorithm provides users with three different actions for each packet: a conform action, an exceed action, and an optional violate action. Traffic entering the interface with two-rate policing configured is placed into one of these categories. Within these three categories, users can decide packet treatments. For example, a user may configure a policing policy as follows:

- Conforming packets are transmitted. Packets that exceed may be transmitted with a decreased priority; packets that violate are dropped.
- The violating condition occurs if there are not enough tokens in the PIR bucket to transmit the packet.
- The exceeding condition occurs if there are enough tokens in the PIR bucket to transmit the packet but not enough tokens in the CIR bucket to transmit the packet. In this case, the packet can be transmitted and the PIR bucket is updated to Tp – B remaining tokens where Tp is the size of the PIR bucket and B is the size of the packet to be transmitted.
- The conforming condition occurs if there are enough tokens in the CIR bucket to transmit the packet. In this case, the packets are transmitted and both buckets (Tc and Tp) are decremented to Tp – B and to Tc – B, respectively, where Tc is the size of the CIR bucket, Tp is the size of the PIR bucket, and B is the size of the packet to be transmitted.

Example: Dual-Rate Token Bucket as a Piggy Bank

Using a dual-rate token bucket is like using two piggy banks and each with a different savings rate. However, you can only take out money from one of the piggy banks at a time.

For example, you can save ten dollars per day into the first piggy bank (PIR = peak spending rate = $10 per day) and then at the same time, you can save five dollars per day into the second piggy bank (CIR = normal average spending rate = $5 per day). However, the maximum amount you can spend is $10 per day, not $15 per day, because you can only take out money from one piggy bank at a time.

In this example, after one day of savings, your first piggy bank (PIR bucket) will contain $10 and your second piggy bank (CIR bucket) will contain $5. The three different spending cases are examined here to show how dual-rate metering operates using the piggy banks example:

- **Case 1:** If you try to spend $11 at once, then you are violating (Tp<B), your peak-spending rate of $10 per day. In this case, you will not be allowed to spend the $11 because $11 is greater than the $10 you have in the first piggy bank (PIR bucket). Remember, you can only take out money from one of the piggy bank at a time.
- **Case 2:** If you try to spend $9 at once, then you are exceeding (Tp>B>Tc), your normal average spending rate of $5 per day. In this case, you will be allowed to spend the $9 and just the first piggy bank (PIR bucket) will be decremented to $10 - $9 or $1.

  After spending $9, the maximum amount you can continue to spend on that day is decremented to $1.
- **Case 3:** If you try to spend $4, then you are conforming (Tp>B and Tc>B) to your normal average spending rate of $5 per day. In this case, you will be allowed to spend the $4, and both piggy banks (PIR and CIR bucket) will be updated.

The first piggy bank (PIR bucket) will be updated to $10 - $4 = $6, and the second piggy bank (CIR bucket) will be updated to $5 - $4 = $1.

Both piggy banks are updated because after spending $4, the maximum amount you can continue to spend on that day is decremented to $6, and the normal spending rate for that same day is decremented to $1.

Therefore, after spending $4, the following will occur:

- If you continue to spend $7 on that same day, then you will be violating your peak-spending rate for that day. In this case, you will not be allowed to spend the $7 because $7 is greater than the $6 you have in the first piggy bank (PIR bucket).
- If you continue to spend $5 on that same day, then you will be exceeding your normal average spending rate for that day. In this case, you will be allowed to spend the $5 and just the first piggy bank (PIR bucket) will be decremented to $6 - $5, or $1.
- If you continue to spend 50 cents on that same day, then you will be conforming to your normal average spending rate for that day. In this case, you will be allowed to spend the 50 cents, and both piggy banks (PIR and CIR bucket) will be updated. The first piggy bank (PIR bucket) will be updated to $6 – $0.5 = $5.5 and the second piggy bank (CIR bucket) will be updated to $1 - $0.5 = $0.5.

Class-Based Traffic Shaping

This topic explains how traffic can be shaped using a single token bucket scheme.

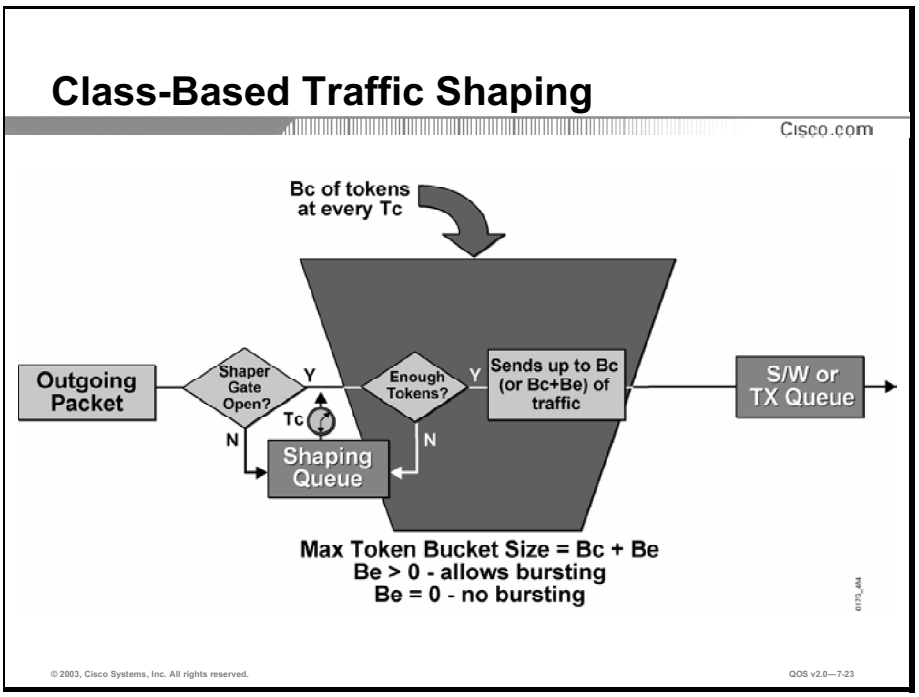

Cisco class-based traffic shaping only applies for outbound traffic.

Class-based traffic shaping uses the basic token bucket mechanism where Bc of tokens are added at every Tc time interval. The maximum size of the token bucket is Bc + Be. You can think of the traffic shaper operation like opening and closing of a transmit gate at every Tc interval. If the shaper gate is opened, the shaper checks to see if there are enough tokens in the token bucket to send the packet. If there are enough tokens, the packet is immediately forwarded. If there are *not* enough tokens, the packet is queued in the shaping queue until the next Tc interval. If the gate is closed, the packet is queued behind other packets in the shaping queue.

For example, on a 128-kbps link, if the CIR is 96 kbps, the Bc is 12 kbps, the Be is 0, and the Tc = 0.125 seconds, then during each Tc (125 ms) interval, the traffic shaper gate opens and up to 12 Kb can be sent. To send 12 Kb over a 128-kbps line, it will only take 91.25 ms. Therefore, this means the router will on the average be sending at three-quarters of the line rate (128 kbps * ¾ = 96 kbps).

Traffic shaping also includes the ability to send more than Bc of traffic in some time intervals after a period of inactivity. This extra number of bits in excess to the Bc is called Be.

Cisco IOS Traffic Policing and Shaping Mechanisms

This topic identifies the key traffic policing and shaping mechanisms available in Cisco IOS software and differentiates between them.

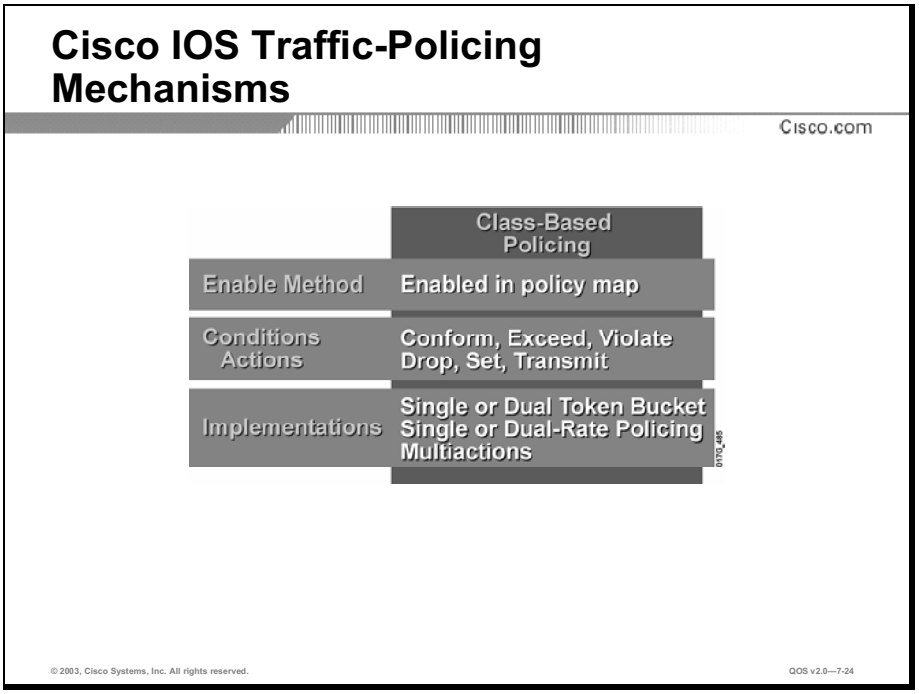

This figure lists the characteristics of the class-based traffic-policing mechanism that is available in Cisco IOS software. Class-based policing is also available on some Cisco Catalyst switches.

Class-based policing supports a single or dual token bucket. Class-based policing also supports single-rate or dual-rate metering and multiactions policing. Multiactions policing allows more than one action to be applied; for example marking the Frame Relay DE bit and also the DSCP value before sending the exceeding traffic.

Class-based policing is configured using Modular QoS command-line interface (CLI) (MQC) with the “police” command under the policy map.

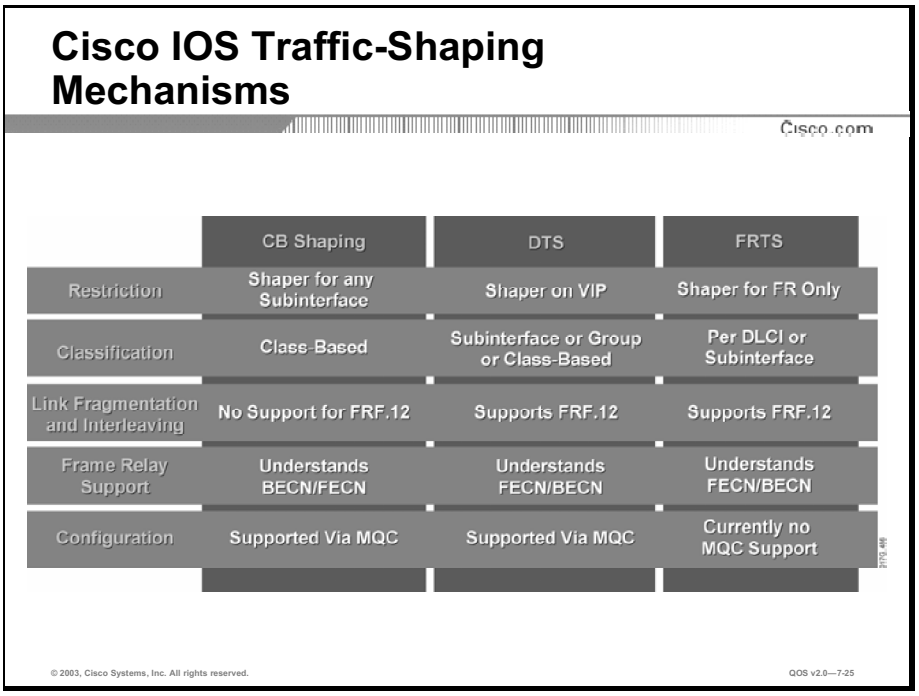

This figure lists some of the different traffic-shaping mechanisms available in Cisco IOS software, the newer class-based traffic shaping, distributed traffic shaping (DTS), and FRTS.

Class-based traffic shaping uses MQC to allow traffic to be shaped per traffic class as defined by the class map. Class-based traffic shaping can be used in combinations with class-based weighted fair queuing (CBWFQ), where the shaped rate is used to define an upper rate limit while the bandwidth statement within the CBWFQ configuration is used to define a minimum rate limit.

DTS is a feature that is specific to the higher-end platforms such as the Cisco 7500 or the Cisco 12000 series routers. These platforms have the ability to offload traffic shaping from the main processor to the individual interface processors (Versatile Interface Processor [VIP] or line card). In networks where distributed Cisco Express Forwarding (dCEF) is the preferred mode of switching, DTS on the VIP or line card is the logical choice for implementing traffic shaping.

FRTS is used to shape Frame Relay traffic only. FRTS allows individual PVC (data-link connection identifier [DLCI]) to be shaped. FRTS can use priority queuing (PQ), custom queuing (CQ) or weighted fair queuing (WFQ) as the shaping queue and only supports WFQ as the software queue.

Both DTS and FRTS support FRF.12 Frame Relay fragmentation while class-based shaping does not support FRF.12 fragmentation for Frame Relay.

All these traffic-shaping mechanisms can interact with a Frame Relay network, adapting to indications of Layer 2 congestion in the WAN. For example, if the BECN bit is received, the router can lower the rate limit to help reduce congestion in the Frame Relay network. And if the forward explicit congestion notification (FECN) bit is received, the router can generate a test frame with the BECN bit set. This enables the sender to notice congestion even if there is no data traffic flowing back from the receiver to the sender.

Only class-based shaping configurations will be discussed in this module.

Applying Traffic Conditioners

This topic identifies the points in a network where rate-limiting can most effectively be employed.

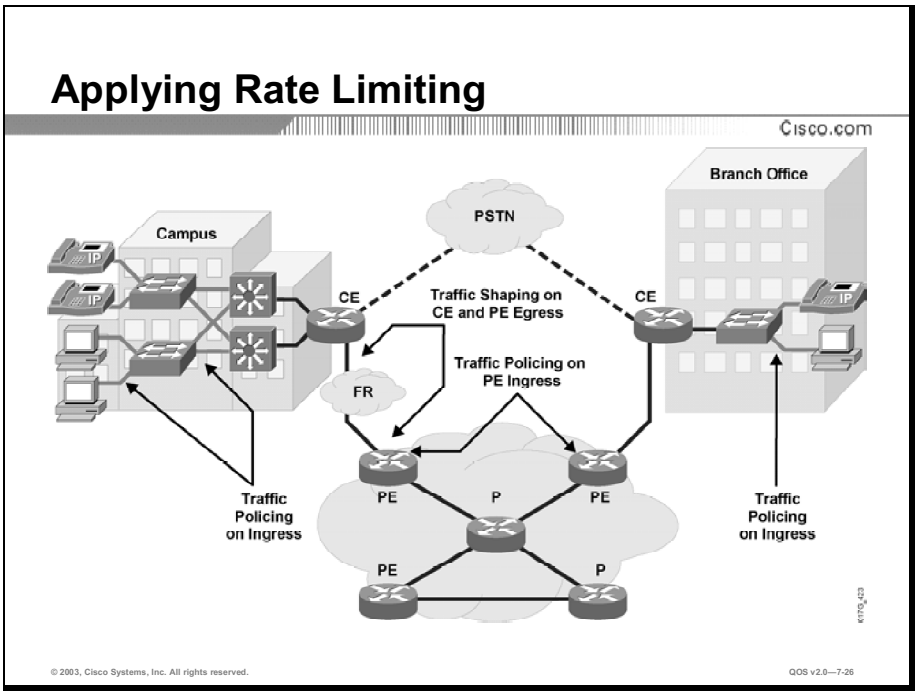

In a typical enterprise network, traffic policing is often implemented at the access or distribution layer to limit certain traffic classes before that traffic exits the campus onto the WAN. Traffic shaping is often implemented at the WAN edge when there are speed mismatches or over-subscription.

In a typical service provider network, traffic policing is often implemented inbound at the PE (provider edge) router to rate-limit incoming traffic from the CE (customer edge) router to ensure the customer traffic rate is not exceeding the contractual rate. Traffic shaping is often implemented outbound at the PE and at the CE to limit the traffic rate between the PE and CE and to allow for FRF.12 fragmentation on Frame Relay connections between the CE and PE.

Summary

This topic summarizes the key points discussed in this lesson.

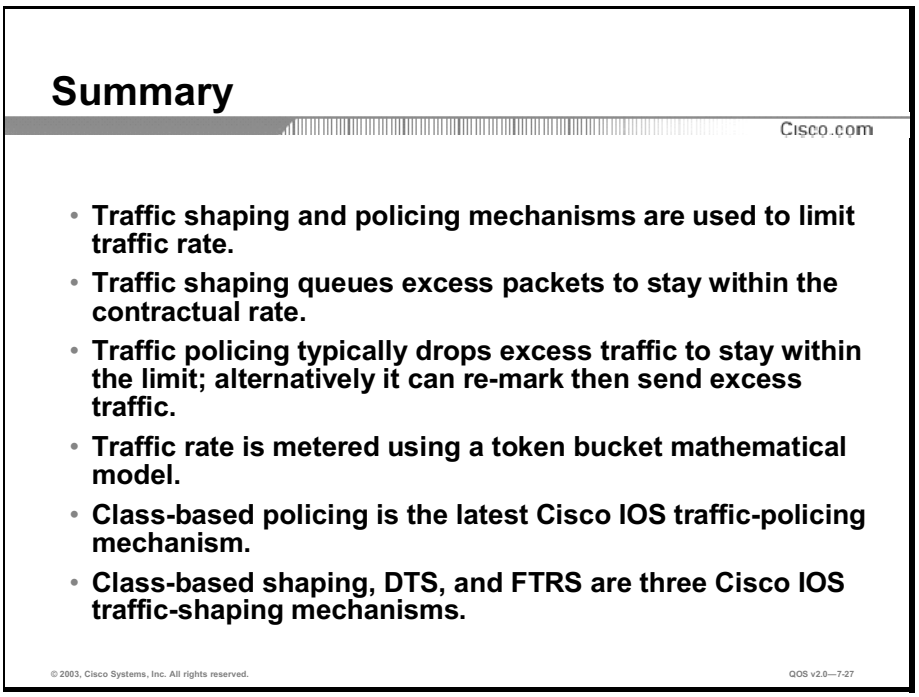

References

For additional information, refer to these resources:

- To learn more about traffic policing and traffic shaping, refer to “Part 4: Policing and Shaping” at the following URL: http://www.cisco.com/univercd/cc/td/doc/product/software/ios123/123cgcr/qos_vcg.htm#1001018
- For information on other traffic shaping mechanisms, refer to the software configuration documentation for your Cisco IOS software release.

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) Which two are common to both traffic policing and traffic shaping? (Choose two.)

A) both use token bucket to meter traffic

B) both use a queue to delay excess traffic

C) both use traffic classification to differentiate traffic

D) both drop all excess traffic

Q2) Traffic policing can be implemented in which directions?

A) Traffic policing can be implemented in both the inbound and outbound directions.

B) Traffic policing can be implemented in the inbound direction only.

C) Traffic policing can be implemented in the outbound direction only.

D) Traffic policing can be applied in the inbound direction for rate-limiting and in the outbound direction for marking purposes only.

Q3) Traffic shaping can be implemented in which directions?

A) Traffic shaping can be implemented in both the inbound and outbound directions.

B) Traffic shaping can be implemented in the inbound direction only.

C) Traffic shaping can be implemented in the outbound direction only.

D) Traffic shaping can be implemented in the outbound direction for rate-limiting and the inbound direction for marking purposes only.

Q4) What are the three possible conditions for the metered traffic when using dual token bucket policing? (Choose three.)

A) accept

B) conform

C) non-confirm

D) obey

E) exceed

F) violate

Q5) What is the mathematical formula between CIR, Bc, and Tc?

A) Tc = CIR * Bc

B) Tc = CIR / Bc

C) Bc = CIR / Tc

D) CIR = Bc / Tc

Q6) What is the main advantage of using a dual token bucket versus a single token bucket when implementing class-based policing?

A) the ability to burst above the Be

B) the ability to distinguish exceeding traffic that exceeds the normal burst rate versus the violating traffic that exceeds the excess burst rate and different actions then can be applied to the exceeding and violating traffic

C) the ability to specify a maximum token bucket size greater than Bc

D) the ability to queue the excess traffic in the second token bucket to reduce packet drops

Q7) Correctly identify each Cisco IOS traffic shaping and traffic-policing mechanism by matching it to the correct term below.

1. traffic shaping
2. traffic policing
3. neither

A) WRED _____

B) CB-Shaping _____

C) DTS _____

D) WFQ _____

E) FRTS _____

F) LLQ _____

G) CB-Policing _____

H) RED _____

I) PBR _____

Quiz Answer Key

Q1) A, C

Relates to: Traffic Policing and Shaping Overview

Q2) A

Relates to: Traffic Policing and Shaping Overview

Q3) C

Relates to: Traffic Policing and Shaping Overview

Q4) B, E, F

Relates to: Measuring Traffic Rates

Q5) D

Relates to: Single Token Bucket Class-Based Policing

Q6) B

Relates to: Dual Token Bucket Class-Based Policing

Q7) A = 3

B = 1

C = 1

D = 3

E = 1

F = 3

G = 2

H = 3

I = 3

Relates to: Cisco IOS Traffic Policing and Shaping Mechanisms

Configuring Class-Based Policing

Overview

Cisco IOS software supports two different traffic-policing mechanisms: CAR and class-based policing. CAR is an older Cisco traffic policing feature and class-based policing is a newer Cisco traffic-policing mechanism based on the MQC. Cisco recommends using MQC features when possible to implement QoS in the network. Traffic policing configurations using CAR, for which no new features or functionality is planned, should be avoided. However, Cisco will continue to support CAR for existing implementations.

This lesson describes the tasks to configure the different options that are used to implement class-based traffic policing to rate-limit certain traffic classes.

Relevance

Traffic policing is a valuable QoS tool to limit the rate at which traffic enters or exits an interface on a Cisco router or switch. Traffic policing using class-based policing is the preferred method of configuring policing.

Objectives

Upon completing this lesson, you will be able to configure class-based policing to rate-limit traffic. This includes being able to meet these objectives:

- Explain the key features of class-based policing
- Identify the Cisco IOS commands required to configure single-rate class-based policing
- Identify the Cisco IOS commands required to configure dual-rate class-based policing
- Identify the Cisco IOS commands required to configure percentage-based class-based policing
- Identify the Cisco IOS commands used to monitor class-based policing

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Knowledge of using MQC to implement Cisco QoS mechanisms
- Knowledge of how to configure class maps to classify traffic
- Knowledge of how traffic is metered using a token bucket
- Knowledge of the operation of a single token bucket versus a dual token bucket

Outline

The outline lists the topics included in this lesson.

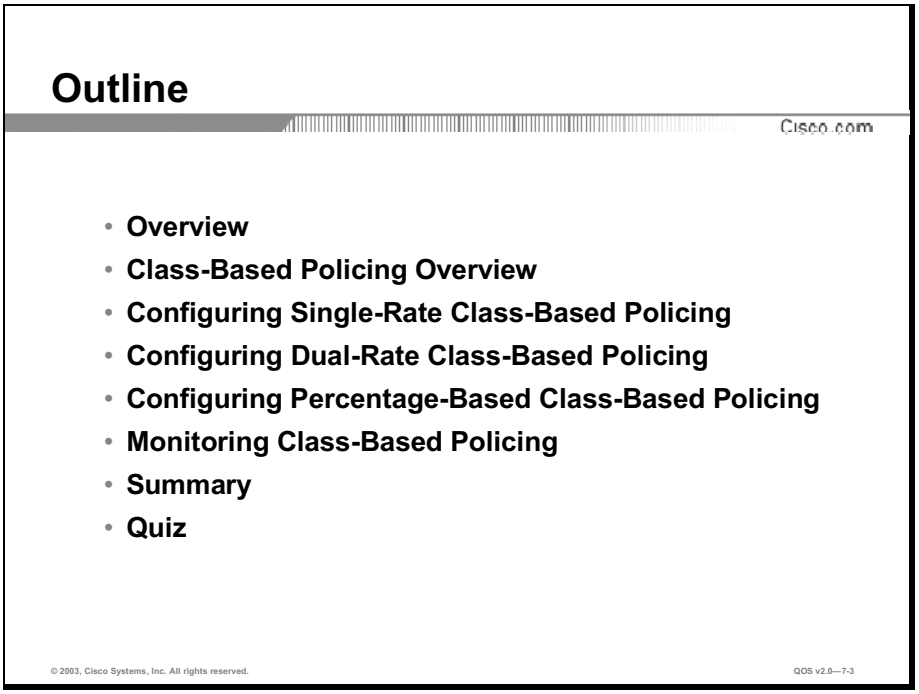

Class-Based Policing Overview

This topic explains the key features of class-based policing.

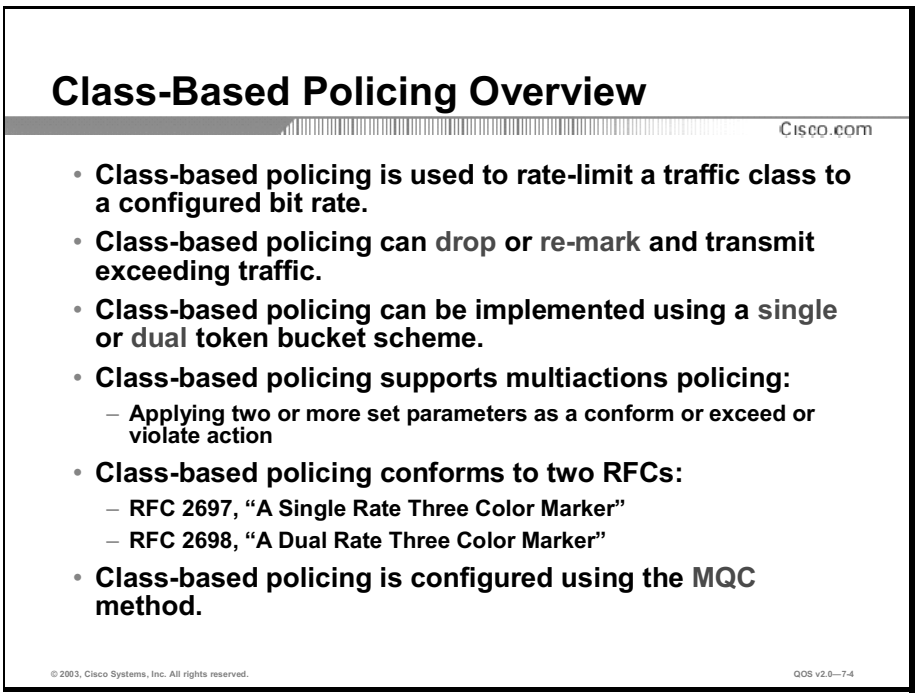

The class-based policing feature performs the following functions:

- Limits the input or output transmission rate of a class of traffic based on user-defined criteria
- Marks packets by setting different Layer 2 or Layer 3 markers or both.

Class-based policing can be implemented using a single or double token bucket method as the metering mechanism. The single token bucket algorithm is used when the violate action option is not specified in the police MQC command.

The dual token bucket algorithm is used when the violate action option is specified in the police MQC command.

Using a dual token bucket, traffic can:

- Conform to the rate limit when it is within the average bit rate
- Exceed the rate limit when it exceeds the average bit rate, but does not exceed the allowed excess burst
- Violate the rate limit when it exceeds both the average rate and the excess bursts

Based on the current packet conforming, exceeding, or violating the rate limit, one or more actions can be taken by class-based policing as follows:

- **Transmit:** The packet is transmitted.
- **Drop:** The packet is dropped.
- **Set precedence (or DSCP value) and transmit:** The IP precedence (type of service [ToS]) or DSCP bits in the packet header are rewritten. The packet is then transmitted. This action can be used to either color (set precedence) or recolor (modify existing packet precedence) the packet.
- **Set QoS group and transmit:** The QoS group can be set and the packet forwarded. Because QoS group is only locally significant within the router (that is, it is not transmitted outside the router), the QoS group setting is used in later QoS mechanisms and performed in the same router, such as CBWFQ, on an outgoing interface.
- **Set MPLS experimental bits and transmit:** The Multiprotocol Label Switching (MPLS) experimental bits can be set. The packet is then transmitted. These are usually used to signal QoS parameters in a MPLS cloud.
- **Set Frame Relay DE bit and transmit:** The Frame Relay DE bit is set in the Layer 2 (Frame Relay) header and the packet is transmitted. This setting can be used to mark excessive or violating traffic (which should be dropped with preference on Layer 2 switches) at the edge of a Frame Relay network.
- **Set ATM cell loss priority (CLP) bit and transmit:** The ATM CLP bit is set in the Layer 2 (ATM) header and the packet is transmitted. This setting can be used to mark excessive or violating traffic (which should be dropped with preference on Layer 2 switches) at the edge of an ATM network.

Multiaction policing is a mechanism that can apply more than one action to a packet; for example, setting the DSCP as well as the CLP bit on the exceeding packets.

Class-based policing also supports single- or dual-rate metering. With the two-rate policer, traffic policing can be enforced according to two separate rates: CIR and PIR. You can specify the use of these two rates, along with their corresponding values, by using two keywords, cir and pir, of the police command.

Cisco class-based policing mechanism conforms to the two following RFCs:

- **RFC 2697, A Single Rate Three Color Marker:** The Single Rate Three Color Marker meters an IP packet stream and marks its packets green (conform), yellow (exceed), or red (violate). Marking is based on a CIR) and two associated burst sizes, a Bc size and a Be size. A packet is marked green if it does not exceed the Bc, yellow if it does exceed the Bc, but not the Be, and red otherwise.
- **RFC 2698, A Two Rate Three Color Marker:** The Two Rate Three Color Marker meters an IP packet stream and marks its packets either green (conform), yellow (exceed), or red (violate). A packet is marked red if it exceeds the PIR. Otherwise it is marked either yellow or green, depending on whether it exceeds or does not exceed the CIR. This is useful, for example, for ingress policing of a service where a peak rate needs to be enforced separately from a committed rate.

Configuring Single-Rate Class-Based Policing

This topic identifies the Cisco IOS commands that are required to configure single-rate class-based policing.

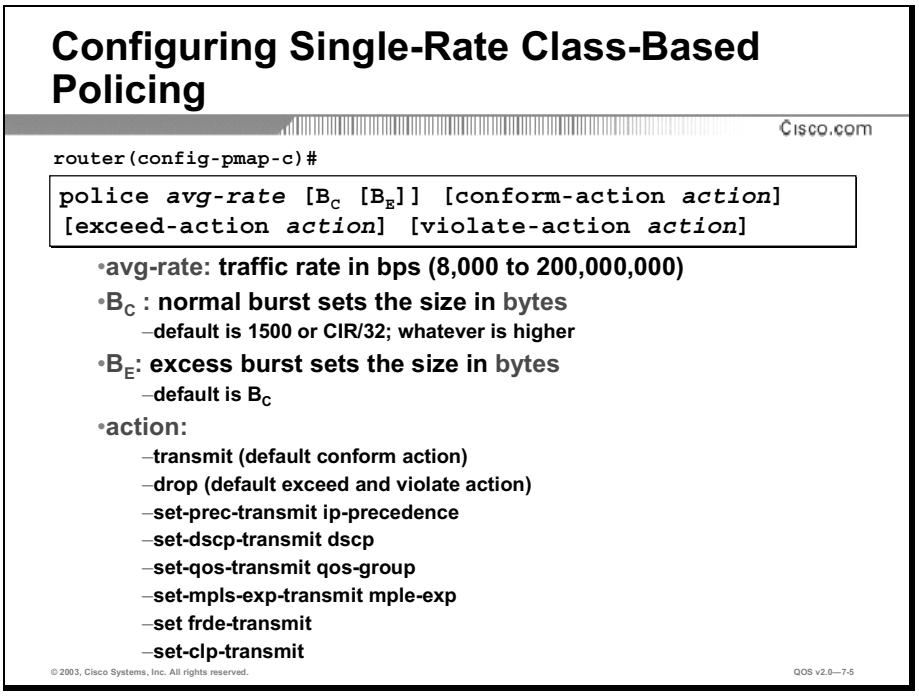

The MQC-based **police** command defines policing parameters for a specific traffic class. The **avg-rate** parameter defines the policed average traffic rate (CIR); Bc and Be define the token bucket sizes in bytes; and the action defines an action for conforming, exceeding, and optionally violating traffic.

If Bc (in bytes) is not specified, it will default to the avg-rate (CIR)/32 or 1500 bytes, whichever is higher. When using the formula CIR/32 to calculate the default Bc (in bytes), Cisco IOS software uses a Tc of 0.25 second where:

- Bc (in bytes) = (CIR x Tc) / 8
- Bc (in bytes) = (CIR x 0.25 seconds) / 8 = CIR / 32

If Be (in bytes) is not specified, it will default to Bc. In a single token bucket case, Cisco IOS software ignores the Be value. This means excess bursting is disabled.

The Be rate can be specified when a violate action is configured, therefore using a dual token bucket. This allows Be to be explicitly configured instead of using the default value of Be=Bc. Be specifies the size of the second (excess) token bucket.

Dual token bucket policing with the violate action was introduced in Cisco IOS 12.1(5)T.

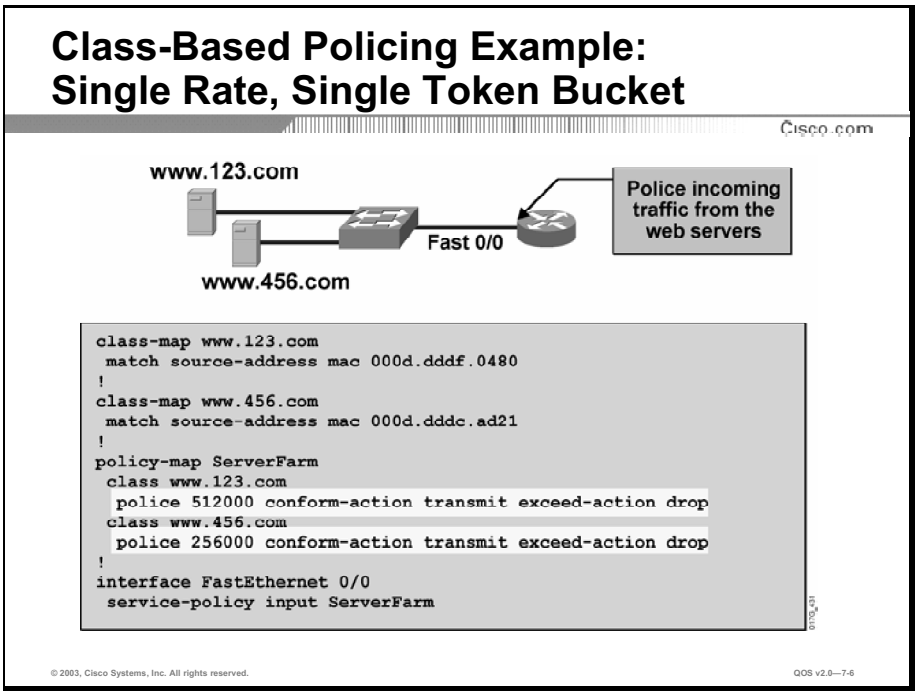

This class-based policing configuration example shows two configured traffic classes based on upstream MAC addresses. Traffic from the particular web server, which is classified by its MAC address, is policed to a fixed bandwidth with no excess burst capability using a single token bucket. Conforming traffic will be sent as-is and exceeding traffic is dropped. In this case, the www.123.com Web server is policed to a rate of 512 kbps and the www.456.com Web server is policed to a rate of 256 kbps.

Because the violate action is not specified, this will use a single token bucket scheme and no excess bursting is allowed.

In this example, the normal burst size (Bc) is not specified, and therefore it will default to the 512000/32 (16000 bytes) and 256000/32 (8000 bytes), respectively.

The default Bc setting can be examined by showing the policy map. Notice that the Be is not displayed because no excess bursting is allowed using a single token bucket with class-based policing:

```
router#show policy-map ServerFarm
  Policy Map ServerFarm
    Class www.123.com
     police cir 512000 bc 16000
       conform-action transmit
       exceed-action drop
    Class www.456.com
     police cir 256000 bc 8000
       conform-action transmit
       exceed-action drop
```

Example: Single Rate, Dual Token Bucket Class-Based Policing

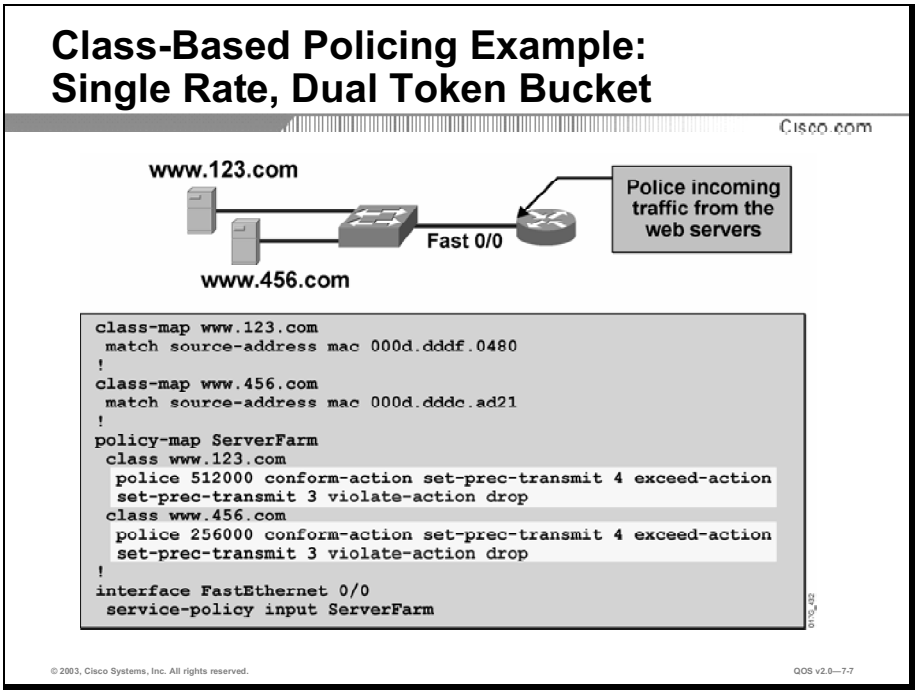

This class-based policing configuration example shows two configured traffic classes based on upstream MAC addresses. Traffic from the particular web server, which is classified by its MAC address, is policed to a fixed bandwidth with excess burst capability using a dual token bucket by configuring a violate action. Conforming traffic will be sent as-is, exceeding traffic will be marked to IP precedence 3 and transmitted, and all violating traffic will be dropped.

In this example because the violate action is specified, this will use a dual token bucket scheme with excess bursting. The normal burst size (Bc) is not specified, and therefore it will default to the 512000/32 (16000 bytes) and 256000/32 (8000 bytes), respectively. The excess burst size (Be) is also not specified, and therefore it will default to Bc.

The default Bc and Be settings can be examined by showing the policy map:

```
router#show policy-map ServerFarm
  Policy Map ServerFarm
    Class www.123.com
     police cir 512000 bc 16000 be 16000
       conform-action transmit
       exceed-action set-prec-transmit 3
       violate-action drop
    Class www.456.com
     police cir 256000 bc 8000 be 8000
       conform-action transmit
       exceed-action set-prec-transmit 3
       violate-action drop
```


Example: Multiactions Class-Based Policing

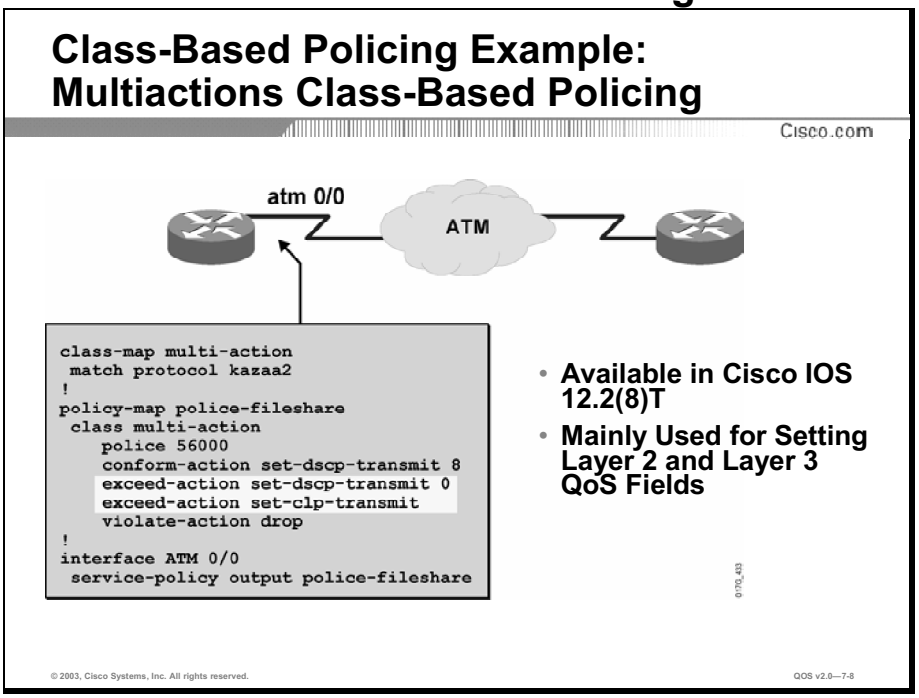

This class-based policing configuration is an example of a multiactions class-based policing. In this case, all Kazaa2 traffic is policed to 56 kbps. All conforming traffic will be marked with a DSCP value of 8, and then transmitted. All exceeding traffic will be marked with a DSCP value of 0, and the CLP bit in the ATM header will also be set before it is transmitted. All violating traffic will be dropped.

The multiactions feature was introduced in Cisco IOS 12.2(8)T and is primarily used so that class-based policing can mark at Layer 2 and also at Layer 3 at the same time.

Configuring Dual-Rate Class-Based Policing

This topic identifies the Cisco IOS commands that are required to configure dual-rate class-based policing.

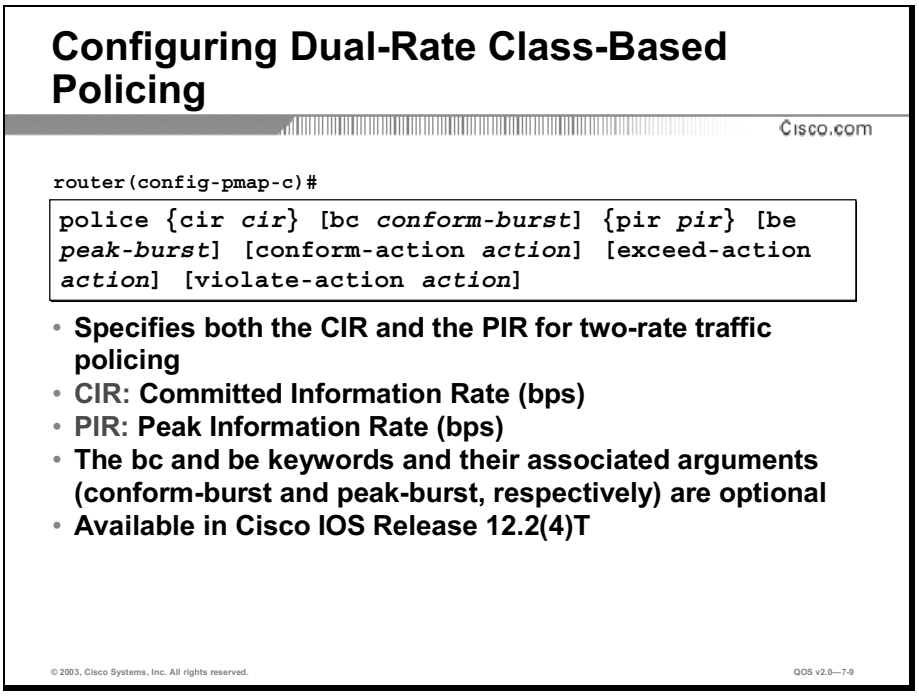

Dual-rate class-based policing was introduced in Cisco IOS 12.2(4)T. With dual-rate policing, traffic policing can be enforced according to two separate rates: CIR and PIR. The use of these two rates can be specified, along with their corresponding values, by using two keywords, cir and pir, of the police command. The Bc and Be keywords and their associated arguments (conform-burst and peak-burst, respectively) are optional. If Bc is not specified, Bc (in bytes) will default to CIR/32 or 1500 bytes, whichever is higher. If Be is not specified, Be (in bytes) will default to PIR/32 or 1500 bytes, whichever is higher.

Example: Dual-Rate Class-Based Policing

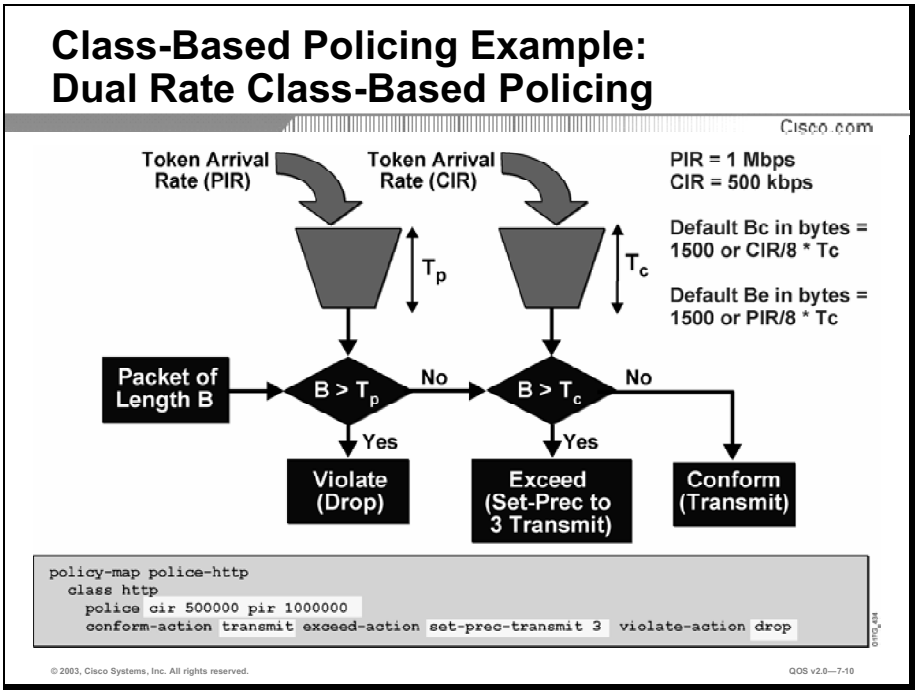

In this example, the two-rate policer is configured on the “http” traffic class to limit http traffic to an average committed rate of 500 kbps and a peak rate of 1 Mbps. Conforming traffic will be sent as-is, traffic exceeding 500 kbps (but not exceeding 1 Mbps) will be marked with IP precedence 3 and then sent, and all traffic exceeding 1 Mbps will be dropped.

Because the Bc and Be is not specified, the default Bc in bytes will be CIR/32 (Tc = 0.25 seconds) or 1500 bytes, whichever is higher. The default Be in bytes will be PIR/32 (Tc = 0.25 seconds) or 1500 bytes, whichever is higher.

In this example, the default Bc = 500000/32 = 15625 bytes and the default Be = 1000000/32 = 31250 bytes.

The **show policy-map** command can be used to display the default settings of the Bc and Be parameters:

```
Router #show policy-map police-http
  Policy Map police-http
    Class silver
     police cir 500000 bc 15625 pir 1000000 be 31250
       conform-action transmit
       exceed-action set-prec-transmit 3
       violate-action drop
```

Configuring Percentage-Based Class-Based Policing

This topic identifies the Cisco IOS commands that are required to configure percentage-based class-based policing.

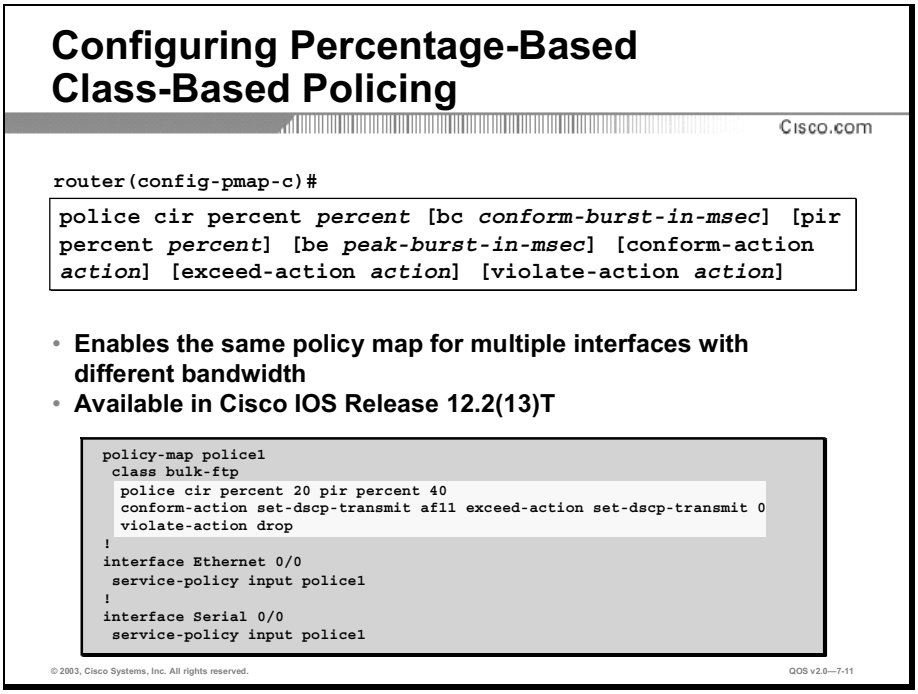

The percent policing feature was introduced in Cisco IOS 12.2(13)T. Before this feature, traffic policing was configured on the basis of a user-specified amount of bandwidth available on the interface. Policy maps were then configured on the basis of that specific amount of bandwidth, meaning that separate policy maps were required for each interface with different bandwidth. This feature provides the ability to configure traffic policing based on a percentage of bandwidth available on the interface. Configuring traffic policing and traffic shaping in this manner enables customers to use the same policy map for multiple interfaces with differing amounts of bandwidth.

Configuring traffic policing based on a percentage of bandwidth is accomplished by using the **police** (percent) command.

This feature also provides the option of specifying burst sizes in ms instead of in bytes when configuring traffic policing based on a percentage of bandwidth.

In this dual-rate percentage-based policing example, the bulk-ftp traffic class is policed to a CIR of 20 percent and a PIR of 40 percent of the interface bandwidth.

Since Bc and Be are not specified, the default Bc in bytes will be CIR/32 or 1500 bytes, whichever is higher. The default Be in bytes will be PIR/32 or 1500 bytes, whichever is higher.

Example: Configuring Percentage-Based Class-Based Policing

In the example, the CIR is 20 percent and the PIR is 40 percent (with no Bc and Be specified) and the policy is applied to an Ethernet (10 Mbps) interface. The CIR and PIR in bps for the Ethernet interface is computed as:

- CIR = 10 Mbps * 0.20 = 2 Mbps
- PIR = 10 Mbps * 0.40 = 4 Mbps

The default values of the Bc and Be parameters will automatically set to:

- Bc = 2 Mbps/32 = 62500 bytes
- Be = 4 Mbps/32 = 125000 bytes

Monitoring Class-Based Policing

This topic identifies the Cisco IOS commands that are used to monitor class-based policing.

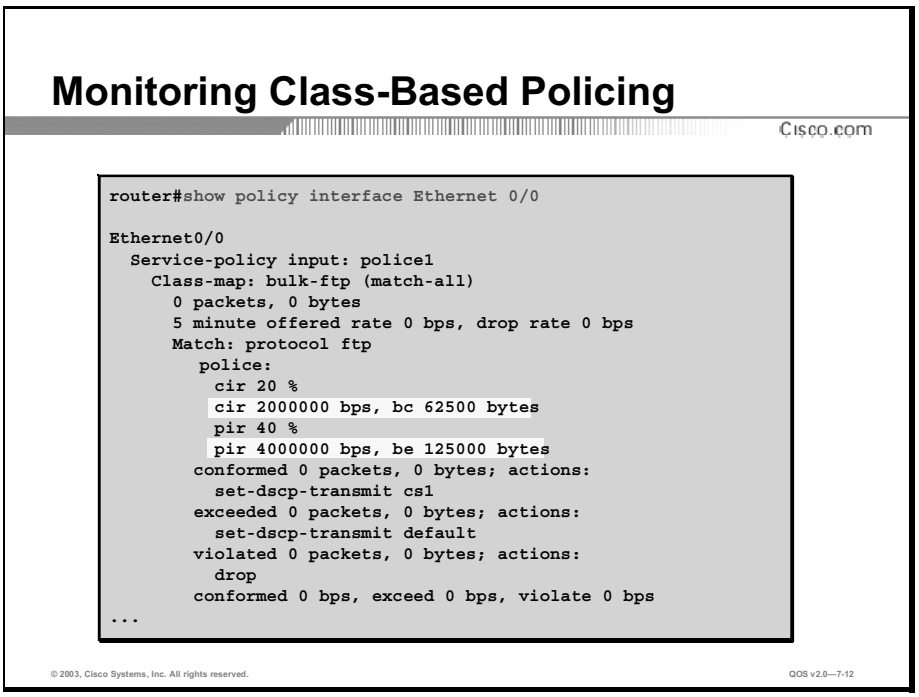

The **show policy-map interface** command is used to examine the policy map that is being applied to an interface.

This show output displays the percent-based policing policy map example shown on the previous slide.

This show output displays the actual Cisco IOS software calculated value of CIR and PIR in bps and the Bc and Be values in bytes, based on the interface bandwidth and the percentage configurations.

Summary

This topic summarizes the key points discussed in this lesson.

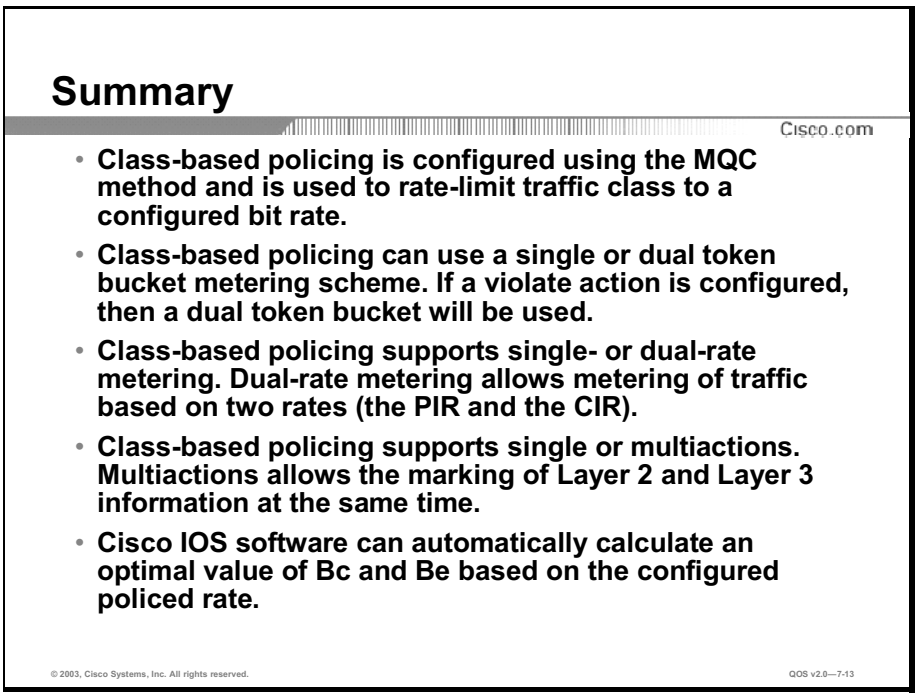

References

For additional information, refer to these resources:

- To learn more about configuring class-based policing, refer to “Configuring Traffic Policing” at the following URL: http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122cgcr/fqos_c/fqcprt4/qcfpoli.htm
- To learn more about metering in class-based policing, refer to “RFC 2697: A Single Rate Three Color Marker” at the following URL: http://www.faqs.org/rfcs/rfc2697.html
- To learn more about metering in class-based policing, refer to “RFC 2698: A Two Rate Three Color Marker” at the following URL: http://www.faqs.org/rfcs/rfc2698.html

Next Steps

For the associated lab exercise, refer to the following section of the course Lab Guide:

- Lab Exercise 7-1: Configuring Class-Based Policing

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) Which parameter in the police MQC command will enable the use of a dual token bucket?

A) Be

B) Bc

C) Exceed action

D) Violate action

Q2) What is the default value of Be if the Be parameter is not specified in the police MQC command when configuring single-rate, dual token bucket policing?

A) The excess burst rate (Be) will default equal to Bc.

B) The excess burst rate (Be) will default equal to 0.

C) The excess burst rate (Be) will default equal to 1500 bytes.

D) The excess burst rate (Be) will default equal to 8000 bytes.

Q3) What is the benefit of using multiaction policing?

A) allows policing to define different conforming, exceeding, and violating actions

B) to allow the use of dual token buckets to distinguish between exceeding and violating traffic

C) allows policing to mark at Layer 2 and also at Layer 3 at the same time

D) allows policing to rate-limit traffic using two rates (PIR and CIR)

Q4) Given the following police command, what will be the default values of the Bc?

```
police cir 250000 pir 500000
```

A) 1500 bytes

B) 7812 bytes

C) 15655 bytes

D) 31250 bytes

Q5) Which of the following statements is true regarding dual-rate class-based policing?

A) One token bucket fills at the PIR and has a maximum size of Be. The other token bucket fills at the CIR and has a maximum size of Bc.

B) Both token buckets fill at a rate equal to CIR.

C) Both token buckets fill at a rate equal to PIR.

D) The PIR will default to the CIR if the PIR is not specified in the police command.

E) The PIR will default to the CIR * 2 if the PIR is not specified in the police command.

Q6) If the same class-based policing configuration is implemented on different interfaces with differing amounts of bandwidth, which class-based configuration should be used?

A) dual-rate class-based policing

B) dual token bucket class-based policing

C) multiactions class-based policing

D) percentage-based class-based policing

Quiz Answer Key

Q1) D

Relates to: Configuring Single-Rate Class-Based Policing

Q2) A

Relates to: Configuring Single-Rate Class-Based Policing

Q3) C

Relates to: Configuring Single-Rate Class-Based Policing

Q4) B

Relates to: Configuring Dual-Rate Class-Based Policing

Q5) A

Relates to: Configuring Dual-Rate Class-Based Policing

Q6) D

Relates to: Configuring Percentage-Based Class-Based Policing

Configuring Class-Based Shaping

Overview

Traffic shaping allows network administrators to control outgoing traffic on an interface to match the speed of transmission to the speed of the remote interface, and to ensure that the traffic conforms to administrative QoS policies. Traffic adhering to a particular profile can be shaped to meet downstream requirements, thereby eliminating bottlenecks in topologies with data-rate mismatches.

This lesson describes the tasks to configure class-based traffic shaping to rate-limit certain traffic classes.

Relevance

The primary reasons to use traffic shaping are to control access to available bandwidth, to ensure that traffic conforms to the policies established for it, and to regulate the flow of traffic so that you can avoid congestion that can occur when the sent traffic exceeds the access speed of its remote, target interface. Traffic shaping is especially important in Frame Relay networks because Frame Relay switches cannot determine which frames take precedence, and therefore which frames should be dropped when congestion occurs.

Objectives

Upon completing this lesson, you will be able to configure class-based shaping to rate-limit traffic. This includes being able to meet these objectives:

- Explain the key features of class-based shaping
- Explain how the two rate limits—average rate and peak rate—can be used to rate-limit traffic
- Identify the Cisco IOS commands required to configure class-based shaping
- Identify the Cisco IOS commands used to monitor class-based shaping

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Knowledge of the MQC configuration method to implement Cisco QoS mechanisms
- Knowledge of how to configure class maps to classify traffic
- Knowledge of how to configure CBWFQ
- Knowledge of how traffic is metered using a token bucket

Outline

The outline lists the topics included in this lesson.

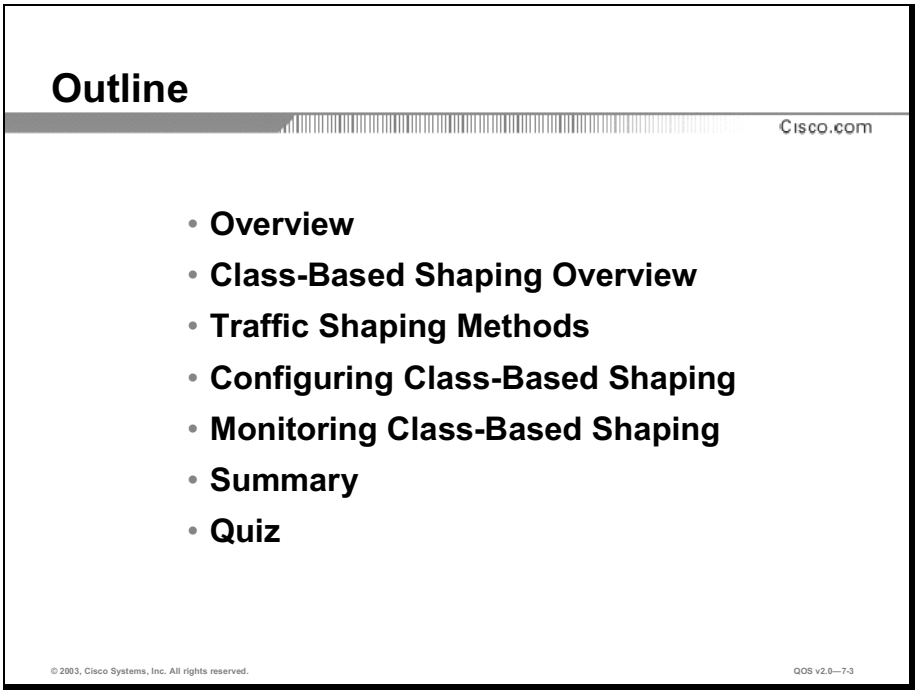

Class-Based Shaping Overview

This topic explains the key features of class-based shaping.

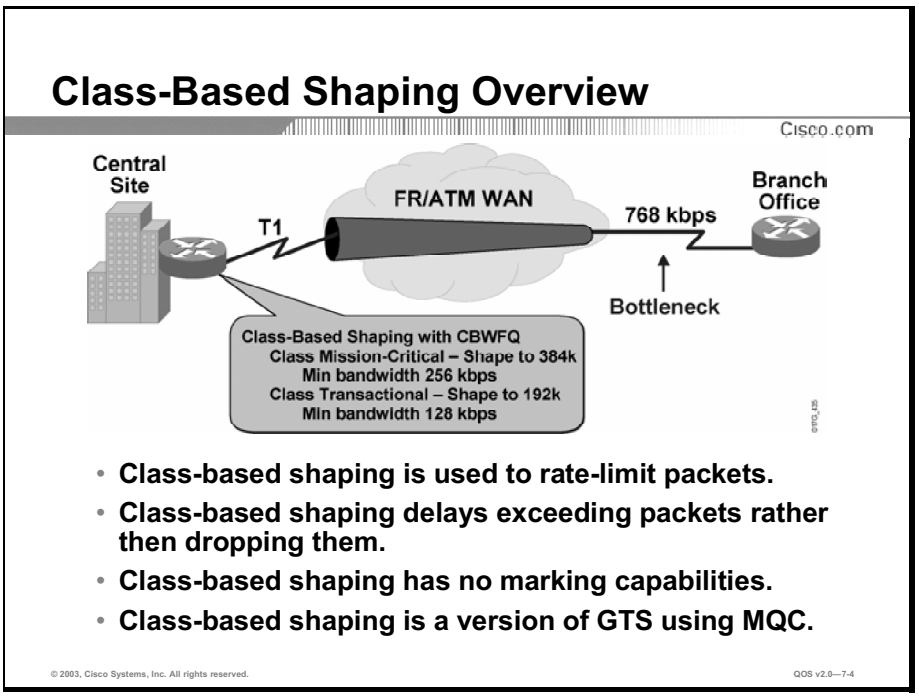

Traffic shaping allows you to control the traffic going out an interface to match its transmission to the speed of the remote, target interface or to ensure that the traffic conforms to policies contracted for it. Traffic adhering to a particular profile can be shaped to meet downstream requirements, thereby eliminating bottlenecks in topologies with traffic-rate mismatches or over-subscriptions. Class-based shaping has these properties:

- Class-based shaping is configured via MQC.
- Class-based shaping has no packet-marking capability.
- Class-based shaping works by queuing excessive packets until they conform to the configured shaped rate.
- Class-based shaping can also be used within the CBWFQ queuing system.
  - When operating within CBWFQ, the shaped rate provides an upper bandwidth limit while the bandwidth statement within CBWFQ provides a minimum bandwidth guarantee.

Traffic Shaping Methods

This topic explains how the two rate limits—average rate and peak rate—can be used to rate-limit traffic.

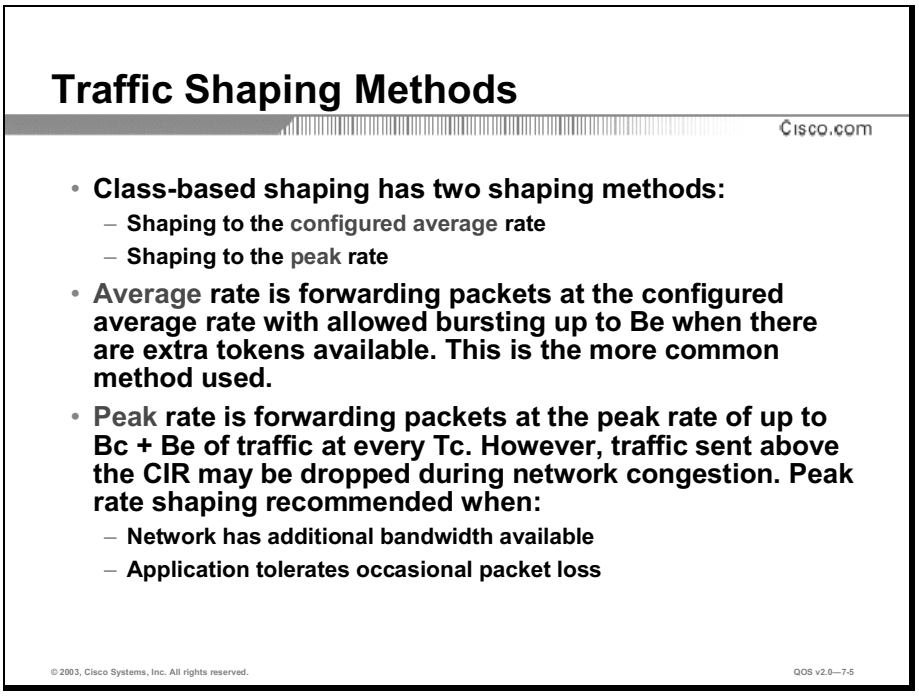

Class-based shaping can be configured in two different ways:

- Shaping to the average rate, which forwards up to Bc of traffic at every Tc time interval with additional bursting capability when enough tokens are accumulated in the bucket. Bc of tokens are added to the token bucket at every Tc time interval. After the token bucket is drained, additional bursting cannot occur until tokens are allowed to accumulate, which can occur only during periods of silence or when transmit rate is lower than average rate. After a period of low traffic activity, up to Bc + Be of traffic can be sent.
- Shaping to the peak rate, which forwards up to Bc + Be of traffic at every Tc time interval. Bc + Be of tokens are added to the token bucket at every Tc time interval. Shaping to the peak rate sends traffic at the peak rate, which is defined as the average rate multiplied by (1 + Be/Bc). Sending packets at the peak rate may result in dropping in the WAN cloud during network congestion. Shaping to the peak rate is only recommended when the network has additional available bandwidth beyond the CIR and applications can tolerate occasional packet drops.

Average rate shaping is the more common approach being deployed.

Configuring Class-Based Shaping

This topic identifies the Cisco IOS commands that are required to configure class-based shaping.

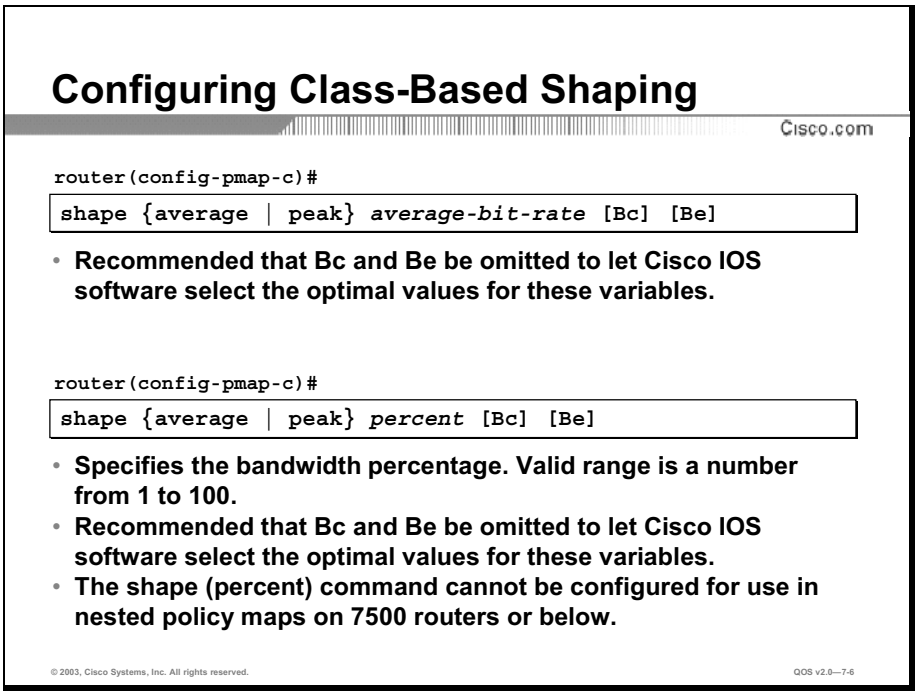

The **shape average** and **shape peak** commands configure average and peak shaping, respectively. The Bc and Be value in bits can be explicitly configured or Cisco IOS software can automatically calculate the optimal value for Bc and Be. It is recommended not to configure the Bc and Be and let the Cisco IOS algorithm determine the best Bc and Be value to use. Class-based traffic shaping uses a single token bucket with a maximum token bucket size of Bc + Be.

The **shape percent** command is often used in conjunction with the **bandwidth** and **priority** commands. The **bandwidth** and **priority** commands can be used to calculate the total amount of bandwidth available on an entity (for example, a physical interface). When the **bandwidth** and **priority** commands calculate the total amount of bandwidth available on an entity, the total bandwidth is the bandwidth on the physical interface.

The **shape max-buffers** command (not shown) specifies the maximum size of the shaping queue. The shaping queue will queue excess packets until they conform to the shaping rate. If the shaping queue is full, packets are tail-dropped. The default max-buffer size is 1000 packets and is usually not changed.

Example: Average Rate, Peak Rate

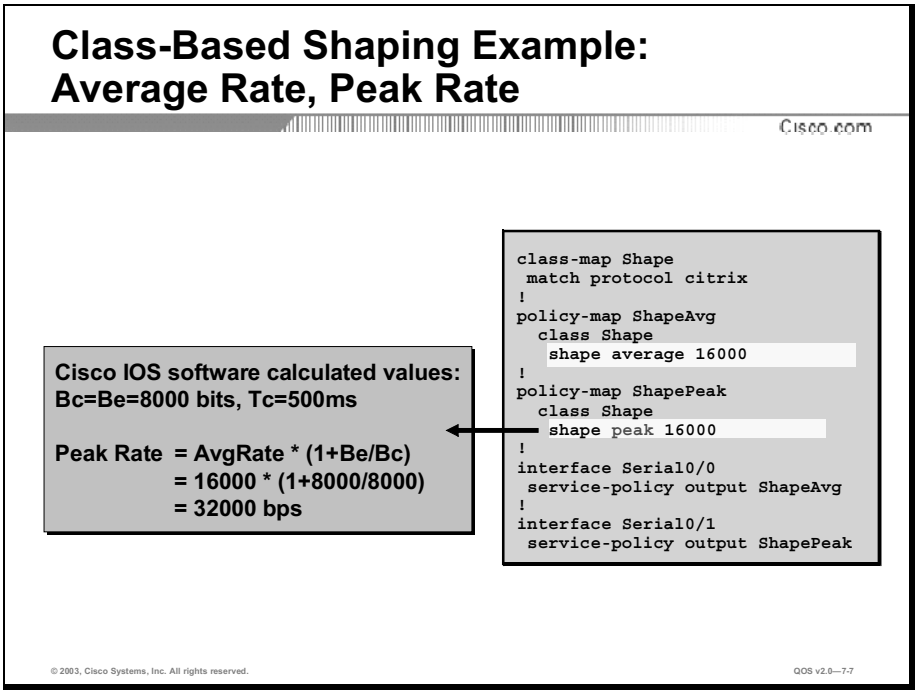

This figure shows an example configuration for standalone class-based shaping (no CBWFQ). Citrix traffic is classified into the “Shape” class. The “Shape” class is then shaped to different rates on two interfaces:

- On the Serial 0/0 interface, traffic is shaped to the average rate, while on the Serial 0/1 interface, traffic is shaped to the peak rate. In both cases, the Bc and Be values are not configured, allowing Cisco IOS software to automatically calculate the optimal values.
  - On the Serial 0/0 interface, the Citrix traffic is shaped to the average rate of 16000 bps. The resulting automatically determined Be and Bc values will both be 8000 bits with a Tc of 500 ms.
- On the Serial 0/1 interface, the Citrix traffic is shaped to the peak rate. Using this peak rate shaping example, because the Be and Bc were not specified, Cisco IOS software will automatically calculate the optimal value for Tc, Bc, and Be. The shape statement is “shape peak 16000”. The resulting automatically determined Be and Bc values will both be 8000 bits with a Tc of 500 ms, therefore, the peak rate will be:

```
peak rate   = ave rate * (1+Be/Bc) = 16000 * (1+8000/8000)
= 16000 * 2 = 32000 bps
```


Showing the policy map on the interface will display the average rate, peak rate, Bc, Be, and Tc values:

```
router#show policy-map interface s0/1
 Serial0/0

  Service-policy output: ShapePeak

    Class-map: Shape (match-all)
      0 packets, 0 bytes
      5 minute offered rate 0 bps, drop rate 0 bps
      Match: protocol citrix
      Traffic Shaping
           Target/Average  Byte   Sustain   Excess    Interval
Increment
             Rate          Limit  bits/int  bits/int  (ms)
(bytes)
            32000/16000     2000   8000      8000      500
2000

        Adapt   Queue   Packets   Bytes     Packets   Bytes
Shaping
        Active  Depth                       Delayed   Delayed
Active
        -          0    0         0         0         0
no

    Class-map: class-default (match-any)
      0 packets, 0 bytes
      5 minute offered rate 0 bps, drop rate 0 bps
      Match: any
```

Example: Class-Based Shaping with CBWFQ

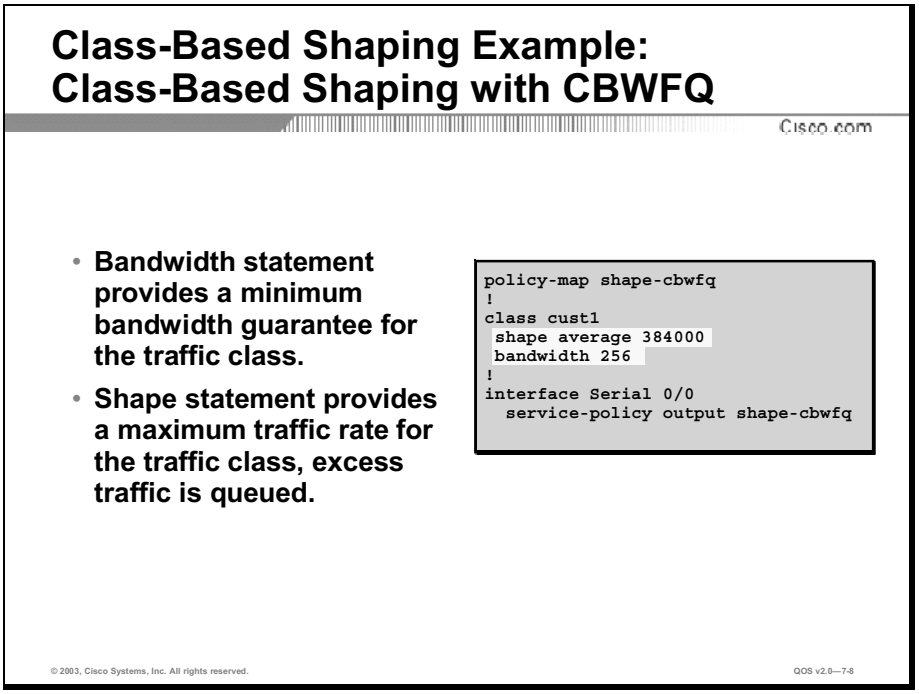

Class-based shaping can be used in combination with CBWFQ. The shape rate provides a maximum rate limit for the traffic class while the bandwidth statement within CBWFQ provides a minimum bandwidth guarantee.

In this example, the “cust1” traffic class is shaped to an average rate of 384 kbps while it is guaranteed a minimum of 256 kbps by the CBWFQ bandwidth statement. Because the Bc and Be are not specified it will be automatically calculated by the Cisco IOS software.

Example: Class-Based Shaping Hierarchical Policy Maps

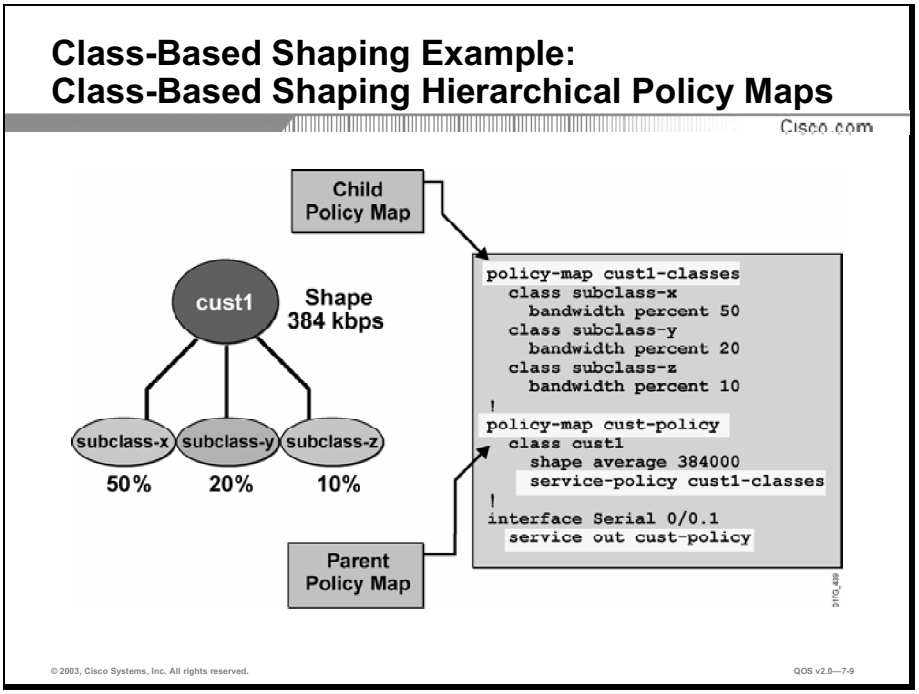

This example uses hierarchical policy maps and configures CBWFQ inside the class-based shaping.

The parent policy is the cust-policy. This parent policy references a child policy named cust1-classes.

The parent policy map cust-policy specifies an average shape rate of 384 kbps for all the "cust1" traffic and assigns the service policy called cust1-classes as the child policy.

Within the "cust1" traffic class, the "cust1" traffic is further classified into three distinct traffic subclasses:

- **Subclass-x:** The subclass-x class is configured to have a minimum guarantee equal to 50 percent of the shaped bandwidth.
- **Subclass-y:** The subclass-y class is configured to have a minimum guarantee equal to 20 percent of the shaped bandwidth.
- **Subclass-z:** The subclass-z class is configured to have a minimum guarantee equal to 10 percent of the shaped bandwidth.

For example, in a service provider implementation, the "cust1" traffic can be all traffic going to a specific customer edge router on a Frame Relay subinterface. The traffic to that specific customer edge router is further classified into:

- Subclass-x class for interactive traffic
- Subclass-y class for business-critical traffic
- Subclass-z class for bulk traffic

The traffic going out to that specific customer edge router over the Frame Relay subinterface is rate-limited to 384 kbps total. The interactive traffic has a minimal guarantee of 192 kbps (384 x 0.5), the business-critical traffic has a minimal guarantee of 76.8 kbps (384 x 0.2), and the bulk traffic has a minimal guarantee of 38.4 kbps (384 x 0.1).

Monitoring Class-Based Shaping

This topic identifies the Cisco IOS commands that are used to monitor class-based shaping.

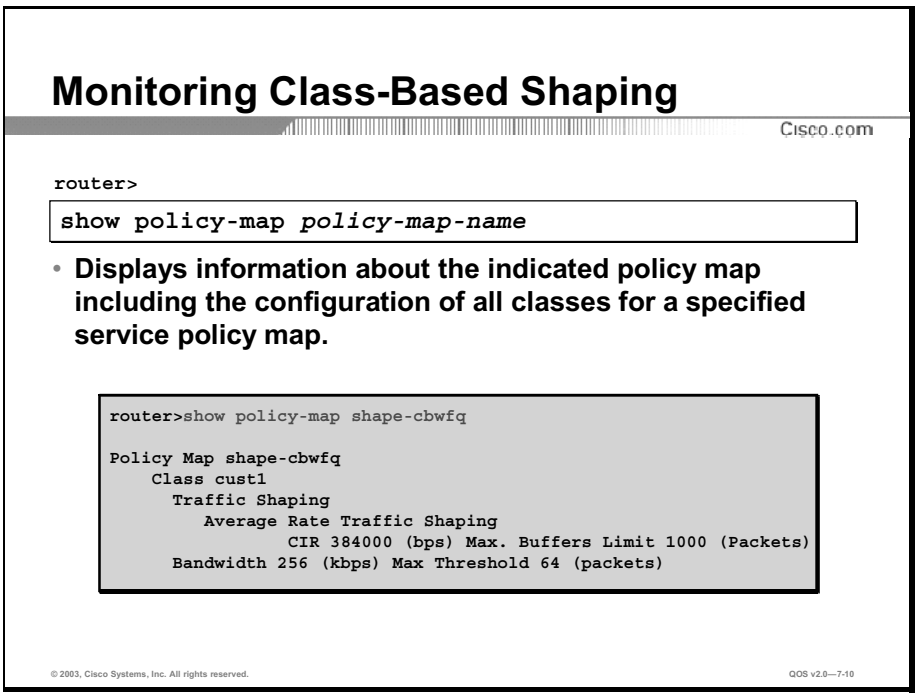

The **show policy-map** command displays the policy-map configuration.

This output represents the CBWFQ in conjunction with the class-based shaping configuration example shown earlier where the cust1 traffic class is shaped to an average rate of 384 kbps, with a default buffer limit of 1000 packets and a minimum bandwidth guarantee of 256 kbps.

Monitoring Class-Based Shaping (Cont.)


```
router#show policy interface s0/0

Serial0/0
  Service-policy output: shape-cbwfq
    Class-map: cust1 (match-all)
     0 packets, 0 bytes
     5 minute offered rate 0 bps, drop rate 0 bps
     Match: access-group 101
     Traffic Shaping
          Target/Average   Byte   Sustain   Excess    Interval  Increment
            Rate           Limit  bits/int  bits/int  (ms)      (bytes)
           38400/38400     1968   7872      7872      207       984
       Adapt  Queue     Packets   Bytes     Packets   Bytes     Shaping
       Active Depth                         Delayed   Delayed   Active
       -      0         0         0         0         0         no
     Queueing
       Output Queue: Conversation 266
       Bandwidth 256 (kbps) Max Threshold 64 (packets)
       (pkts matched/bytes matched) 0/0
       (depth/total drops/no-buffer drops) 0/0/0
   Class-map: class-default (match-any)
     0 packets, 0 bytes
     5 minute offered rate 0 bps, drop rate 0 bps
     Match: any
```


The **show policy-map interface** command displays all service policies that are applied to the interface.

This output represents the CBWFQ in conjunction with the class-based shaping configuration example shown earlier, where the cust1 traffic class is shaped to an average rate of 384 kbps and has a minimum bandwidth guarantee of 256 kbps.

Among the settings shown are the class-based shaping parameters like Be (excess bits/int), Bc (sustain bits/int), and Tc (interval), and other statistics.

In this example, the shape average rate is 384 kbps and the automatically calculated values for Bc and Be are 7872 bits, with a Tc of 207 ms.

Summary

This topic summarizes the key points discussed in this lesson.

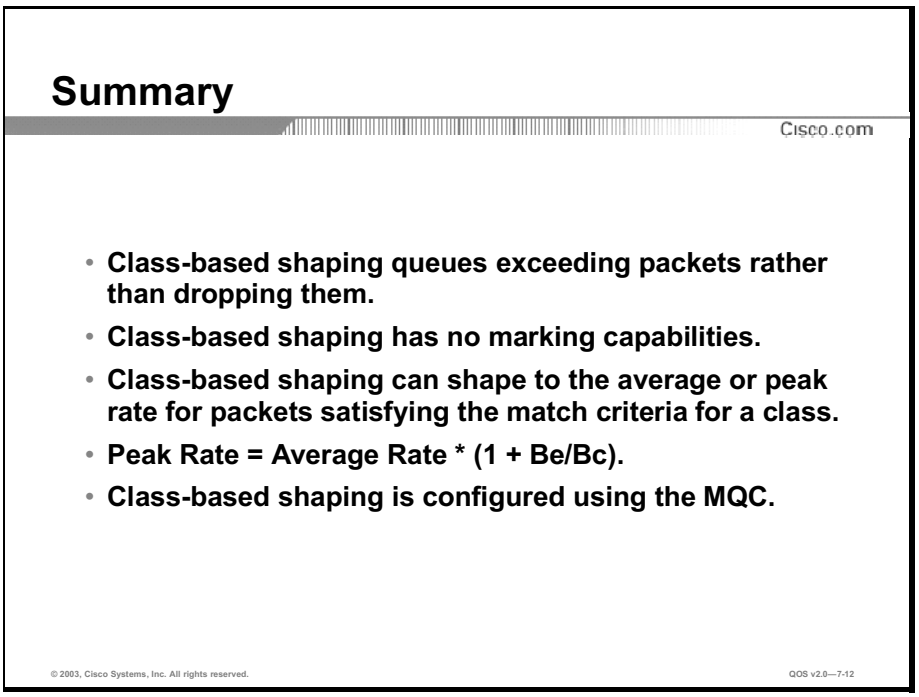

References

For additional information, refer to these resources:

- To learn more about class-based shaping, refer to “Configuring Class-Based Shaping” at the following URL: http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122cgcr/fqos_c/fqcprt4/qcfcbshp.htm
- To learn more about DTS, refer to “Configuring Distributed Traffic Shaping” at the following URL: http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122cgcr/fqos_c/fqcprt4/qcfdts.htm
- To learn more about FRTS, refer to “Configuring Frame Relay and Frame Relay Traffic Shaping” at the following URL: http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/12cgcr/qos_c/qcpart4/qcfrts.htm
- To learn more about FRTS, refer to “Traffic Shaping” at the following URL: http://www.cisco.com/warp/public/125/21.pdf

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) If the average rate is 64 kbps, the Bc is 8000 bits, and the Be is 4000 bits, what will the peak rate be?

A) 96 kbps

B) 128 kbps

C) 192 kbps

D) 224 kbps

Q2) Which one of the following statements is true regarding class-based shaping configurations?

A) Specifying the Bc and Be values are mandatory.

B) Specifying the Be value is only required for peak-rate shaping.

C) Specifying the Be value is only required to implement a dual token bucket.

D) Specifying the Bc and Be values are optional. If Bc and Be are not specified, IOS software will calculate an optimal value for them.

Q3) Which one of the following statements is true when implementing CBWFQ with class-based shaping?

A) CBWFQ defines the token bucket size used within class-based shaping.

B) CBWFQ defines the maximum rate limit for a traffic class using class-based shaping.

C) CBWFQ provides a minimum bandwidth guarantee to the traffic class, while class-based shaping provides a maximum rate limit for the traffic class.

D) CBWFQ can be used in conjunction with class-based shaping only if the shaping queue is configured as a FIFO queue.

Q4) Complete the following Cisco IOS MQC configuration to enable hierarchical class-based shaping based on the following requirements:

Shape the customer-one traffic class to an average rate of 64 kbps.

Within the customer-one traffic class, configure the gold class to use 40 percent of the bandwidth, the silver class to use 20 percent of the bandwidth, and the bronze class to use 10 percent of the bandwidth.

policy-map customer-one-classes

class mission-critical

_____

class transactional

_____

class bulk

_____

!

policy-map customer-one-policy

class customer-one

_____

_____

!

Interface serial 0/0.1

Service-policy output customer-one-policy

Quiz Answer Key

Q1) A

Relates to: Configuring Class-Based Shaping

Q2) D

Relates to: Configuring Class-Based Shaping

Q3) C

Relates to: Configuring Class-Based Shaping

Q4)

```
policy-map customer-one-classes
class mission-critical
bandwidth percent 40
class transactional
bandwidth percent 20
class bulk
bandwidth percent 10
!
policy-map customer-one-policy
class customer-one
shape average 64000
service-policy customer-one-classes
!
Interface serial 0/0.1
Service-policy output customer-one-policy
```

Relates to: Configuring Class-Based Shaping

Configuring Class-Based Shaping on Frame Relay Interfaces

Overview

FRTS builds upon existing Frame Relay support on Cisco routers by adding new capabilities that improve the scalability and performance of a Frame Relay network. FRTS can help eliminate bottlenecks in Frame Relay networks consisting of high-speed connections at a central hub site and low-speed connections at branch spoke sites. Using FRTS, network administrators can configure rate enforcement to either the CIR or some other defined value such as the excess information rate on a per-virtual circuit (VC) basis. One limitation of FRTS is that it applies only to Frame Relay PVCs and switched virtual circuits (SVCs).

Class-based shaping can be configured to behave the same as FRTS by allocating one DLCI per subinterface and using BECN support. This lesson describes the tasks to configure class-based traffic shaping to rate-limit certain traffic classes on Frame Relay interfaces.

Relevance

Traffic shaping is an important addition to Frame Relay networks because Frame Relay switches cannot determine which frames take precedence, and therefore which frames should be dropped when congestion occurs. By enabling traffic shaping over Frame Relay interfaces, especially when there are speed mismatches or network over-subscriptions, drops in the Frame Relay network can be controlled.

Objectives

Upon completing this lesson, you will be able to configure class-based shaping on Frame Relay WAN interfaces to rate-limit traffic. This includes being able to meet these objectives:

- Explain the purpose of the Frame Relay FECN, BECN, and the DE bit
- Explain the use of FECN and BECN as a Frame Relay congestion control mechanism
- Explain how class-based shaping can adapt dynamically to available Frame Relay bandwidth by integrating BECN signals
- Explain the FECN to BECN propagation mechanism
- Identify the Cisco IOS commands required to configure Frame Relay adaptive class-based shaping on Frame Relay interfaces
- Identify the Cisco IOS commands used to monitor class-based shaping on FR interfaces

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Basic knowledge of Frame Relay operation and configuration
- Knowledge of class-based shaping operation and configuration

Outline

The outline lists the topics included in this lesson.

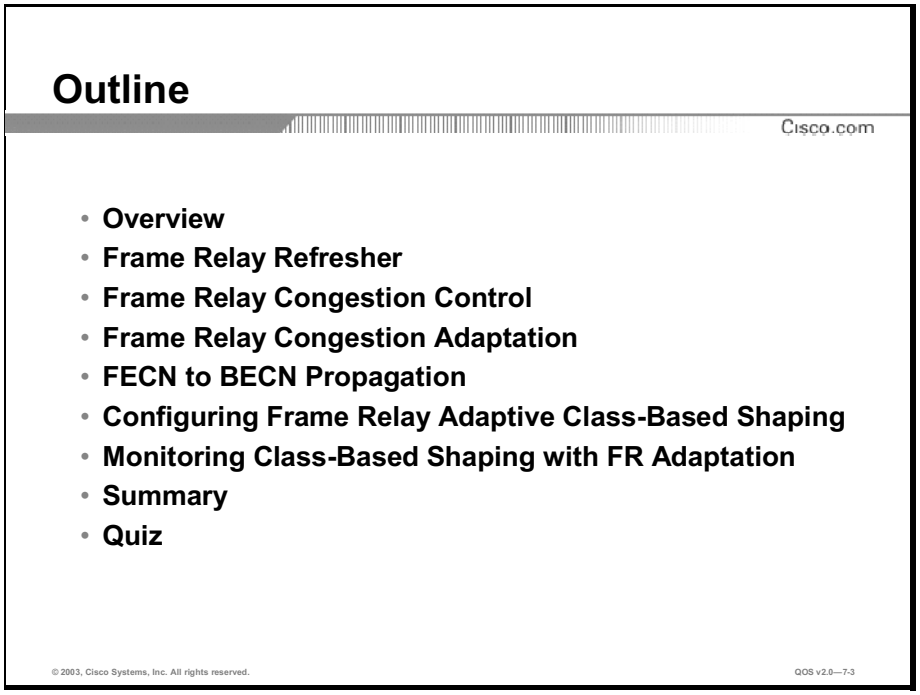

Frame Relay Refresher

This topic explains the purpose of the Frame Relay FECN, BECN, and the DE bit.

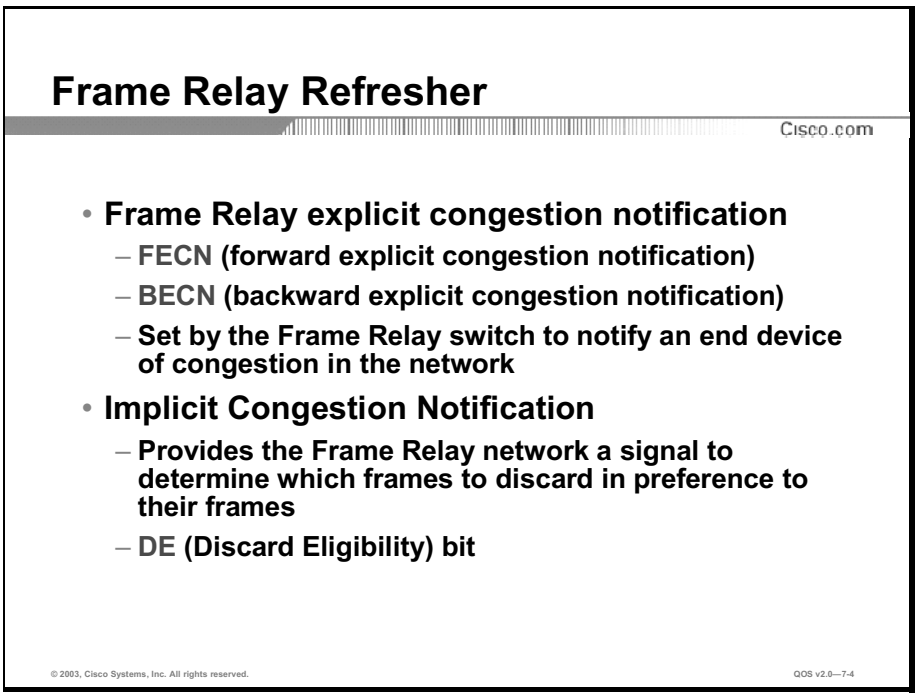

When the Frame Relay network becomes congested to the point that it cannot process new data transmissions, it begins to discard frames. These discarded frames are detected and then retransmitted by higher layer protocols, thus causing more congestion. In an effort to prevent this situation, several mechanisms have been developed to notify user devices at the onset of congestion, so that the offered load may be reduced.

Frame Relay performs congestion notification to its Layer 2 endpoints (FR DTE) by including congestion signaling inside the Layer 2 frame headers. The FECN, BECN,and DE bits in the Q.922 header of the frame provide in-band congestion signaling.

The FECN and BECN bit is set by a Frame Relay switch to notify an end device (FR DTE, which may be a router) that it should initiate congestion avoidance procedures.

When there is congestion, the network must decide which frames to discard. The DE bit provides the network with a signal to help determine which frames to discard. The DE bit indicates that a frame may be discarded in preference to other frames, if congestion occurs, to maintain the committed QoS within the network. The DE bit may be set by the user on lower-priority frames. Alternatively, the network may set the DE bit to indicate to other nodes that a frame should be preferentially selected for discard, if necessary.

Frame Relay Congestion Control

This topic explains the use of FECN and BECN as a Frame Relay congestion control mechanism.

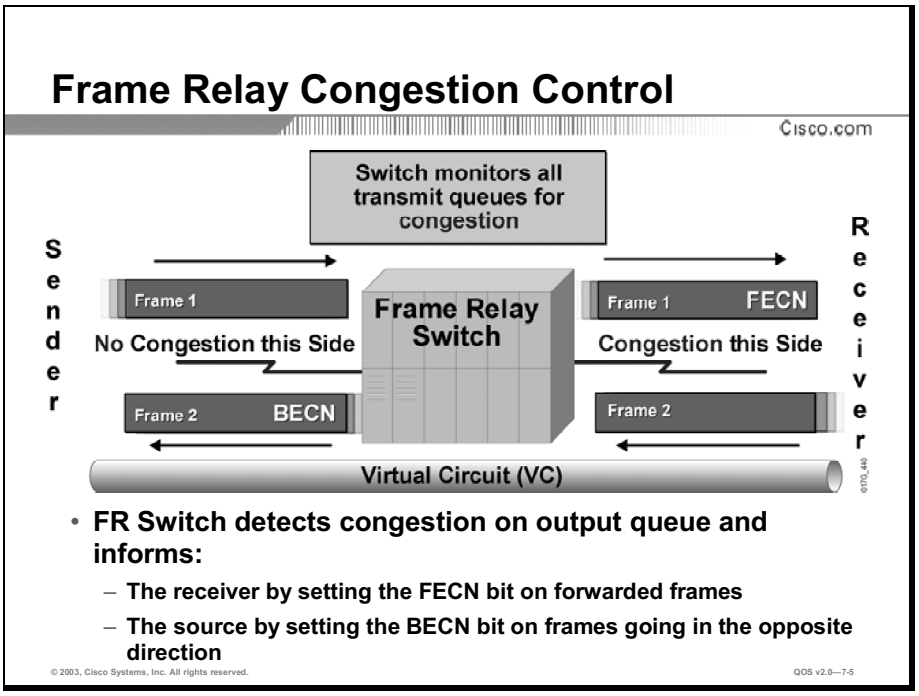

A Frame Relay switch can explicitly report congestion in two directions: forward and backward.

When a frame queue inside a switch is congested, the switch will generate congestion signals based on the FECN and BECN bits.

If congestion occurs in a queue towards the main receiver of traffic, FECN signals are sent to the receiving Layer 2 endpoint and BECN signals are sent to the sending Layer 2 endpoint.

FECN and BECN bits are not sent as separate frames, but are piggybacked inside data frames.

Frame Relay Congestion Adaptation

This topic explains how class-based shaping can adapt dynamically to available Frame Relay bandwidth by integrating BECN signals.

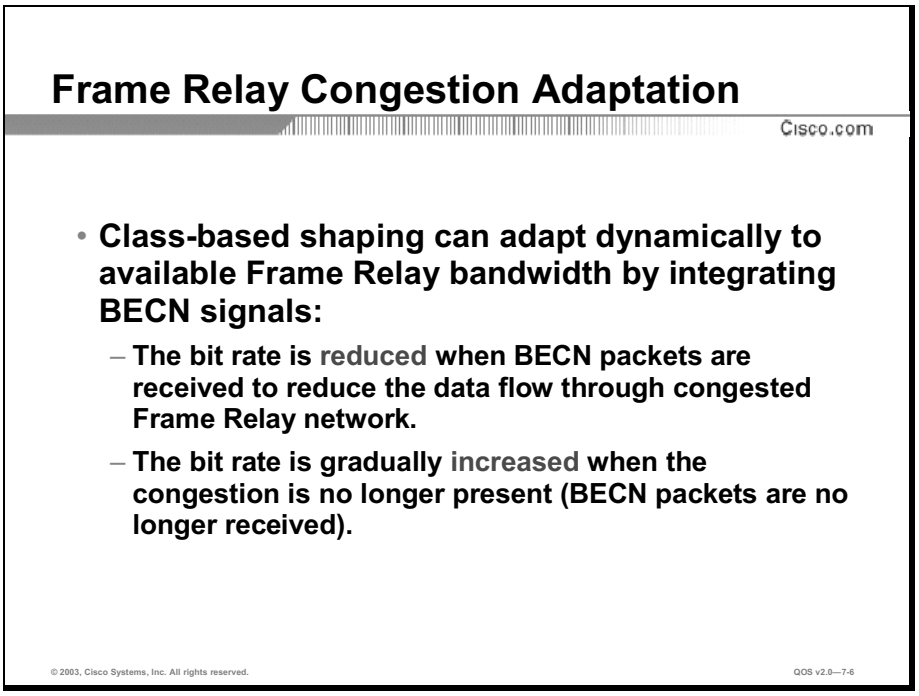

Class-based shaping is able to respond to Layer 2 congestion in the Frame Relay network by reducing its shaping rate to three-quarters of the current rate, until the Layer 2 Frame Relay network recovers from congestion. When BECN flags are no longer received, the rate is slowly ramped up again to the original shaping rate.

FECN to BECN Propagation

This topic explains the FECN to BECN propagation mechanism.

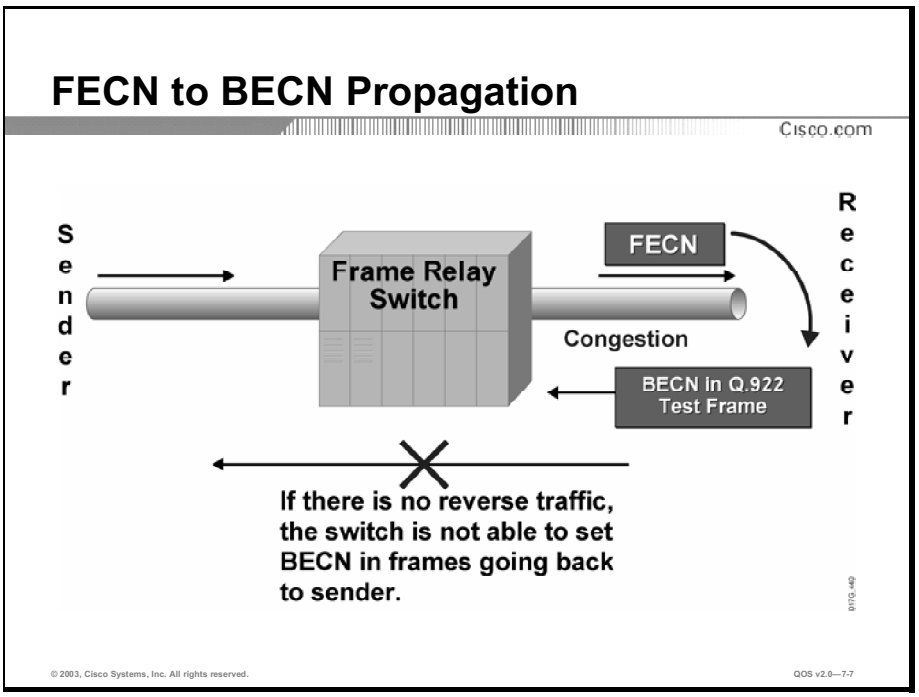

Another adaptation method, FECN to BECN propagation, configures a router Frame Relay subinterface to reflect received FECN bits as BECN in Q.922 Test Response messages. This enables the sender to notice congestion in the Layer 2 Frame Relay network, even if there is no data traffic flowing from the receiver back to the sender.

Configuring Frame Relay Adaptive Class-Based Shaping

This topic identifies the Cisco IOS commands that are required to configure Frame Relay adaptive class-based shaping on Frame Relay interfaces.

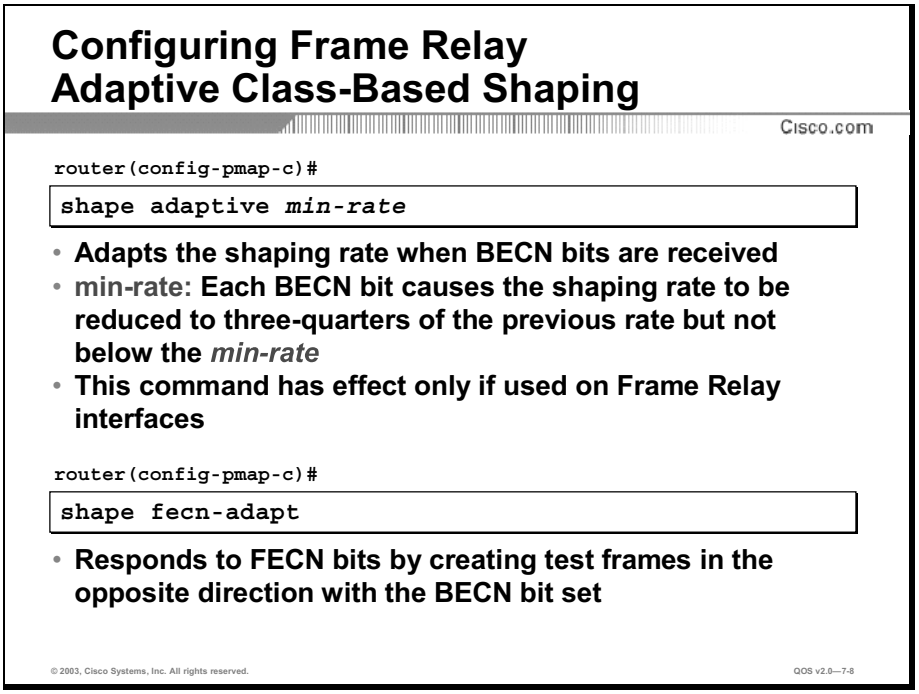

The **shape adaptive** command configures the class-based shaping system to adapt the shaping rate to BECN indications.

The **min-rate** (also called min-cir) parameter specifies the minimum-shaping rate that is allowed. Typically, this is set to the CIR guaranteed for the traffic class when there is congestion in the network. Taking a closer look at BECN integration:

- If router receives any BECNs during current time interval, it decreases the transmit rate by 25 percent. It will continue to drop with each BECN (limit one drop per time interval) until the traffic rate gets to min-rate where it stops.
- After the traffic rate has decreased, it takes 16 time intervals of receiving no BECNs to start to increase traffic again. The amount it increases by is (Be + Bc)/16. Thus it takes much longer to get back to CIR than it did to drop to min-rate (similar to slow start in TCP/IP).

The **shape fecn-adapt** command configures the class-based shaping system to respond to FECN-marked frames with BECN test frames.

Example: Class-Based Shaping with Frame Relay Adaptation

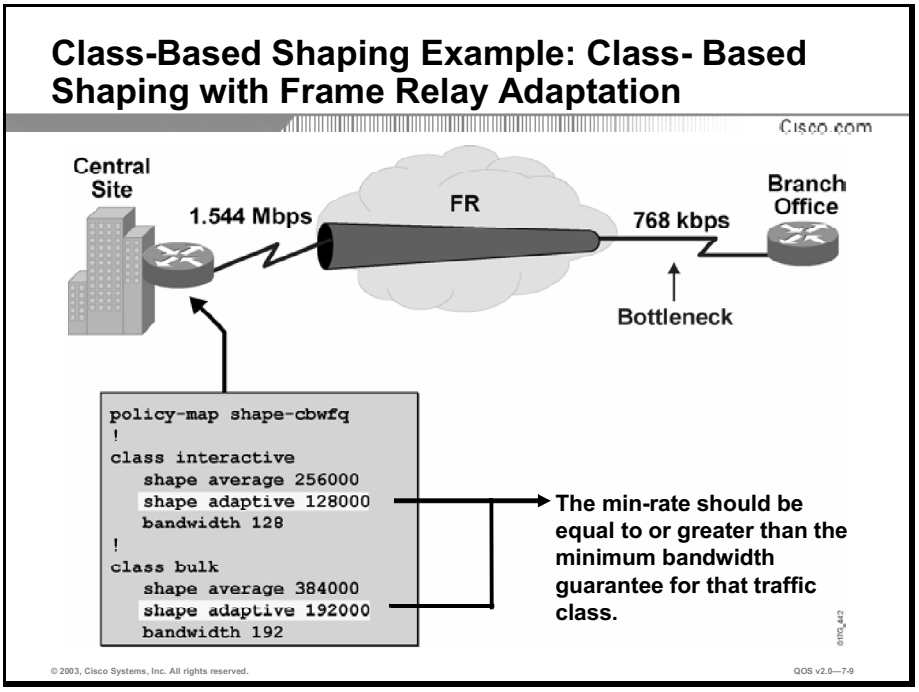

This example shows a typical class-based adaptive shaping in conjunction with CBWFQ configuration.

In this case, two traffic classes are defined with different shape rate and minimum bandwidth guarantee.

The interactive traffic class is shaped to an average rate of 256 kbps with a minimum bandwidth guarantee of 128 kbps. When BECN bit is detected, the shape rate can start throttling down to a min-rate of 128 kbps.

The bulk traffic class is shaped to an average rate of 384 kbps with a minimum bandwidth guarantee of 192 kbps. When the BECN bit is detected, the shape rate will start throttling down to a min-rate of 192 kbps.

When using class-based traffic shaping with CBWFQ, the min-rate should be equal to or greater than the minimum bandwidth guarantee for that traffic class.

Monitoring Class-Based Shaping with FR Adaptation

This topic identifies the Cisco IOS commands that are used to monitor class-based shaping on FR interfaces.

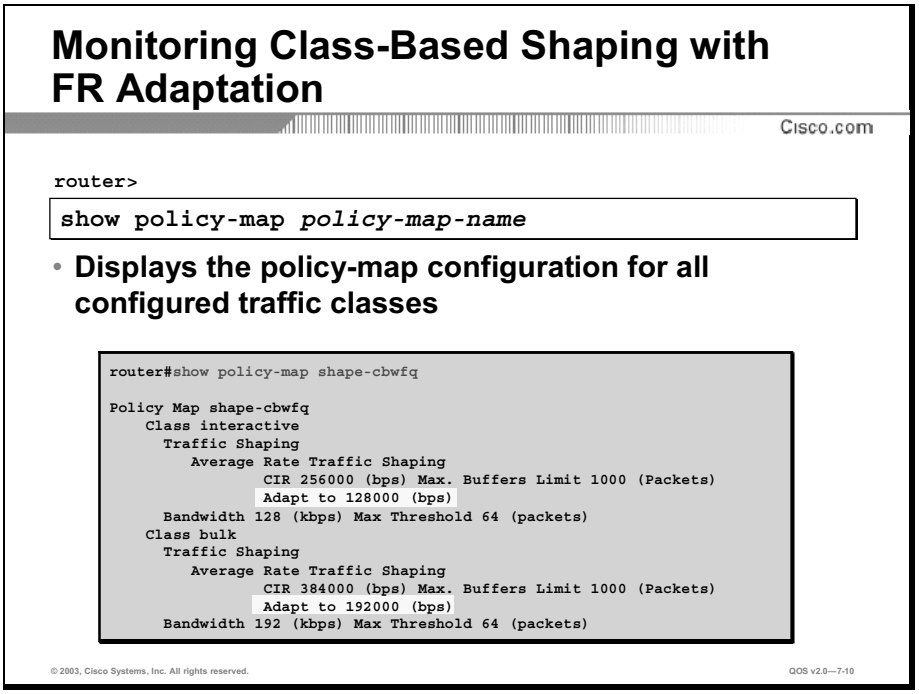

The **show policy-map** command displays the policy-map configuration.

In this example, class-based adaptive shaping is used in conjunction with CBWFQ on the interactive and bulk traffic classes.

The interactive traffic class is shaped to an average rate of 256 kbps with a minimum bandwidth guarantee of 128 kbps and a min-rate of 128 kbps.

The bulk traffic class is shaped to an average rate of 384 kbps with a minimum bandwidth guarantee of 192 kbps and a min-rate of 192 kbps.

Summary

This topic summarizes the key points discussed in this lesson.

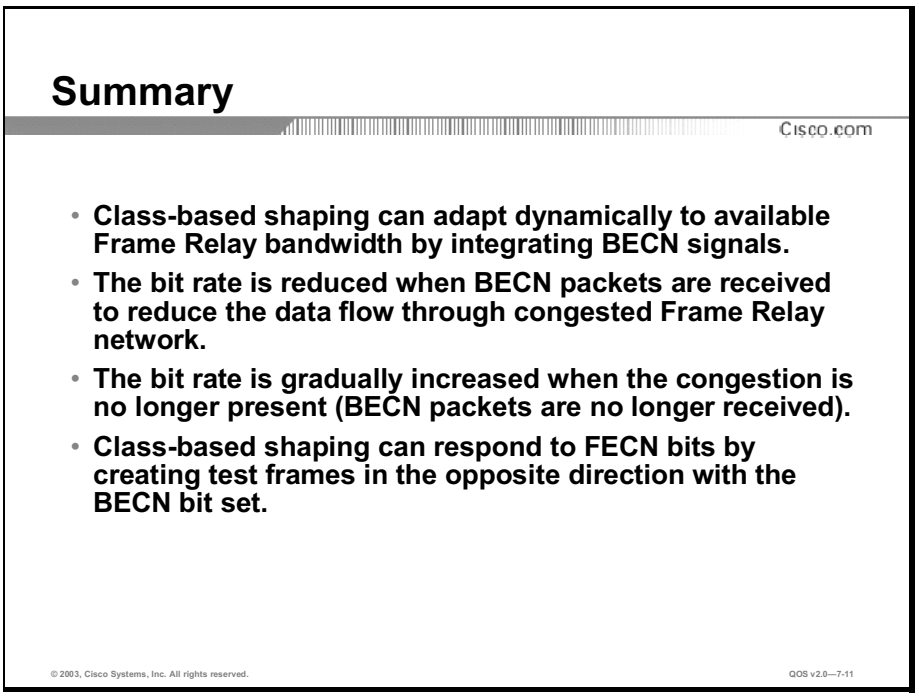

References

For additional information, refer to these resources:

- To learn more about class-based shaping, refer to “Configuring Class-Based Shaping” at the following URL: http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122cgcr/fqos_c/fqcprt4/qcfcbshp.htm
- For information about Frame Relay configuration, refer to “Configuring Frame Relay and Frame Relay Traffic Shaping” at the following URL: http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/12cgcr/qos_c/qcpart4/qcfrts.htm
- For information about Frame Relay configuration, refer to “Configuring and Troubleshooting Frame Relay” at the following URL: http://www.cisco.com/en/US/tech/tk713/tk237/technologies_tech_note09186a0080094372.shtml

Next Steps

For the associated lab exercise, refer to the following section of the course Lab Guide:

- Lab Exercise 7-2: Configuring Class-Based Shaping

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) Which three of the following bits in the Frame Relay header are used for in-band congestion signaling? (Choose three.)

A) CLP

B) DE

C) FECN

D) BECN

E) DLCI

F) PTI

G) LMI

Q2) When a router receives a frame with the FECN bit set, but the router has no reverse traffic to send back to the source, what can be enabled so that the router can send a test frame back to the source with the BECN bit set?

A) FECN-to-BECN propagation (shape fecn-adapt)

B) BECN-to-FECN propagation (shape becn-adapt)

C) adaptive traffic shaping (shape adaptive min-rate)

D) adaptive FECN traffic shaping (shape fecn min-rate)

Q3) Which one of the following statements correctly describes the BECN and FECN bit operations?

A) FECN and BECN bits are set by the Frame Relay switch to notify the Frame Relay DTE of congestion in the network using special Q.922 test frame.

B) FECN and BECN bits are set by the Frame Relay DTE to indicate to the network that a frame may be discarded in preference to other frames when the network is congested.

C) If congestion occurs in a queue towards the main receiver of traffic, BECN signals are sent to the receiving Layer 2 endpoint and FECN signals are sent to the sending Layer 2 endpoint.

D) If congestion occurs in a queue towards the main receiver of traffic, FECN signals are sent to the receiving Layer 2 endpoint and BECN signals are sent to the sending Layer 2 endpoint.

Q4) Which of the following correctly explains how the traffic rate adapts to the BECN congestion notification signal when adaptive Frame Relay traffic shaping is enabled on the router?

A) When a router receives a frame with the BECN bit set, it decreases the transmit rate directly to the min-rate. When a FECN bit is then received, the router increases the transmit rate by 25 percent. It will continue to increase the transmit rate with each FECN until the traffic rate gets to the average or peak shape rate.

B) When a router receives any BECNs, it decreases the transmit rate by 25 percent. It will continue to drop the transmit rate with each BECN until the traffic rate gets to the min-rate where it stops. When the traffic rate has been decreased, after 16 time intervals of receiving no BECNs, the traffic rate will start to increase by (Be + Bc)/16.

C) When a router receives a frame with the BECN bit set, it decreases the transmit rate directly to the min-rate. When a FECN bit is then received, the router increases the transmit rate back to the average or peak shape rate.

D) When a router receives any BECNs, it decreases the transmit rate by 25 percent. It will continue to drop the transmit rate with each BECN until the traffic rate gets to the min-rate where it stops. The traffic rate will increase back to the average or peak shape rate once a FECN bit is received.

Q5) Which of the following is the correct configuration to enable adaptive Frame Relay traffic shaping on the gold traffic class where the gold traffic will be shaped to an average rate of 768 kbps, and will adapt to a minimum rate of 256 kbps when there is congestion in the Frame Relay network?

A)
```
Class gold
Shape average 768000
Shape adaptive 256000
Bandwidth 256000
```

B)
```
Class gold
Shape average 256000
Bandwidth 768000
```

C)
```
Class gold
Shape average 768000
Shape fecn-adapt 256000
Bandwidth percent 50
```

D)
```
Class gold
Shape average 256000
Shape fecn-adapt
Bandwidth percent 50
```


Quiz Answer Key

Q1) B, C, D

Relates to: Frame Relay Refresher

Q2) A

Relates to: FECN to BECN Propagation

Q3) D

Relates to: Frame Relay Congestion Control

Q4) B

Relates to: Frame Relay Congestion Adaptation

Q5) A

Relates to: Configuring Frame Relay Adaptive Class-Based Shaping

Module Assessment

Overview

Use this assessment to test what you learned in this module. The correct answers and solutions are found in the Module Assessment Answer Key.

Quiz: Traffic Policing and Shaping

Complete the Quiz to assess what you have learned in the module.

Objectives

This activity tests your knowledge on how to meet these objectives:

- Explain how traffic policing and traffic shaping can be used to rate-limit traffic
- Configure class-based policing to rate-limit traffic
- Configure class-based shaping to rate-limit traffic
- Configure class-based shaping on Frame Relay WAN interfaces to rate-limit traffic

Instructions

Complete these steps:

Step 1 Answer all questions in this quiz by selecting the best answer(s) to each question,

Step 2 Verify your results against the answer key located at the end of this section.

Step 3 Review the topics in this module that relate to the questions that you answered incorrectly.

Q1) Which of the following is a major difference between traffic policing versus traffic shaping?

A) Traffic policing drops excess traffic while traffic shaping delays excess traffic by queuing them.

B) Traffic policing is only applied in the outbound direction while traffic shaping can be applied to both the inbound and outbound directions.

C) Traffic policing is not available on the Catalyst switches like the 2950. Traffic shaping is available on Catalyst switches like the 2950.

D) Traffic policing requires policing queues to buffer excess traffic while traffic shaping does not require any queues to buffer excess traffic.

Q2) What mathematical model is used by traffic policing mechanisms to meter traffic?

A) token bucket

B) RED

C) FIFO metering

D) predictor or stacker

Q3) When configuring single-rate class-based policing, what configuration parameter is used to enable a dual token bucket?

A) configuring a “violate” action

B) configuring an “exceed” action

C) configuring the PIR in addition to the CIR

D) configuring Be

Q4) What is the main advantage of using multiaction policing?

A) to distinguish between exceeding and violating traffic

B) to distinguish between conforming and exceeding traffic

C) to allow the setting of both Layer 2 and Layer 3 QoS markers at the same time

D) to allow marking of the traffic before transmission

Q5) Which three of the following situations typically requires the use of a rate-limiting mechanism? (Choose three.)

A) Frame Relay connection with mismatch speeds at the two end points

B) hub-and-spoke Frame Relay topology with over-subscription at the hub site

C) service provider providing sub-rate access

D) in the service provider high-speed backbone to avoid congestion

Q6) What is the min-rate?

A) the minimum bandwidth guaranteed by CBWFQ for the traffic class

B) the minimum bandwidth guaranteed by the LLQ for the traffic class

C) the minimum speed the traffic rate will throttle down to when the BECN bits are being received

D) the minimum speed the traffic rate will throttle down to when the FECN bits are being received

E) the minimum class-based policing traffic rate for a traffic class

Q7) Which two of the following statements is true when class-based shaping is used in conjunction with CBWFQ? (Choose two.)

A) The **bandwidth** command defines the minimum guaranteed bandwidth for the traffic class.

B) The **bandwidth** command defines the maximum guaranteed bandwidth for the traffic class.

C) The **shape peak** command defines the maximum rate limit for the traffic class.

D) The **shape average** command defines the minimum bandwidth guaranteed for the traffic class.

Q8) What is the function of FECN to BECN propagation?

A) to allow a receiving router to send back a test frame with the BECN bit set to the sender when the receiver has no data to send back to the sender after receiving the FECN bit

B) to allow a router to throttle down the traffic rate to the min-rate when either the FECN or BECN bit is received

C) to allow a router to throttle down the traffic rate to the min-rate when the BECN bit is received

D) to allow a router to throttle down the traffic rate to the min-rate when the FECN bit is received

E) to allow a receiving router to set the BECN bit on a user data frame going back to the sender after receiving the FECN bit

Q9) What are the two configuration options when configuring class-based traffic shaping? (Choose two.)

A) shape average
B) shape peak
C) single or dual token bucket
D) single or multiactions traffic shaping
E) single- or dual-rate traffic shaping

Q10) Which three are features of class-based policing? (Choose three.)

A) single or dual token bucket
B) single or multiactions policing
C) single- or dual-rate policing
D) single or dual FIFO queuing
E) single or dual drop threshold

Scoring

You have successfully completed the quiz for this lesson when you earn a score of 80 percent or better.

Module Assessment Answer Key

Q1) A

Relates to: Traffic Policing and Traffic Shaping Overview

Q2) A

Relates to: Traffic Policing and Traffic Shaping Overview

Q3) A

Relates to: Configuring Class-Based Policing

Q4) C

Relates to: Configuring Class-Based Policing

Q5) A, B, C

Relates to: Traffic Policing and Traffic Shaping Overview

Q6) C

Relates to: Configuring Class-Based Shaping on Frame Relay Interfaces

Q7) A, C

Relates to: Configuring Class-Based Shaping

Q8) A

Relates to: Configuring Class-Based Shaping on Frame Relay Interfaces

Q9) A, B

Relates to: Configuring Class-Based Shaping

Q10) A, B, C

Relates to: Configuring Class-Based Policing

Module Summary

This topic summarizes the key points discussed in this module.

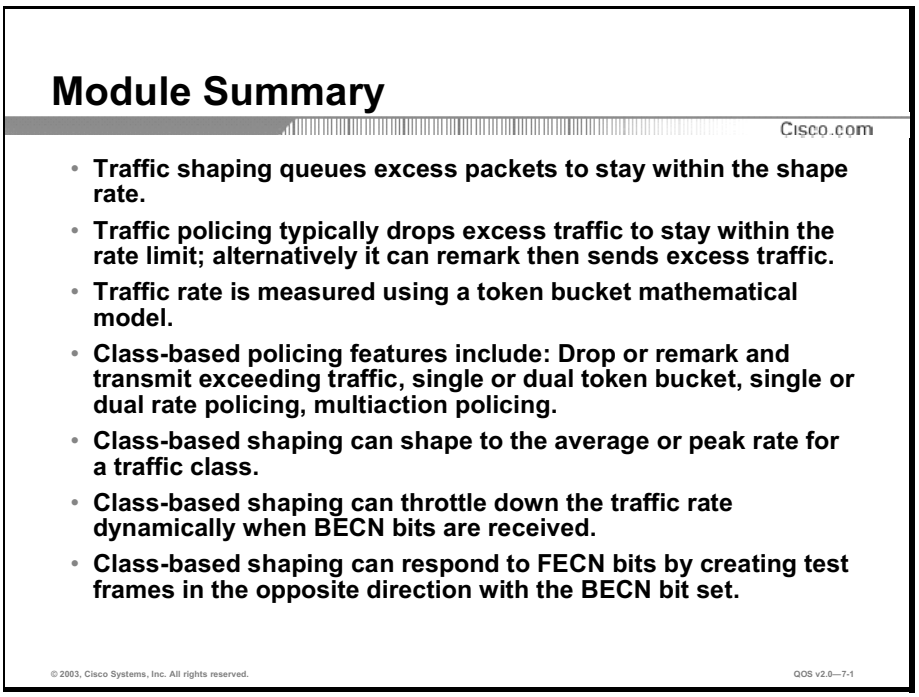

Rate-limiting is the ability to prevent excess traffic from entering or leaving the network. Rate-limiting is required because of speed mismatches, oversubscriptions, sub-rate access, and to prevent unwanted traffic from causing network congestions.

To measure the traffic rate, a token bucket is used. Parameters that define the operations of a token bucket includes the CIR, Bc (normal burst size), Be size, and Tc.

The CIR can be calculated by using the following formula:

- CIR (in bps) = Bc (in bits)/ Tc (in seconds)

Both traffic policing and traffic shaping are QoS mechanisms that are used to rate-limit a traffic class. Traffic policing operates by dropping excess traffic while traffic shaping delays excess traffic with the aid of queuing.

In Cisco IOS software, the most current rate-limiting mechanisms are class-based policing and class-based shaping. Both of these rate-limiting mechanisms are configured using the MQC. On Frame Relay interfaces, class-based shaping can be used in conjunction with CBWFQ and supports BECN adaptation and FECN-to-BECN propagation.

Module 8

Link Efficiency Mechanisms

Overview

In many networks consisting of a large WAN covering many sites, the monthly recurring costs of even the smallest link upgrades can be too high. In some cases, the only option for these low bandwidth remote sites is to upgrade their wide area circuit. In other cases, a set of quality of service (QoS) techniques can be used that can improve the efficiency of these low-speed WAN links. Header compression and payload compression mechanisms reduce the size of packets, reducing delay and increasing available bandwidth on a link. Other QoS link efficiency techniques such as link fragmentation and interleaving (LFI) allow fragile traffic types, such as voice and interactive traffic, to be sent either ahead or interleaved with larger, more aggressive flows. These techniques decrease latency and assist in meeting the service level requirements of delay-sensitive traffic types such as voice.

This module discusses the different link efficiency mechanisms (LEMs) that are available in Cisco IOS software to implement header compression, payload compression, and LFI.

Module Objectives

Upon completing this module, you will be able to use Cisco link efficiency mechanisms to improve the bandwidth efficiency of low-speed WAN links.

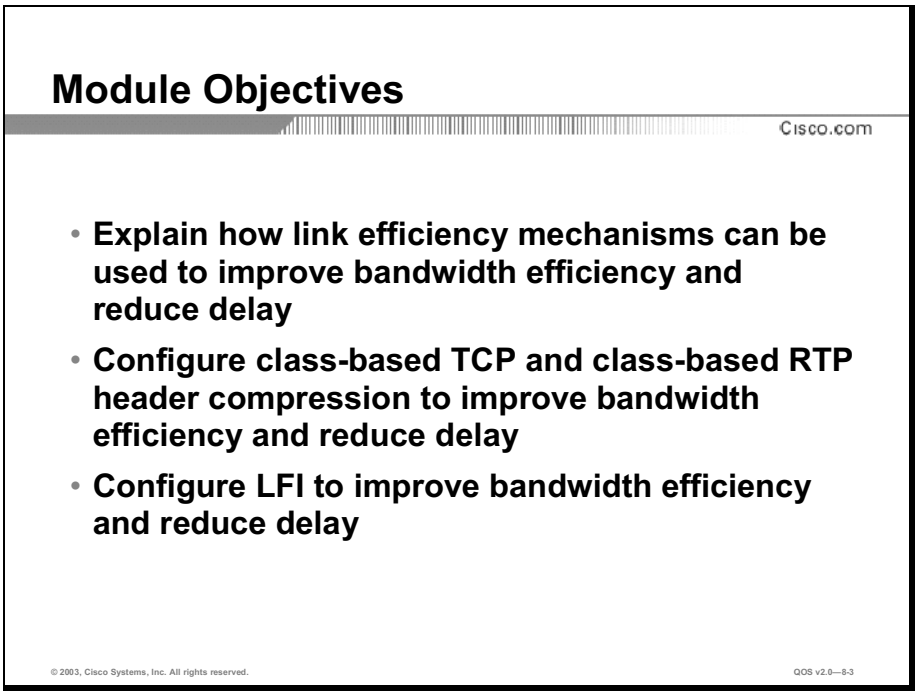

Module Outline

The outline lists the components of this module.

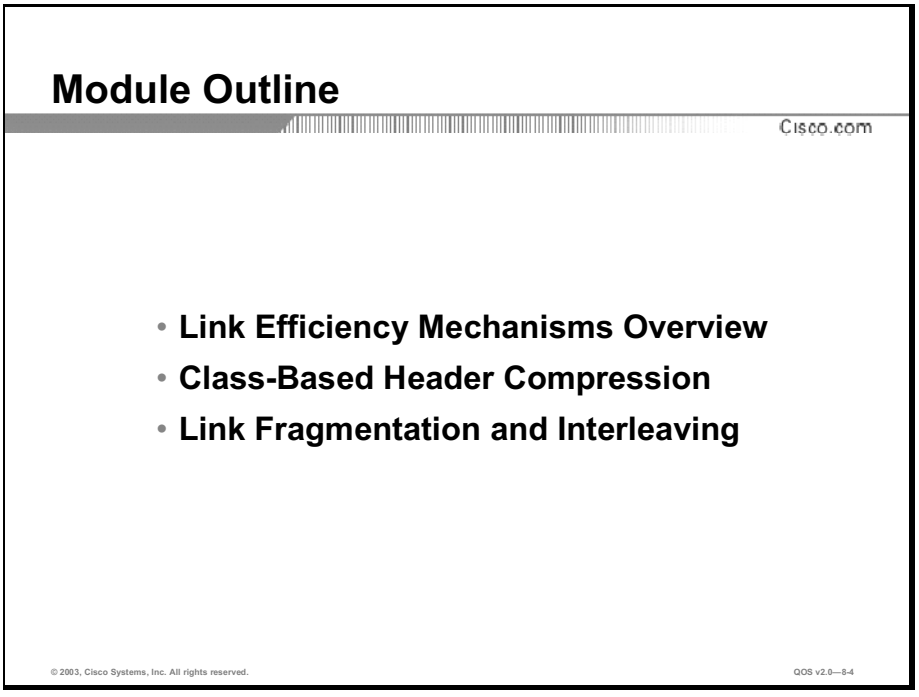

Link Efficiency Mechanisms Overview

Overview

Interactive traffic such as Telnet and Voice over IP (VoIP) are susceptible to increased latency when network processes using large packets, such as bulk FTP, traverse WAN links. Packet delay is especially significant when FTP packets are queued on slower links within the WAN. To solve delay problems on slow bandwidth links, a method for fragmenting larger frames, and then queuing the smaller frames between fragments of the large frames, is required. To this requirement, Cisco IOS software supports PPP multilink LFI as well as Frame Relay fragmentation (FRF.12 and FRF.11.C). In addition, other tools, such as header and payload compression techniques, can be used to reduce the size of frames that are sent on WAN links.

This lesson describes different approaches to improving the efficiency of WAN links. It discusses LEMs that either compress the payload of packets (stacker and predictor) or reduce packet headers overhead by compressing the headers (TCP and Real-Time Transport Protocol [RTP] header compression). It also discusses the different Layer 2 link LFI mechanisms (Multilink PPP [MLP] LFI and Frame Relay Fragmentation – FRF.12 and FRF.11.C) that are available in Cisco IOS software.

Relevance

WAN links are an expensive resource in all enterprise and service provider networks. Improving the efficiency of these expensive and vital resources is important because it can improve response times, reduce delays, and extend the longevity of some slower WAN links.

Objectives

Upon completing this lesson, you will be able to explain how link efficiency mechanisms can be used to improve bandwidth efficiency and reduce delay. This includes being able to meet these objectives:

- Explain the various link efficiency mechanisms and their functions
- Describe the purpose of Layer 2 payload compression and how Layer 2 payload compression affects throughput and delay
- Describe the purpose of header compression and how header compression affects throughput and delay
- Explain how VoIP packets are susceptible to increased latency when large packets such as FTP transfers traverse slow WAN links
- Explain LFI operation and how LFI reduces the delay and jitter of VoIP packets
- Identify the points in a network where link efficiency mechanisms can most effectively be employed

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Basic knowledge of internetworking with TCP/IP concepts

Outline

The outline lists the topics included in this lesson.

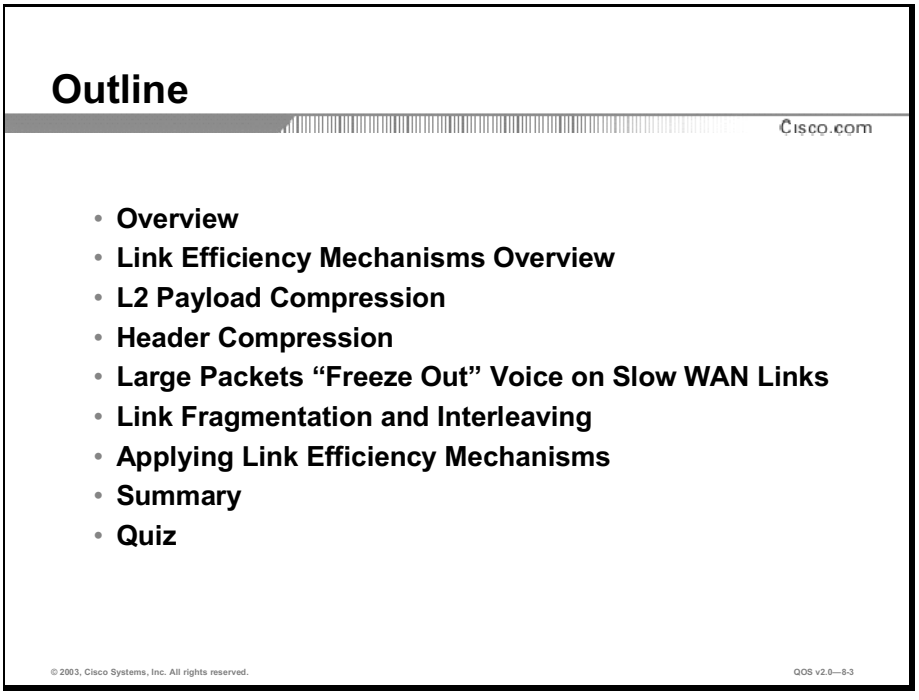

Link Efficiency Mechanisms Overview

This topic explains the various LEMs and their functions.

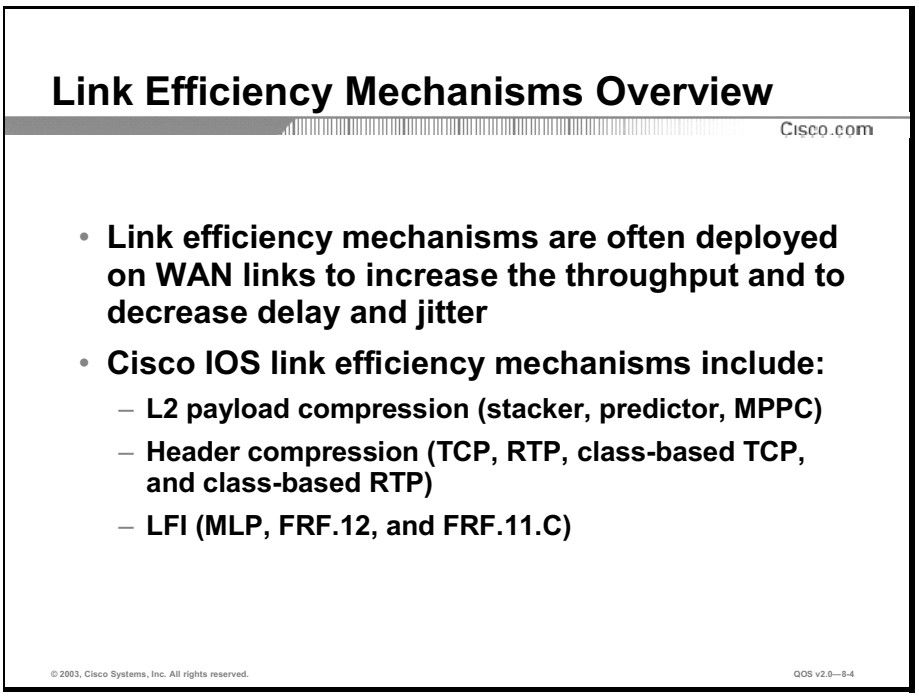

While many QoS mechanisms exist for optimizing throughput and reducing delay in network traffic, QoS mechanisms do not create bandwidth. QoS mechanisms optimize the use of existing resources, and enable the differentiation of traffic according to a policy.

Link efficiency QoS mechanisms like payload compression, header compression, and LFI are deployed on WAN links to optimize the use of WAN links.

Payload compression squeezes packet payloads, and therefore increases the amount of data that can be sent through a transmission resource in a given time period. Payload compression is primarily performed on Layer 2 frames and therefore compresses the entire Layer 3 packet. The Layer 2 payload compression methods available in Cisco IOS software include stacker, predictor, and Microsoft Point-to-Point Compression (MPPC). These algorithms differ vastly in their compression efficiency and in utilization of router resources.

RFC 2393, IP PCP (Payload Compression Protocol) is a fairly new technique for compressing the IP payload at Layer 3. (IP PCP will not be discussed in this lesson.)

All compression methods are based on eliminating redundancy when sending the same or similar data over a transmission medium. One piece of data, which is often repeated, is the protocol header. In a flow, the protocol header information of packets in the same flow does not change much over the lifetime of that flow. Therefore, using header compression mechanisms, most of header information could be sent only at the beginning of the session, stored in a dictionary, and then referenced in later packets by a short dictionary index. The header compression methods available in Cisco IOS include: TCP and RTP header compression, and class-based TCP and RTP header compression.

LFI is a Layer 2 technique where large Layer 2 frames are broken into small, equal-size fragments, and transmitted over the link in an interleaved fashion. Using LFI, large frames waiting in the queuing system are fragmented, smaller frames are prioritized, and a mixture of fragments is sent over the link. LFI reduces the queuing delay of small frames, as they are sent almost immediately. Link fragmentation therefore reduces delay and jitter by expediting the transfer of smaller frames through the hardware transmit (Tx) queue. The LFI methods available in Cisco IOS software include: MLP LFI, FRF.12 and FRF.11.C.

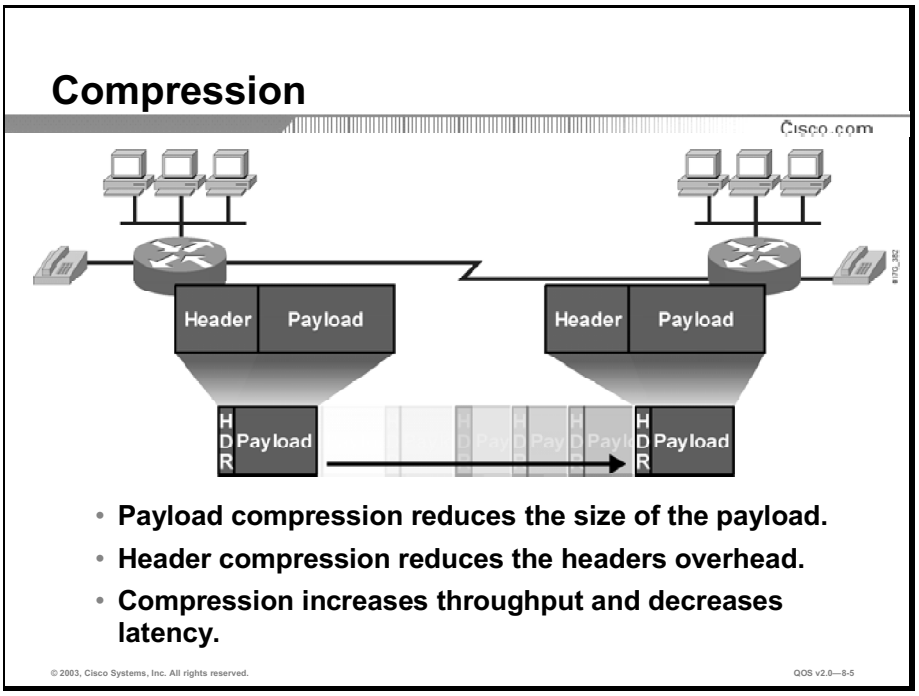

Layer 2 payload compression squeezes Layer 2 payloads (the entire Layer 3 packet). Layer 2 payload compression both increases the throughput and latency in transmission because smaller packets (with compressed payloads) take less time to transmit than the larger, uncompressed packets. Layer 2 payload compression is performed on a link-by-link basis.

Header compression methods work by *not* transmitting repeated information in packet headers throughout a session. The two peers on a PPP Layer 2 connection (such as a dial-up link) agree on session indices, which index a dictionary of packet headers. The dictionary is built at the start of every session and is used for all subsequent (non-initial) packets. Only changing (non-constant) parameters in the headers are actually sent along with the session index.

It is important to note that header compression is performed on a link-by-link basis. Header compression cannot be performed across multiple routers because routers need full Layer 3 header information to be able to route packets to the next hop.

Example: Indexing Operations

Compressed RTP (cRTP) maintains the state for session contexts. A session context is defined by the combination of the IP source and destination addresses, the User Datagram Protocol (UDP) source and destination ports, and the RTP synchronization source (SSRC) field. A compressor implementation might use a hash function on these fields to index a table of stored session contexts. The compressed packet carries a small integer, called the session context identifier, to indicate in which session context that packet should be interpreted. The decompressor can use the context identifier to index its table of stored session contexts directly.

Note: Refer to RFC 2508 for more information on cRTP operations.

L2 Payload Compression

This topic describes the purpose of Layer 2 payload compression and how Layer 2 payload compression affects throughput and delay.

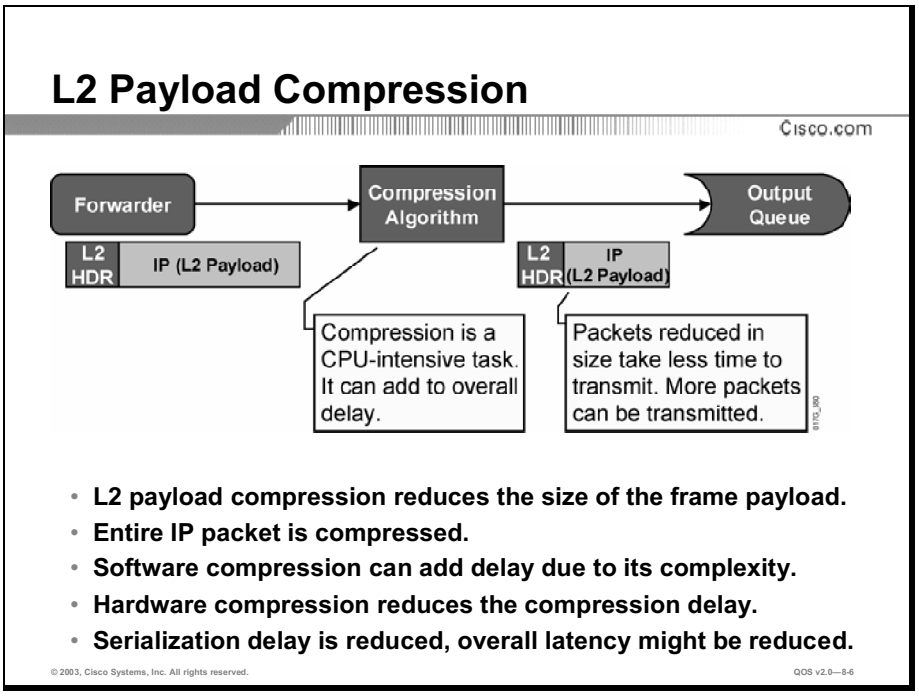

The figure shows a basic block diagram of a Layer 2 payload compression method. When a router forwards a packet, it is subjected to the Layer 2 compression method after it has been encapsulated at the output. The compression method squeezes the payload of the Layer 2 frame (the entire Layer 3 packet), and transmits the packet on the interface.

Layer 2 payload compression is a CPU-intensive task and can add per-packet compression delay due to the application of the compression method to each frame. The serialization delay, however, is reduced, because the resulting frame is smaller. Serialization delay is the fixed delay that is required to clock the frame onto the network interface. Depending on the complexity of the Layer 2 payload compression algorithm, overall latency might be reduced, especially on low-speed links.

Cisco routers support hardware-assisted compression to reduce the CPU load and the Layer 2 payload compression delay.

Example: L2 Payload Compression Results

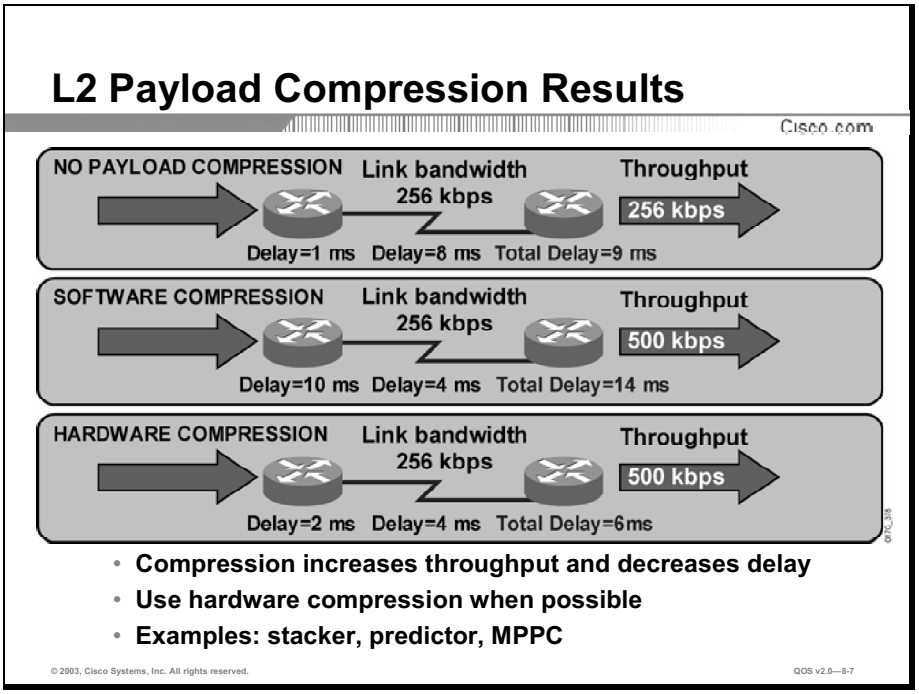

The figure compares three throughput/latency scenarios on a PPP link.

If no compression is used, the throughput is limited by the link bandwidth, and the average delay is influenced only by the forwarding/buffering delay, the serialization, and the propagation delay.

If compression is enabled—even if the serialization delay is now shorter because the frame is smaller—the compression/decompression delay may increase the overall latency between the two hops. The perceived throughput is generally increased because the size of the Layer 2 payload is reduced, therefore allowing more Layer 2 frames to be sent through a transmission resource in a given time period. The throughput is limited by the effectiveness of the Layer 2 payload compression algorithm and may be significantly higher than the link bandwidth limit.

If hardware-assisted Layer 2 payload compression is used, the compression/decompression delays may become insignificant compared to forwarding and serialization delays, and overall latency may decrease. The throughput is again limited by the effectiveness of the Layer 2 payload compression method and may be significantly higher than the link bandwidth limit.

Note: Layer 2 payload compression configuration will not be covered in this module. Refer to the latest Cisco IOS documentation for configuration details.

Header Compression

This topic describes the purpose of header compression and how header compression affects throughput and delay.

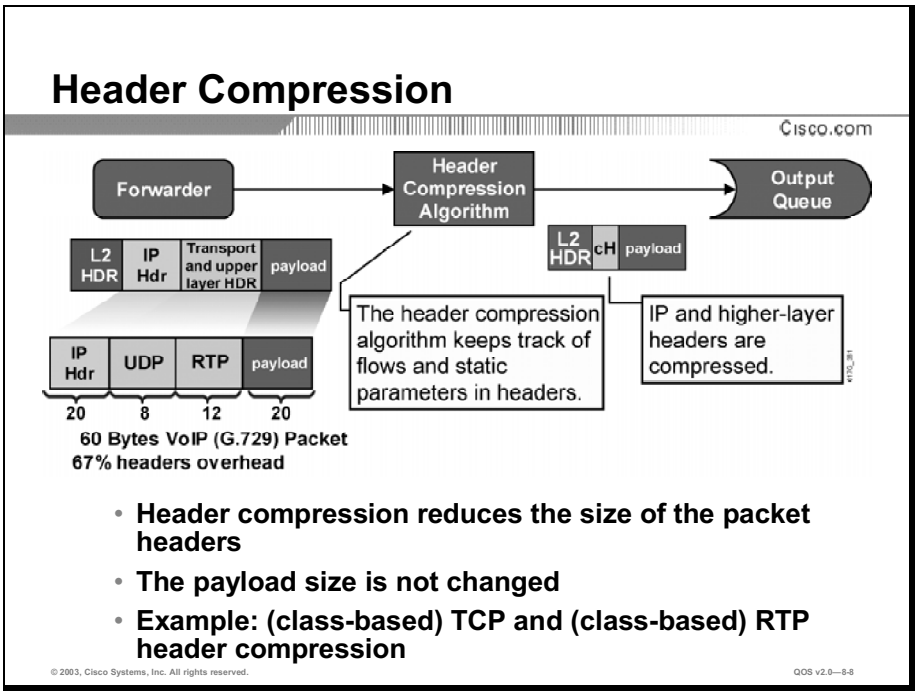

Header compression increases the perceived throughput and reduces the delay by compressing the protocol headers. Header compression is most useful for applications that generate small payloads because the protocol headers of such applications consume a significant percentage of available bandwidth on a link relative to their payload. Real-time applications typically generate small payloads. Target applications for header compression include Telnet and RTP applications, such as VoIP.

TCP and RTP header compression applies to all TCP and RTP flows. For example, if TCP compression is enabled on a link, there is no mechanism to restrict its function to specific application types. TCP header compression for bulk data transfer (packets with a large payload) yields little bandwidth savings. Using the newer class-based TCP header compression, TCP header compression can be performed only on a certain traffic class; for example, only perform TCP header compression on the Telnet traffic class.

The figure shows a block diagram of a header compression method. The header compression algorithm tracks active transport-layer connections over an interface. After the packet has been forwarded, the header compression algorithm compresses the Layer 3 and Layer 4 headers within the frame, and replaces them with a session index from the session dictionary (table). Only the non-constant parameters in the headers will be sent along with the session index. The packet is then sent to the output queue, and transmitted to the remote peer. When the remote peer receives the packet, the header is decompressed using the local session table, and passed to the forwarding process.

For example, without RTP header compression, the IP/UDP/RTP header overhead of the voice packet shown in the figure is about 67 percent (40/60 x 100 percent). With RTP header compression, the IP/UDP/RTP header can be reduced to 2 or 4 bytes (without and with checksum, respectively) for most packets, thus the IP/UDP/RTP header overhead can be reduced to about 9 percent (2/22 x 100 percent) or 17 percent (4/24 x 100 percent).

Example: Header Compression Results

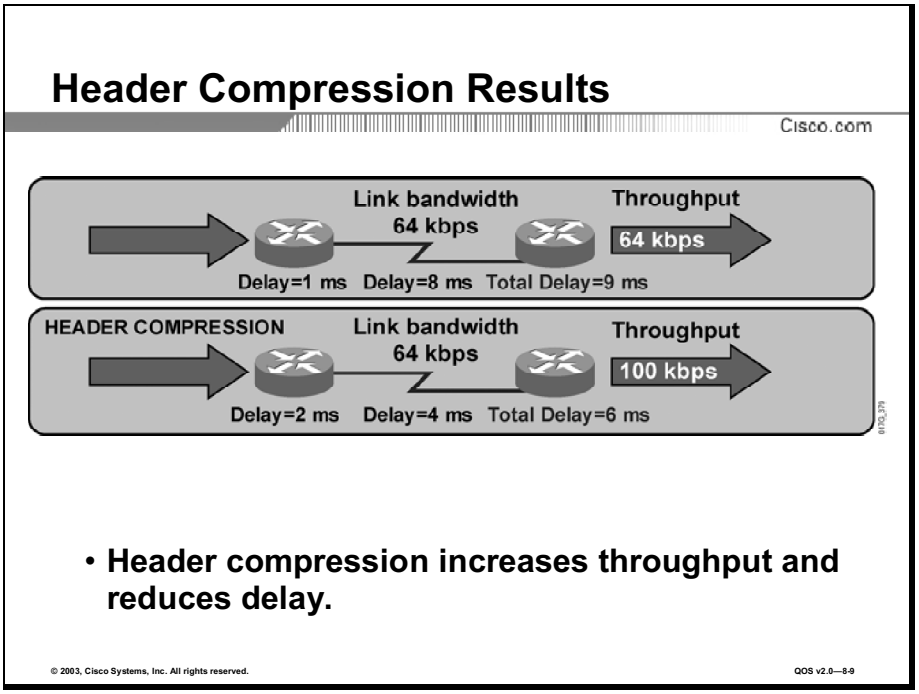

The figure compares two throughput/latency scenarios on a PPP link.

If header compression is not used, the throughput is limited by the link bandwidth, and the average delay is influenced only by the forwarding/buffering delay, the serialization, and the propagation delay.

If header compression is enabled, by compressing the protocol headers, the packet gets smaller, therefore allowing more packets to be sent through a transmission resource in a given time period to increase the throughput. Because the packet size is smaller, the serialization delay also becomes smaller, thus reducing the overall delay. Header compression has low CPU overhead and a low compression delay.

Large Packets “Freeze Out” Voice on Slow WAN Links

This topic explains how VoIP packets are susceptible to increased latency when large packets such as FTP transfers traverse slow WAN links.

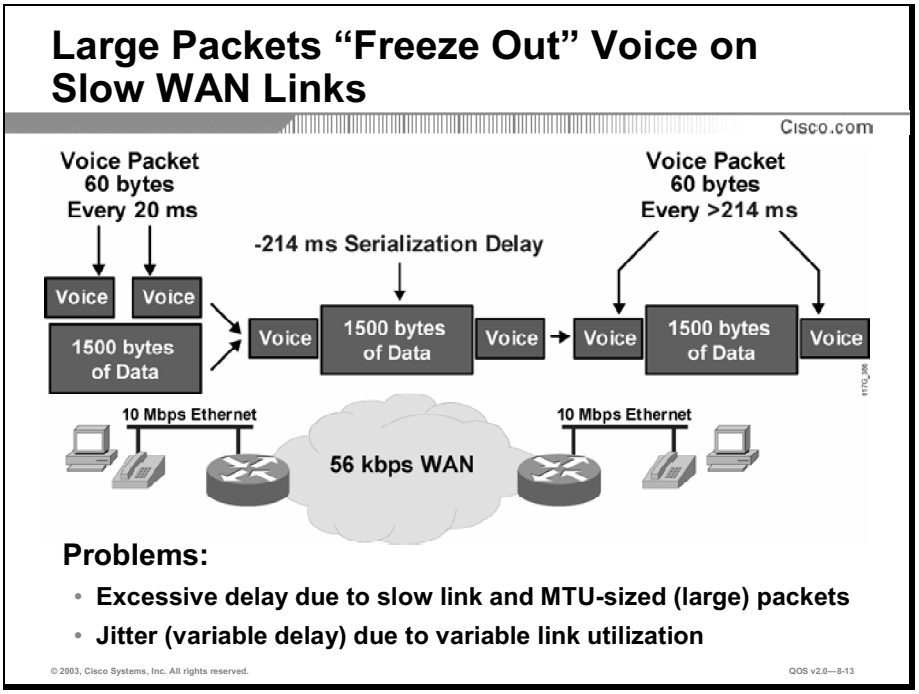

When considering delay between two hops in a network, queuing delay in a router must be considered because it may be comparable to—or even exceed—the serialization and propagation delay on a link. In an empty network, an interactive or voice session experiences low or no queuing delay because it does not compete with other applications on an interface output queue. Also, the small delay does not vary enough to produce considerable jitter on the receiving side.

In a congested network, interactive data and voice applications compete in the router queue with other applications. Queuing mechanisms may prioritize voice traffic in the software queue, but the hardware queue (Tx ring) always uses a FIFO scheduling mechanism. Therefore, after packets of different applications leave the software queue, they will mix with other packets in the hardware tx-queue (TxQ), even if their software queue processing was expedited. Thus, a voice packet may be immediately sent to the hardware TxQ where two large FTP packets may still be waiting for transmission. The voice packet must wait until the FTP packets are transmitted, thus producing an unacceptable delay in the voice path. Because links are variably utilized, this delay varies with time and may produce unacceptable jitter in jitter-sensitive applications such as voice.

Serialization Delays

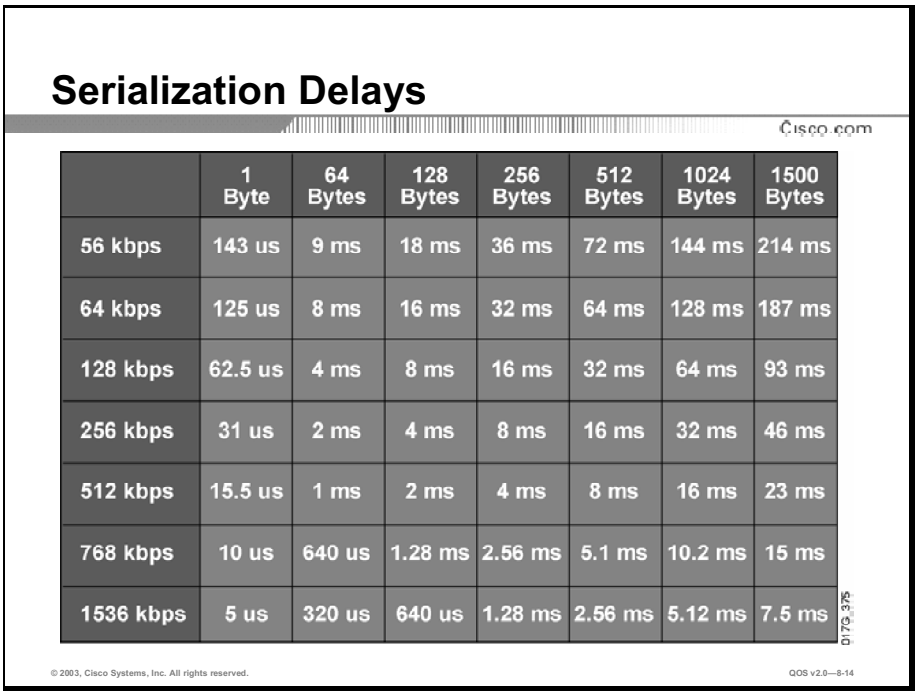

Serialization delay is the fixed delay that is required to clock a voice or data packet onto the network interface. Serialization delay is directly related to the link speed and the size of the packet.

This figure shows the serialization delay as a function of the link speed and packet size.

For example, the serialization delay for a 1500-byte packet over a 56-kbps link will be 214 ms, while the serialization delay is only 7.5 ms over a 1.536-Mbps link for the same 1500-byte packet.

Link Fragmentation and Interleaving

This topic explains LFI operation and how LFI reduces the delay and jitter of VoIP packets.

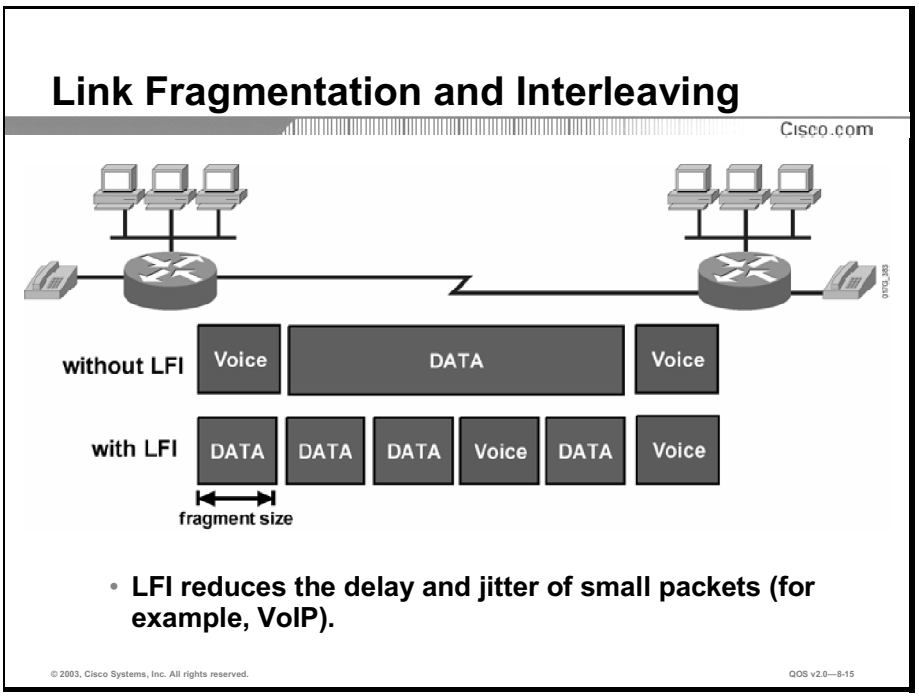

Using a hybrid queuing method like LLQ can provide low latency and low jitter for VoIP packets while servicing other data packets in a fair manner. But even if VoIP packets are always sent to the front of the software queue, there is still the issue of serialization delay. A large packet may be on its way out of the hardware TxQ, which is FIFO when a VoIP packet is sent to the front of the software queue. The serialization of the large packet can cause the VoIP packet to wait for a long time before it can be transmitted out.

For example, the serialization delay of a 1500-byte packet over a 56-kbps link will be 214 ms. For VoIP traffic, the maximum recommended end-to-end delay is 150 ms. Therefore, having a 1500-byte packet ahead of a VoIP packet in the hardware TxQ on a 56-kbps link can cause the end-to-end delay of the voice packet to be over the budget of 150 ms.

The solution to this problem is to fragment the large packets so that they never cause a VoIP packet to wait for more than a predefined amount of time. The VoIP packets must also be allowed to transmit in between the fragments of the larger packets (interleaving), or there will be no point in doing the fragmenting.

When configuring the proper fragment size to use on a link, a typical goal is to have a maximum serialization delay of around 10 to 15 milliseconds. Depending on the LFI mechanisms being configured, the fragment size is either configured in bytes or in milliseconds.

Applying Link Efficiency Mechanisms

This topic identifies the points in a network where link efficiency mechanisms can most effectively be employed.

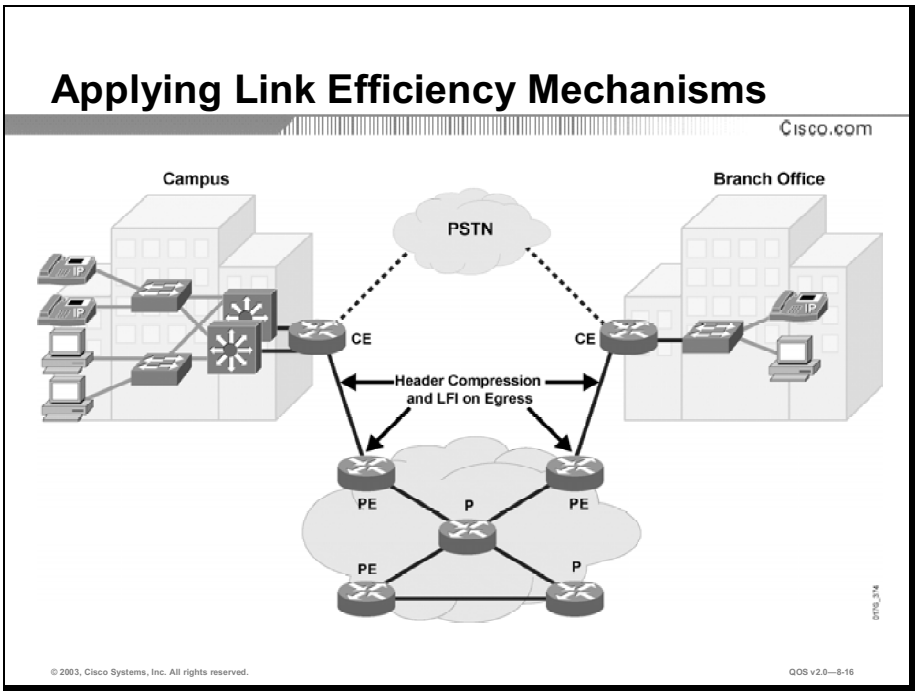

Header compression and LFI are typically configured at the WAN edge for WAN links below T1 or E1 speeds to optimize the use of the WAN link and to prevent long serialization delay.

Layer 2 payload compression is less commonly being deployed on WAN links, especially without the use of hardware-assisted payload compression.

Summary

This topic summarizes the key points discussed in this lesson.

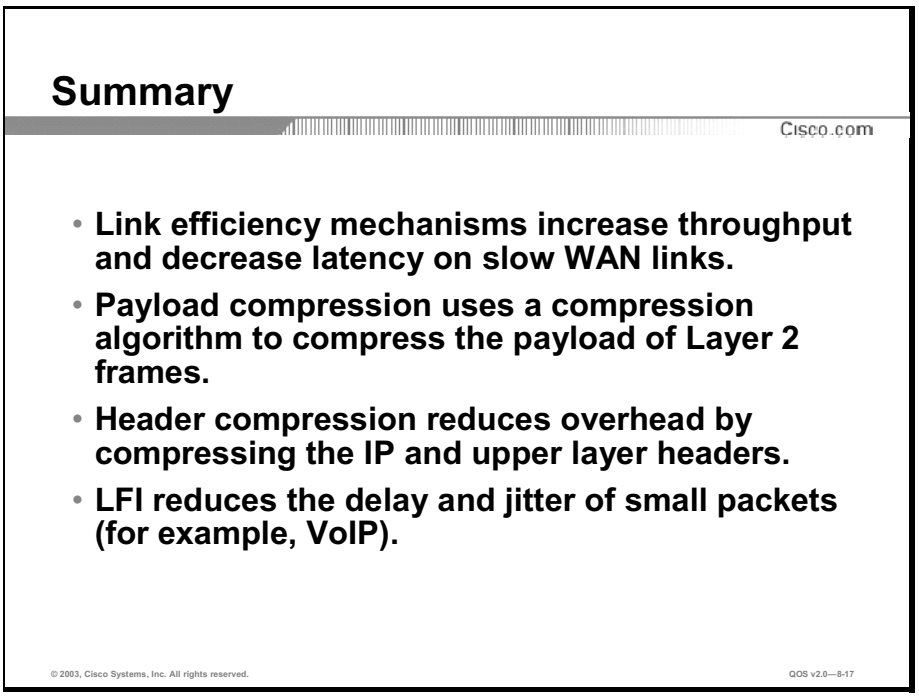

References

For additional information, refer to these resources:

- For more information on the link efficiency mechanisms, refer to “Link Efficiency Mechanisms Overview” at the following URL: http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122cgcr/fqos_c/fqcprt6/qcflem.htm
- For more information on the Layer 2 payload compression, refer to “Understanding Data Compression” at the following URL: http://www.cisco.com/warp/public/116/compress_overview.html

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) What are stacker, predictor, and MPPC?

A) Layer 2 payload compression methods

B) link fragmentation and interleaving methods

C) TCP header compression methods

D) RTP header compression methods

Q2) RTP header compression is used to reduce the overhead of which protocol header(s)?

A) only the RTP header

B) only the UDP and RTP headers

C) only the TCP and RTP headers

D) the IP, UDP, and RTP headers

E) the IP, TCP, and RTP headers

Q3) Using Layer 2 payload compression over a slow WAN link, which two statements are true about the resulting overall total delay and perceived throughput? (Choose two.)

A) The overall total delay will always decrease because the packet sizes are smaller resulting in lower serialization delay.

B) The overall total delay may increase due to compression/decompression delays if hardware-assisted compression is not used.

C) The perceived throughput is generally increased because the size of the Layer 2 payload is reduced, therefore allowing more Layer 2 frames to be sent through a transmission resource in a given time period.

D) The perceived throughput may be decreased due to compression/decompression delays if hardware-assisted compression is not used.

Q4) How many bytes are in the RTP and UDP headers?

A) UDP = 8 bytes. RTP = 12 bytes.

B) UDP = 8 bytes, RTP = 20 bytes.

C) UDP = 12 bytes, RTP = 8 bytes.

D) UDP = 12 bytes. RTP = 20 bytes.

Q5) What problem does LFI solve?

A) A large packet in the software queue delaying a VoIP packet because the software queue is FIFO.

B) The serialization delay of VoIP packets over a slow WAN link.

C) The large overhead of the VoIP packet protocol headers.

D) A large packet may be on its way out of the hardware tx-queue which is FIFO when a VoIP packet is sent to the front of the software queue. The serialization of the large packet can cause the VoIP packet to wait for a long time before it can be transmitted out.

Q6) What should be the target budget for the end-to-end delay of VoIP packets to ensure high voice quality?

A) 10 to 15 ms

B) 50 ms

C) 150 ms

D) 300 ms

Quiz Answer Key

Q1) A

Relates to: Link Efficiency Mechanisms Overview

Q2) D

Relates to: L2 Payload Compression

Q3) B, C

Relates to: L2 Payload Compression

Q4) A

Relates to: Header Compression

Q5) D

Relates to: Link Fragmentation and Interleaving

Q6) C

Relates to: Large Packets “Freeze Out” Voice on Slow WAN Liinks

Class-Based Header Compression

Overview

Headers exist on almost every communication layer of the Open System Interconnection (OSI) stack. When data is sent between workstations, headers will typically be applied at the session, transport, network, and data link layers.

RTP is a protocol for the transport of real-time data. RTP includes a data portion and a header portion. The data portion of RTP is a thin protocol that provides support for the real-time properties of applications, such as continuous media, including timing reconstruction, loss detection, and content identification. The header portion of RTP is considerably large, especially when compared to the payload data that it supports. To avoid the unnecessary consumption of available bandwidth, RTP header compression (referred to as cRTP) is used on a link-by-link basis.

TCP header compression is also supported on Cisco routers to reduce the overhead that is associated with TCP and IP headers in TCP/IP packets. TCP header compression is most effective for interactive traffic with small packet size such as Telnet on slow WAN links.

Using the MQC method, Cisco IOS software supports class-based RTP and TCP header compression so that the header compression can be performed on a specific traffic class.

This lesson discusses both class-based RTP and class-based TCP header compression, including configuration and monitoring.

Relevance

To improve the WAN link efficiency, header compression can be used to reduce the overhead of the protocol headers. This is especially effective for packets having small payload, like VoIP packets; and interactive TCP packets, like Telnet packets. Header compression is essential on low-speed WAN links where voice transport is a requirement.

Objectives

Upon completing this lesson, you will be able to configure class-based TCP and class-based RTP header compression to improve bandwidth efficiency and reduce delay. This includes being able to meet these objectives:

- Explain the purpose of header compression and provide some common reasons for its use
- Roughly calculate the overhead that can be saved by using class-based TCP header compression
- Roughly calculate the overhead that can be saved on VoIP packets using class-based RTP header compression
- Identify the Cisco IOS commands required to configure class-based header compression
- Identify the Cisco IOS commands used to monitor class-based header compression

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Basic knowledge of internetworking with TCP/IP concepts
- Knowledge of using MQC to implement Cisco QoS mechanisms

Outline

The outline lists the topics included in this lesson.

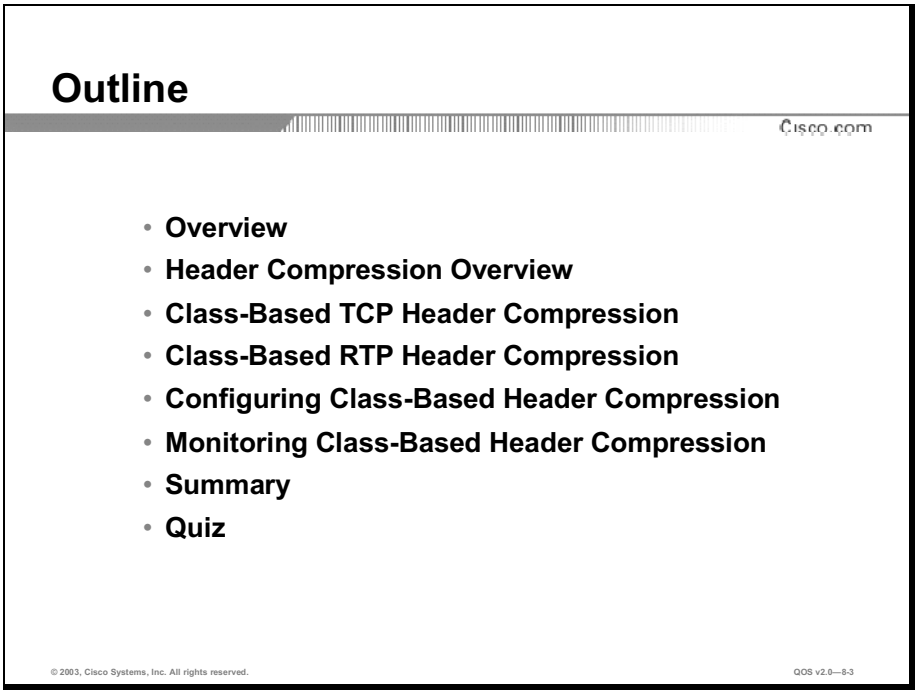

Header Compression Overview

This topic explains the purpose of header compression and provides some common reasons for its use.

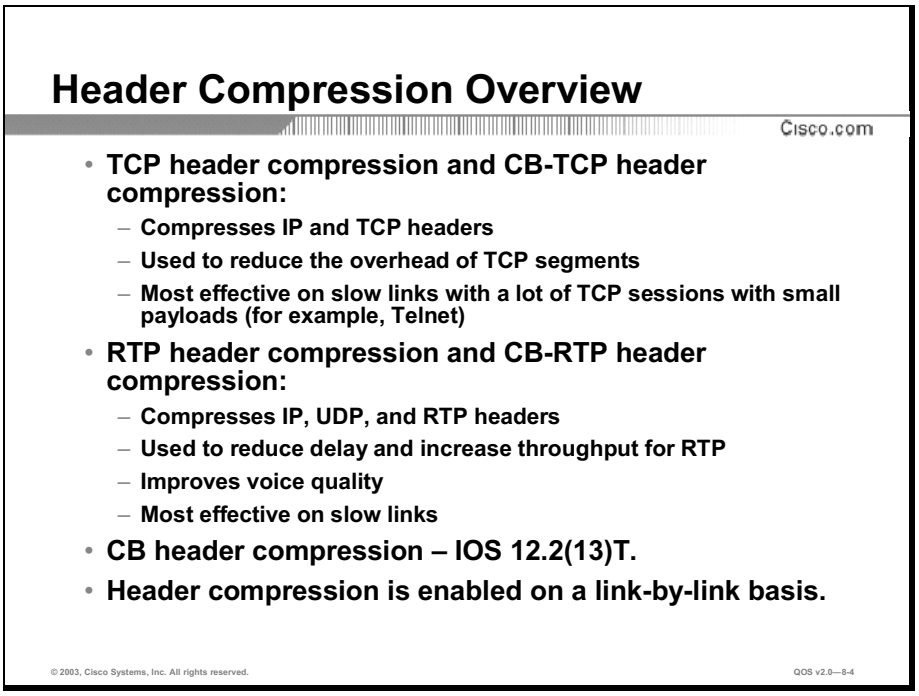

All compression methods are based on eliminating redundancy when sending the same or similar data over a transmission medium. One piece of data, which is often repeated, is the protocol header. In a flow, the header information of packets in the same flow does not change much over the lifetime of that flow. Therefore, most of header information could be sent only at the beginning of the session, stored in a dictionary, and then referenced in later packets by a short dictionary index.

Two methods were standardized by the Internet Engineering Task Force (IETF) for use with IP protocols:

- TCP header compression (also known as the Van Jacobson or VJ header compression) is used to compress the packet IP and TCP headers over slow links, thus considerably improving the interactive application performance.
- RTP header compression is used to compress the packet IP, UDP, and RTP headers, thus lowering the delay for transporting real-time data, such as voice and video over slower links.

When TCP and RTP header compression is enabled, it occurs by default in the fast-switched path or the Cisco Express Forwarding (CEF)-switched path, depending on which switching method is enabled on the interface.

Class-based header compression enables RTP or TCP header compression on a per-class basis. This feature was introduced in Cisco IOS release 12.2(13)T.

Note: It is important to note that header compression is performed on a link-by-link basis. Header compression cannot be performed across multiple routers because routers need full Layer 3 header information to be able to route packets to the next hop.

Example: RTP Header Compression

RTP header compression is often implemented on WAN links to reduce the IP/UDP/RTP header overhead for VoIP packets. Reducing the protocol header overhead will result in increased throughput and decrease in serialization delay.

Class-Based TCP Header Compression

This topic roughly calculates the overhead that can be saved by using class-based TCP header compression.

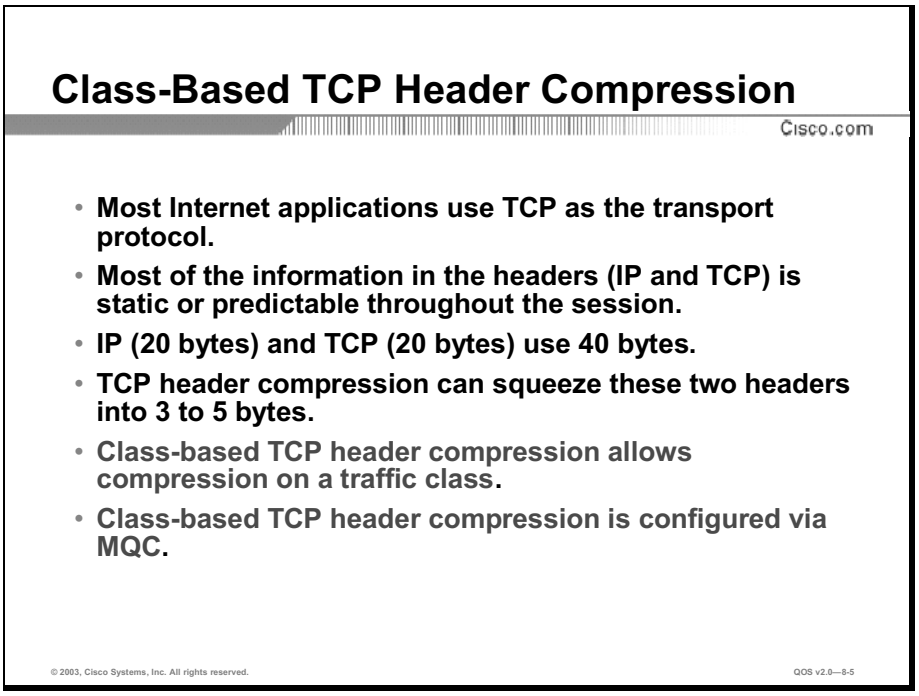

With TCP header compression, the IP and TCP headers, which normally use 20 bytes each, is reduced to a session index, and the changing part of the header. With all optimizations, the combined header length of 40 bytes can be reduced to a 3 to 5-byte compressed header.

Class-based TCP header compression enables TCP header compression on a per-class basis, when a class is configured within a policy map. Policy maps are created using the modular Quality of Service (QoS) Command-Line Interface (CLI) (MQC). The policy maps are attached to an interface by using the **service-policy** command. The **service-policy** command gives customers the option of specifying either an input service policy (for input interfaces), or an output service policy (for output interfaces). For this feature, only output service policies can be specified.

Enabling TCP header compression on an interface applies header compression to all TCP flows out of the interface. If TCP compression is enabled on a link, there is no mechanism to restrict its function to specific application types. TCP header compression for bulk data transfer (packets with a large payload) yields little bandwidth savings. Using the newer class-based TCP header compression, TCP header compression can be performed only on certain traffic classes; for example, only perform TCP header compression on the interactive traffic class where the packet payload sizes are small.

Example: Class-Based TCP Header Compression

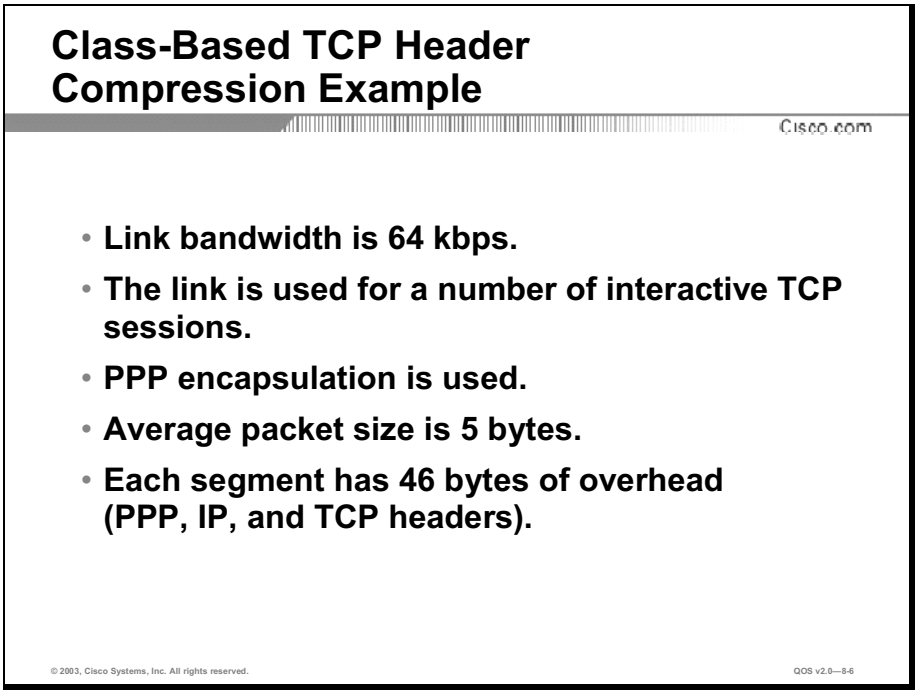

This example illustrates the benefits of TCP header compression on slow links. A 64-kbps link is used to transport a TCP-based application using PPP as the Layer 2 framing protocol. For the case study application (Telnet), the average packet payload size is 5 bytes. Because PPP has 6 bytes of frame header, the total header overhead is 6 + 20 + 20 = 46 bytes, counting the PPP, IP, and TCP headers.

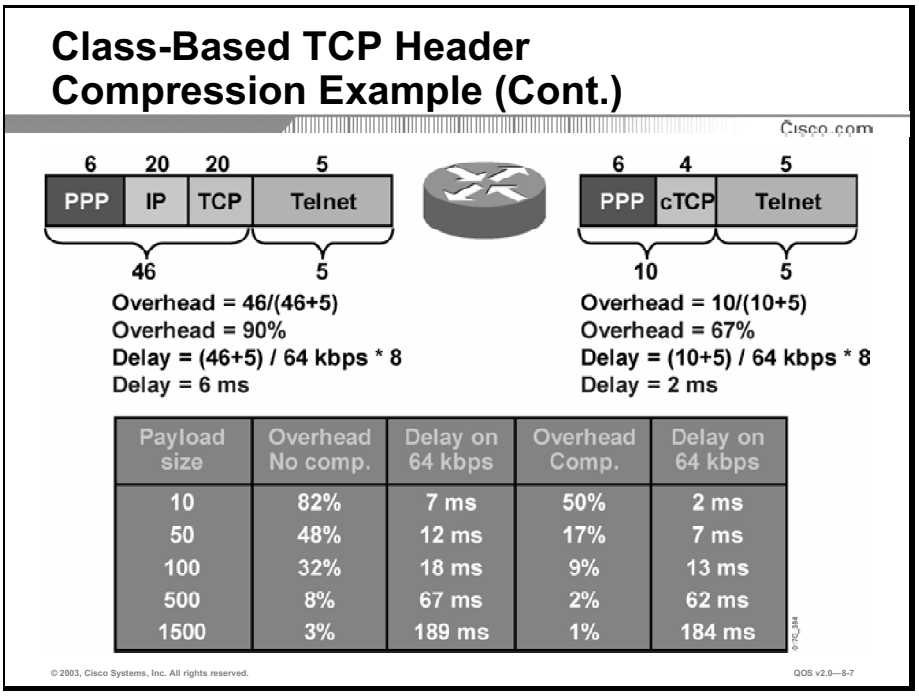

The figure shows the packet size before and after TCP header compression. After TCP header compression, the IP and TCP headers are reduced to 4 bytes, resulting in 10 bytes of overall headers. The overhead is reduced from 90 percent to 67 percent, when small packets are used. Because of the packet size reduction, the serialization delay decreases from 6 ms to 2 ms on the same 64-kbps link.

The table in the figure shows how TCP header compression impacts performance when different packet sizes are used. TCP header compression is most effective on small packets, and is often used on slow links.

Class-Based RTP Header Compression

This topic roughly calculates the overhead that can be saved on VoIP packets using class-based RTP header compression.

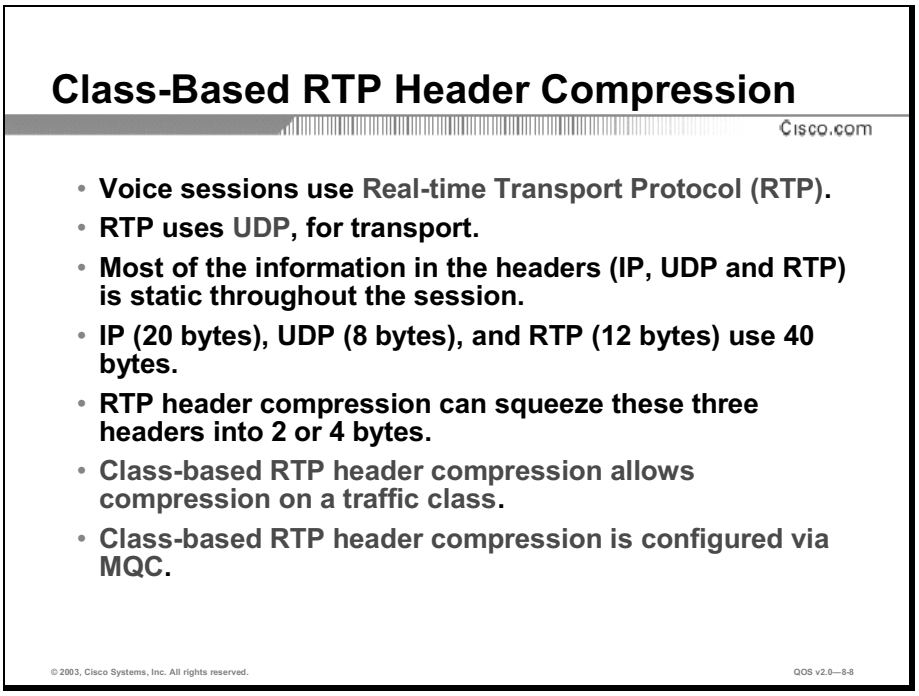

RTP is the Internet standard (RFC 1889) protocol for the transport of real-time data. It is intended to provide end-to-end network transport functions for applications that support audio, video, or simulation data over multicast or unicast network services. RTP is used in most VoIP applications to transport voice packets.

RTP includes a data portion and a header portion. The data portion of RTP is a thin protocol that provides support for the real-time properties of applications, such as continuous media, and includes timing reconstruction, loss detection, and content identification. RTP contains a relatively large-sized header. The 12 bytes of the RTP header, combined with 20 bytes of IP header and 8 bytes of the UDP header, create a 40-byte IP/UDP/RTP header. For compressed-payload audio applications, the RTP packet typically has a 20-byte to 160-byte payload depending on the audio compression codec. Given the size of the IP/UDP/RTP header combinations, it is inefficient to send the IP/UDP/RTP header without compressing it.

To avoid the unnecessary consumption of available bandwidth, the RTP header compression feature (cRTP) is used on a link-by-link basis. cRTP can reduce the header from 40 bytes to a 2- or 4-byte header, which significantly reduces delay on slow links.

Class-based RTP header compression enables RTP header compression on a per-class basis, when a class is configured within a policy map. Policy maps are created using the MQC. The policy maps are attached to an interface by using the **service-policy** command. The **service-policy** command gives customers the option of specifying either an input service policy (for input interfaces), or an output service policy (for output interfaces). For this feature, only output service policies can be specified.

Example: Class-Based RTP Header Compression

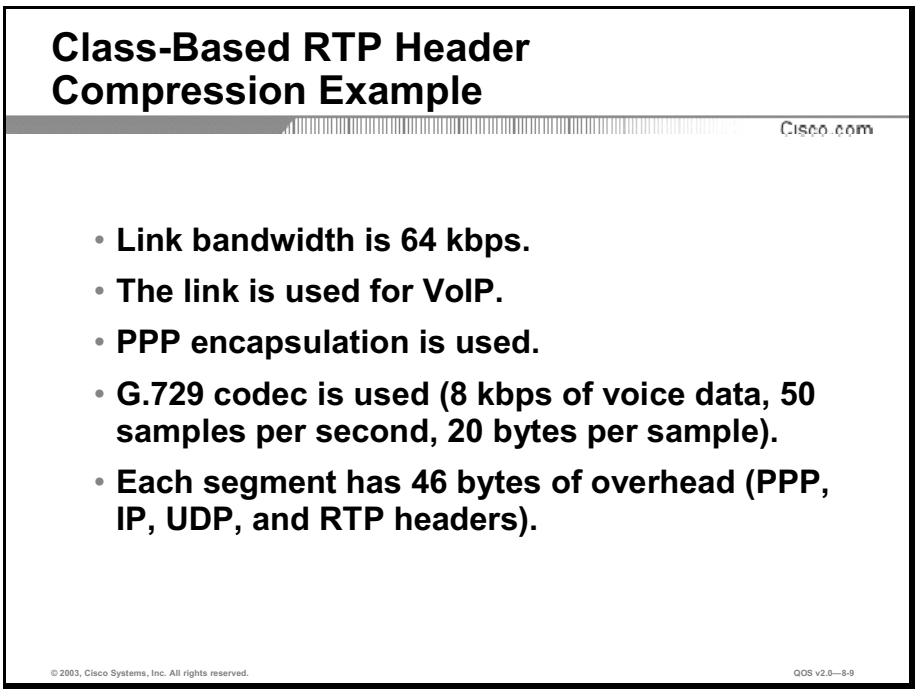

This example illustrates the benefits of RTP header compression on slow links.

A 64-kbps link is used to transport VoIP using PPP as the Layer 2 framing protocol.

For the case study application (voice using the G.729 audio compression codec), the payload size is 20 bytes. Because PPP has 6 bytes of frame header, the total header overhead is 6 + 20 + 8 + 12 = 46 bytes, counting the PPP, IP, UDP, and RTP headers.

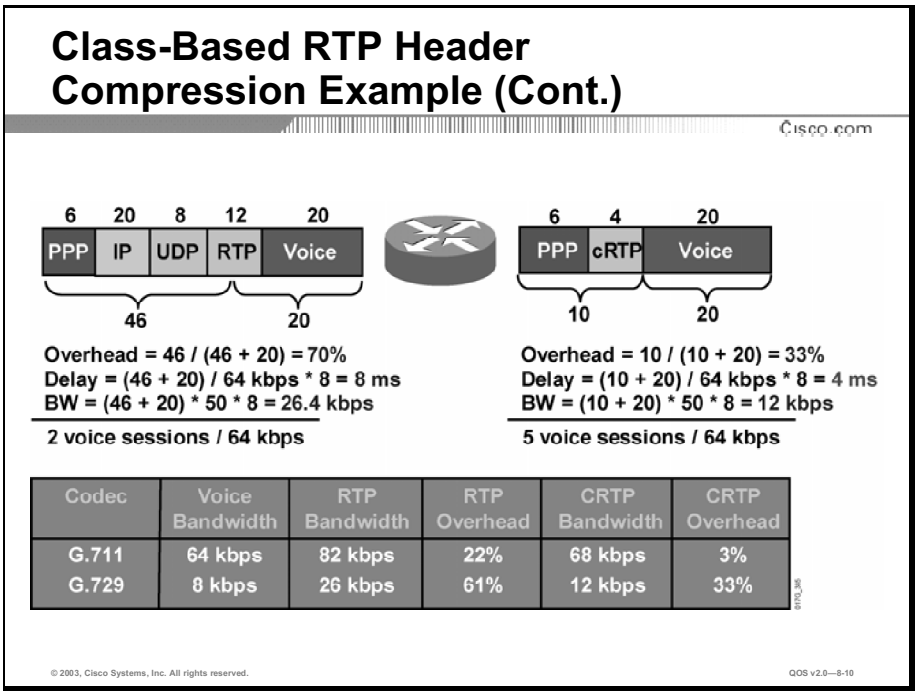

The figure shows the packet size before and after RTP header compression. The IP, UDP, and RTP headers are reduced to 4 bytes, resulting in 10 bytes of overall headers. The overhead is reduced from 70 percent to 33 percent, when small packets are used. Because of the packet size reduction, the serialization delay decreases from 8 ms to 4 ms, and the bandwidth that is used to transport a single voice call (using the G.729 codec) is reduced from 26.4 kbps (66 bytes/frame * 50 frames/sec * 8 bits/byte) to 12 kbps (30 bytes/frame * 50 frames/sec * 8 bits/byte). Therefore, a 64-kbps link can support up to two G.729 voice calls without cRTP, but up to five G.729 voice calls with cRTP.

The table in the figure shows how RTP header compression impacts performance when a different audio codec is used. For the traditional G.711 voice codec, RTP header compression still optimizes its transmission over slow links. However, the difference is more obvious when using advanced, low-bandwidth codecs.

Configuring Class-Based Header Compression

This topic identifies the Cisco IOS commands that are required to configure class-based header compression.

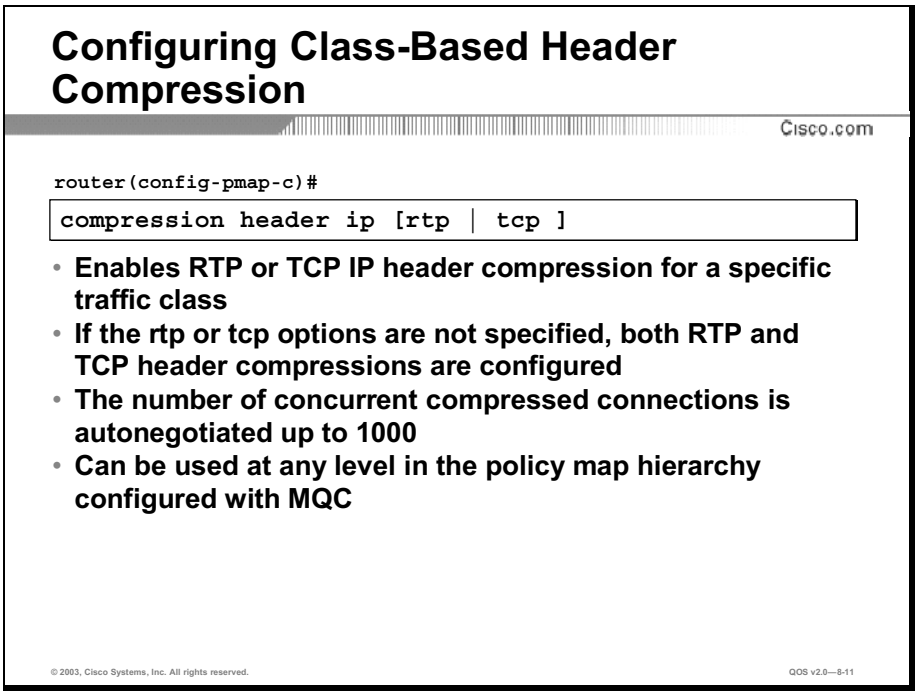

Class-based TCP and RTP header compression is configured within a policy map using the **compression header ip** command.

If you do not specify either RTP or TCP header compression (that is, you simply enter a carriage return after the **compression header ip** command), then both RTP and TCP header compression will be configured.

Example: Configuring Class-Based TCP Header Compression

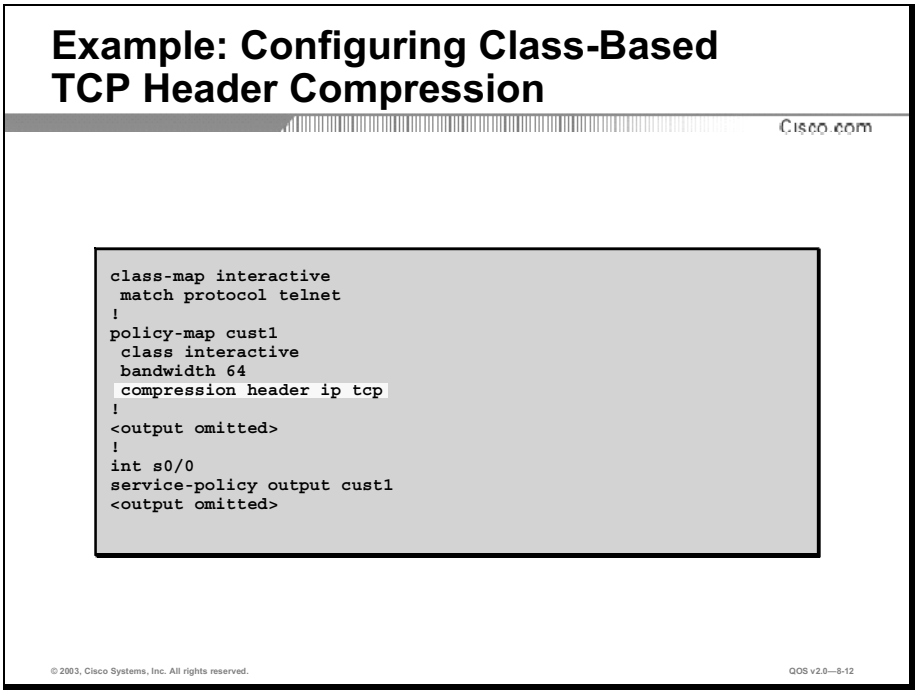

In the figure, the **compression header ip** command has been configured to use TCP header compression for a traffic class called “interactive”. The “interactive” traffic class is part of policy map called “cust1”. This “cust1” policy map is applied to the s0/0 interface in the outbound direction.

This policy provides a minimum bandwidth guarantee of 64 kbps for the “interactive” traffic class and will perform TCP header compression on the “interactive” traffic class (all Telnet packets in this example) leaving the s0/0 interface.

Example: Configuring Class-Based RTP Header Compression

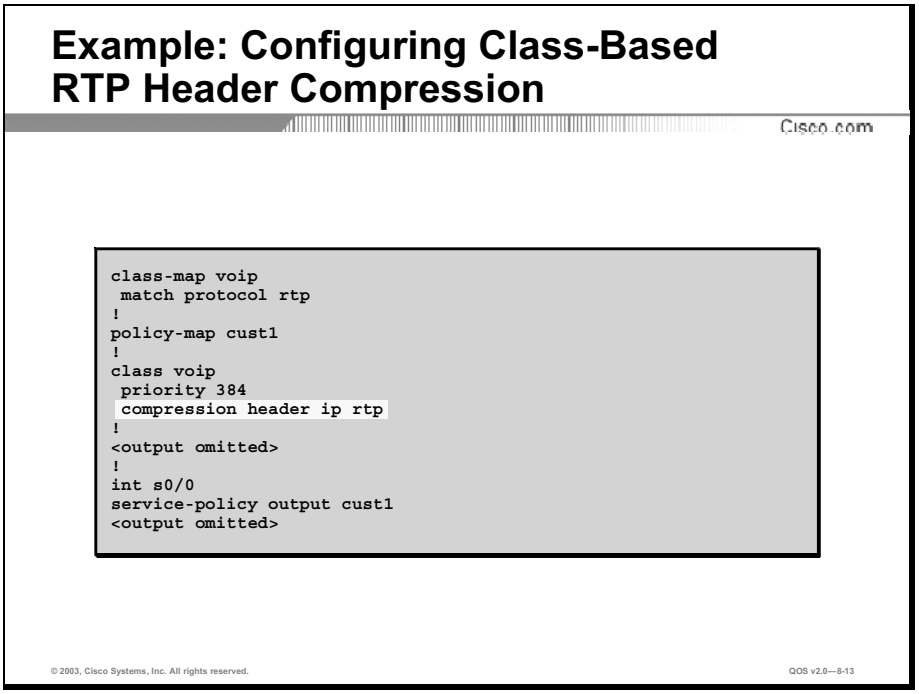

In the figure, the **compression header ip** command has been configured to use RTP header compression for a traffic class called “voip”. The “voip” traffic class is part of policy map called “cust1”. This “cust1” policy map is applied to the s0/0 interface in the outbound direction.

This policy provides a maximum bandwidth guarantee of 384 kbps for the “voip” traffic class and will perform RTP header compression on the “voip” traffic class (all RTP packets in this example) leaving the s0/0 interface.

Monitoring Class-Based Header Compression

This topic identifies the Cisco IOS commands that are used to monitor class-based header compression.

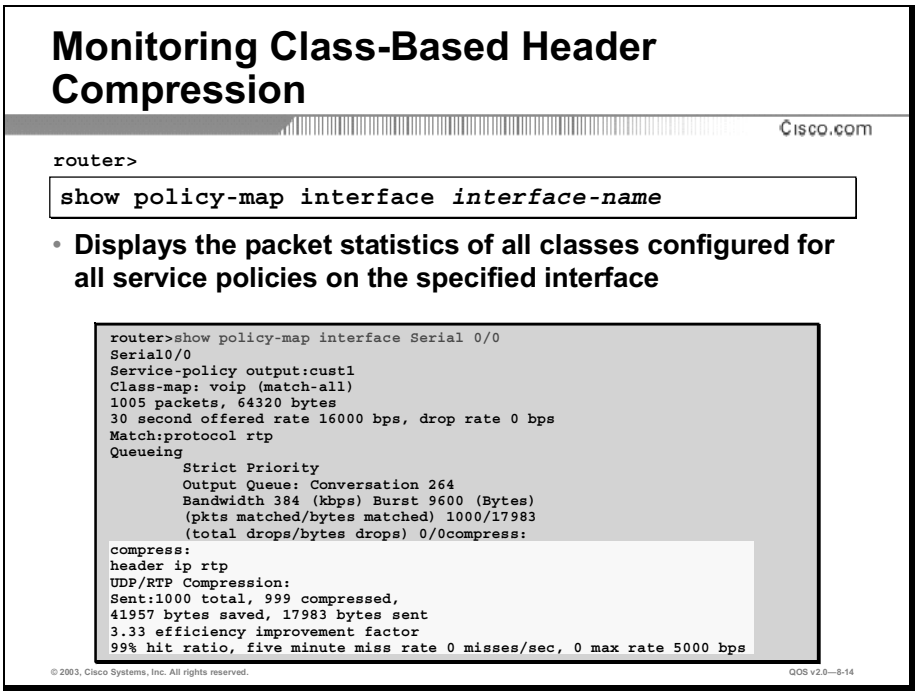

The **show policy-map interface** command output displays the type of header compression configured (RTP in this example), the interface to which the policy map called "cust1" is attached (serial 0/0), the number of packets sent, the number of packets compressed, the number of bytes saved, and the number of bytes sent.

Other statistical information provided in the output includes the "efficiency improvement factor," which indicates the percentage of increased bandwidth efficiency as a result of header compression. For example, an "efficiency improvement factor" of 3.33 means 330 percent efficiency improvement. The "hit ratio" is the percentage of packets that are found in the context database. In most instances, this percentage should be high. The "five-minute miss rate" is the number of traffic flows in the last five minutes that were *not* found in the context database. The rate is the actual traffic rate after the packets are compressed.

Summary

This topic summarizes the key points discussed in this lesson.

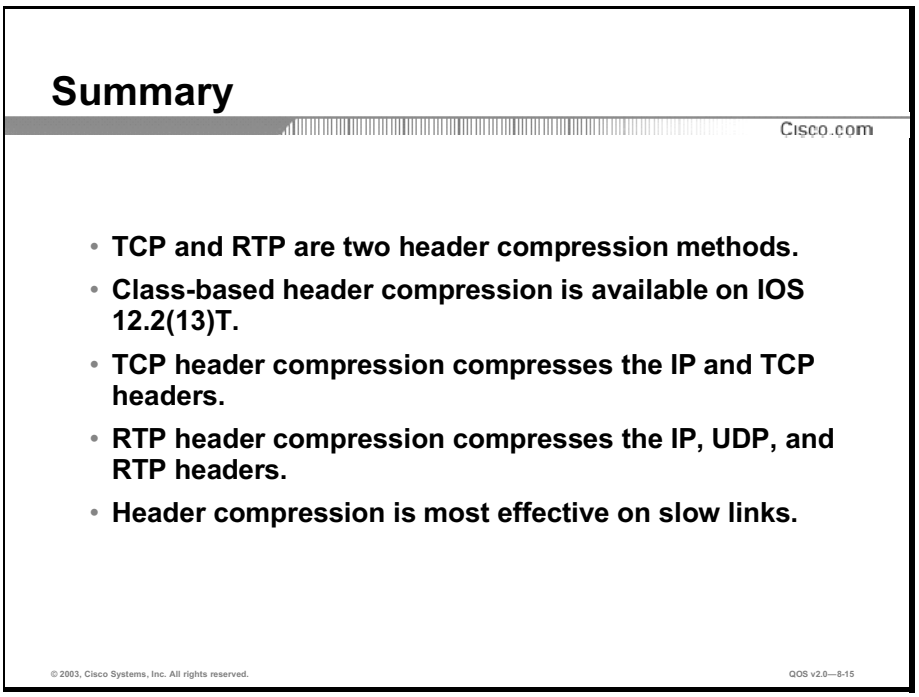

References

For additional information, refer to these resources:

- For more information on TCP and RTP header compression, refer to the following RFCs:
  - RFC 2508 (RTP header compression)
  - RFC 1144 (TCP header compression)
- For more information on Cisco class-based TCP and RTP header compression, refer to "Class-Based RTP and TCP Header Compression" at the following URL: http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122t/122t13/fthdrcmp.htm#wp1044256

Next Steps

For the associated lab exercise, refer to the following section of the course Lab Guide:

- Lab Exercise 8-1: Configuring Class-Based Header Compression

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) By default, if RTP or TCP is not specified with the **compression header IP** command, which header compression will be enabled?

A) none

B) RTP header compression only

C) TCP header compression only

D) both RTP and TCP header compression.

Q2) Typically, class-based TCP header compression should be performed on WAN links for which traffic class?

A) interactive traffic class

B) VoIP traffic class

C) bulk traffic class

D) RTP traffic class

Q3) RTP uses which transport layer protocol?

A) TCP

B) UDP

C) SPX

D) RTCP

Q4) cRTP can reduce the IP/UDP/RTP header to how many byte(s)?

A) 1 byte

B) 2 or 4 bytes

C) 4 or 8 bytes

D) 6 or 12 bytes

Q5) From the **show policy-map interface** command, what is the “hit ratio”?

A) the percentage of packets found in the context database

B) the percentage of packets matched by the class map

C) the percent of packets being policy routed

D) the percent of packets not compressed

Q6) What is incorrect about the following class-based RTP header compression configuration?

```
class-map voip
match protocol rtp
!
policy-map compress-me
class voip
      priority 384
      compression header ip rtp
!
Interface serial 0
      service-policy input compress-me
!
```

A) The **tcp** parameter is missing in the **compression header ip rtp** command.

B) The **priority** command can not be used in conjunction with the **compression header ip** command in the policy map.

C) The **service-policy** should be applied to the output not the input.

D) Nothing is incorrect about the configuration.

Quiz Answer Key

Q1) D

Relates to: Configuring Class-Based Header Compression

Q2) A

Relates to: Class-Based TCP Header Compression

Q3) B

Relates to: Header Compression Overview

Q4) B

Relates to: Class-Based RTP Header Compression

Q5) A

Relates to: Monitoring Class-Based Header Compression

Q6) C

Relates to: Configuring Class-Based Header Compression

Link Fragmentation and Interleaving

Overview

Because the hardware TxQ is a FIFO queue, having a large packet like a FTP packet in front of a small packet like a VoIP packet can cause excessive delay for the VoIP packet. LFI allows the large packets to be fragmented into smaller fragments, and then the small unfragmented packets are interleaved in between the fragments. LFI reduces the delay and jitter of small packets like VoIP packets over a WAN link.

This lesson discusses the configuration and monitoring of two different LFI mechanisms supported in Cisco IOS software, MLP with interleaving and FRF.12.

Relevance

Improving the efficiency of WAN links is important in that it can improve response times, reduce delays, and extend the longevity of some slower WAN links. To improve the WAN link efficiency, link fragmentation and interleaving techniques can be used to reduce the delays that are incurred by small interactive traffic types that are waiting in queues as large data frames are transmitted. LFI techniques are especially effective for VoIP packets and interactive TCP packets like Telnet packets.

Objectives

Upon completing this lesson, you will be able to configure LFI to improve bandwidth efficiency and reduce delay. This includes being able to meet these objectives:

- Identify the different options available for link fragmentation
- Given a list of link speeds and a specific delay requirement, determine the proper fragment size to use at each link speed and identify the typical delay requirement for VoIP packets
- Identify the Cisco IOS commands required to configure MLP with interleaving
- Identify the Cisco IOS commands used to monitor MLP with interleaving
- Explain when FRF.12 can be used and explain how FRF.12 affects VoIP packets
- Identify the Cisco IOS commands required to configure FRF.12
- Identify the Cisco IOS commands required to monitor FRF.12

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Knowledge of basic PPP operation and configuration on Cisco IOS routers
- Knowledge of basic Frame Relay operation and configuration on Cisco IOS routers

Outline

The outline lists the topics included in this lesson.

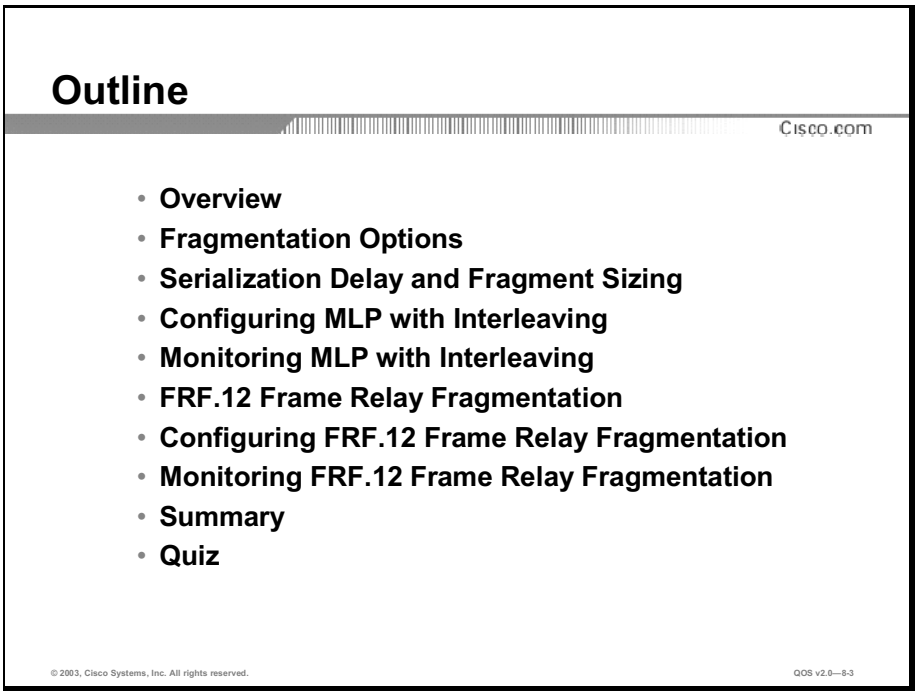

Fragmentation Options

This topic identifies the different options that are available for link fragmentation.

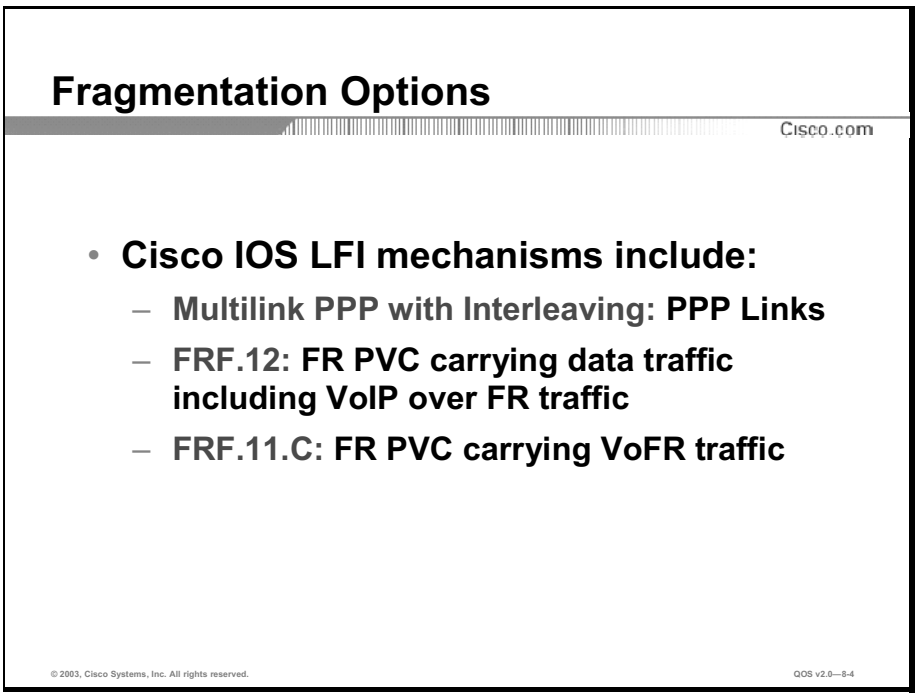

LFI is a Layer 2 technique, where all Layer 2 frames are broken into small, equal-size fragments, and transmitted over the link in an interleaved fashion. LFI reduces delay and jitter by expediting transfer of smaller frames through the hardware TxQ.

There are three LFI mechanisms implemented in Cisco IOS software:

- MLP with interleaving is by far the most common and widely used form of LFI.
- FRF.12 Frame Relay LFI is used with Frame Relay data connections.
- FRF.11 Annex C LFI is used with Voice over Frame Relay (VoFR).

Note: For configurations information on FRF.11.C, refer to the latest Cisco IOS documentation.

Serialization Delay and Fragment Sizing

This topic describes the proper fragment size to use on links based on their speed and identifies the typical delay requirements for VoIP packets.

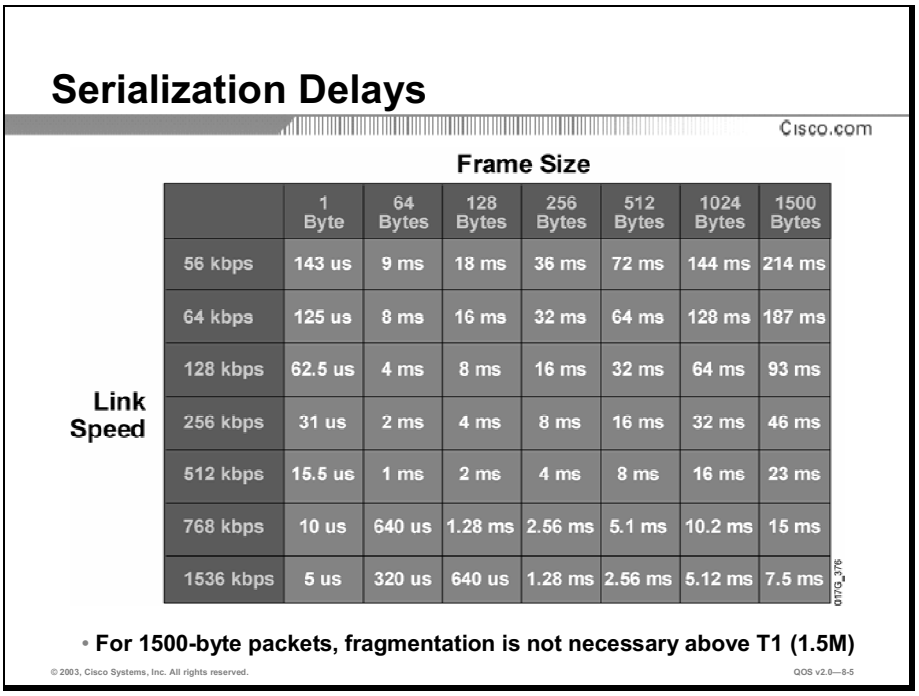

Serialization delay is the fixed delay required to clock a voice or data packet onto the network interface. Serialization delay is directly related to the link speed and the size of the packet.

This figure shows the serialization delay as a function of the link speed and packet size. For example, the serialization delay for a 1500-byte packet over a 64-kbps link will be:

- ([1500 bytes * 8 bits/byte]/64 kbps) = 187 ms

To ensure low delay and jitter for voice packets on slow links, the recommended standard goal for serialization delay is around 10 to 15 ms. Therefore, on a T1 or higher speed link, LFI is not necessary since the serialization delay for a 1500-byte packet is only 7.5 ms on a T1 link. T1 link has a bandwidth of 1536 kbps excluding the T1 framing overheads.

Fragment Size Recommendation for Voice

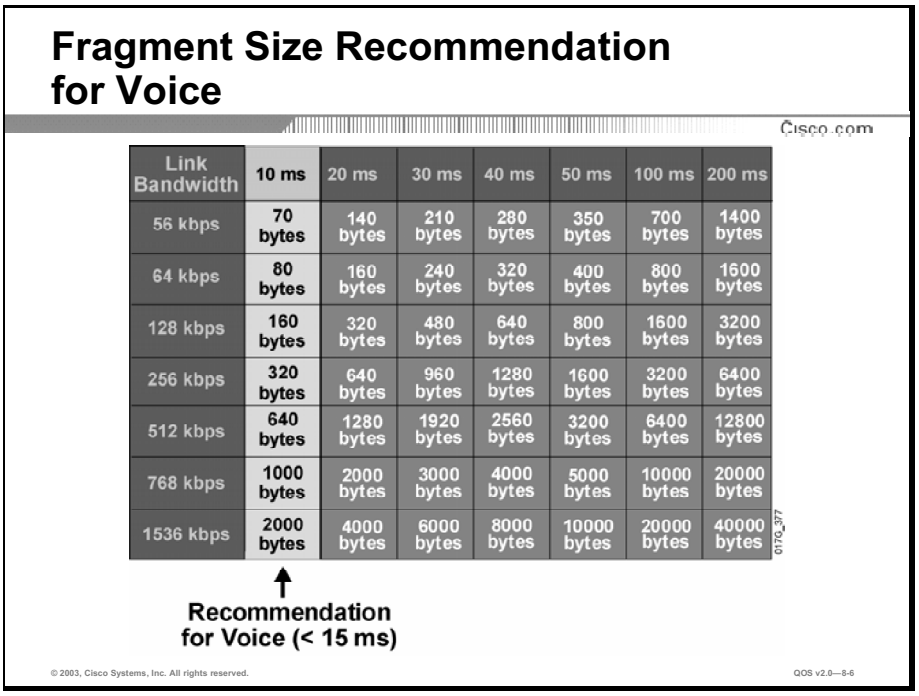

To meet the standard serialization delay goal of 10 to 15 ms to ensure low delay and jitter for voice packets, a fragment size of about 80 bytes per every 64 kbps of the clocking rate for the interface should be configured.

Depending on the LFI mechanism, the fragment size is either configured in bytes or in ms. For example, MLP LFI maximum fragment size is configured in ms while FRF.12 maximum fragment size is configured in bytes.

Example: Determining the Proper Fragment Size

- On a 64-kbps link, the proper fragment size to use is 80 bytes.
- On a 128-kbps link, the proper fragment size to use is 160 bytes (80 * 2).
- On a 192-kbps link, the proper fragment size to use is 240 bytes (80 * 3).
- On a 256-kbps link, the proper fragment size to use is 320 bytes (80 * 4).
- And so on.

Configuring MLP with Interleaving

This topic identifies the Cisco IOS commands required to configure MLP with interleaving.

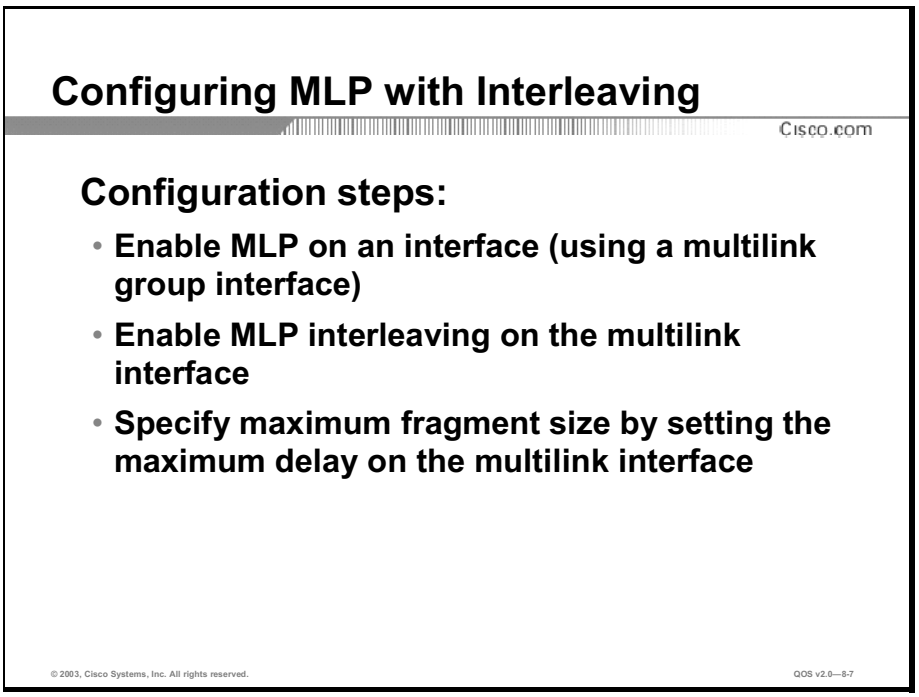

To configure MLP with interleaving, the following configuration steps must be performed:

Step 1 Enable MLP on a PPP interface.

Step 2 On the multilink interface, enable interleaving within MLP.

Step 3 In the multilink interface configuration, specify the maximum fragment size by specifying the maximum desired serialization delay in ms.

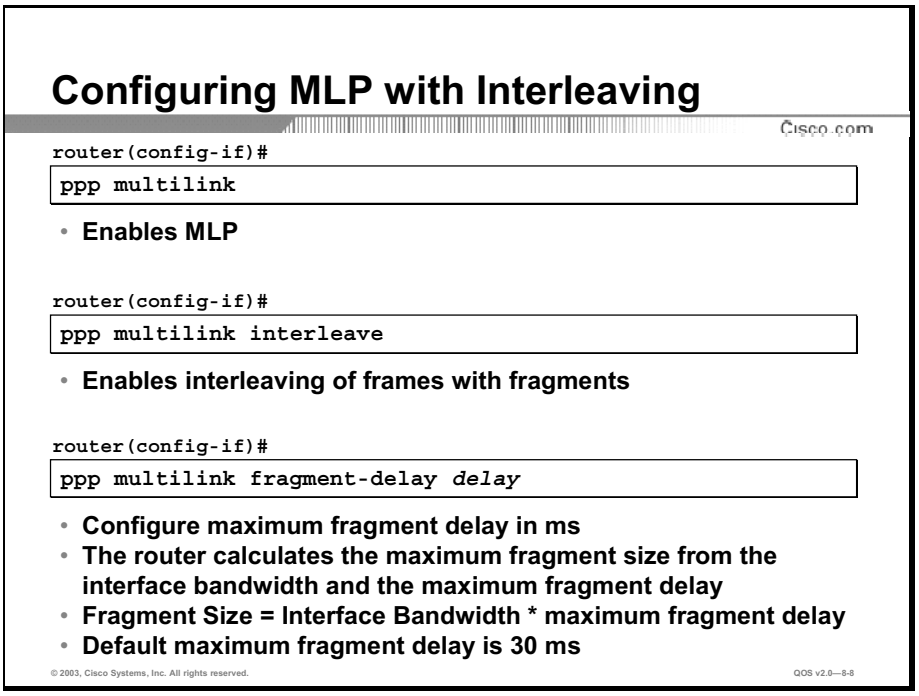

The **ppp multilink** command enables MLP on a PPP interface.

The **ppp multilink interleave** command enables interleaving of fragments within the multilink connection.

The **ppp multilink fragment-delay** command specifies the maximum desired fragment delay for the interleaved multilink connection. The maximum fragment size is calculated from the interface bandwidth and the specified maximum delay. The default is set at 30 ms. To support voice packets, a maximum fragment size of 10 to 15 ms should be used.

If distributed Cisco Express Forwarding (dCEF) is configured on a Versatile Interface Processor (VIP) interface, MLP with interleaving will run in distributed mode on the VIP.

Example: MLP with Interleaving

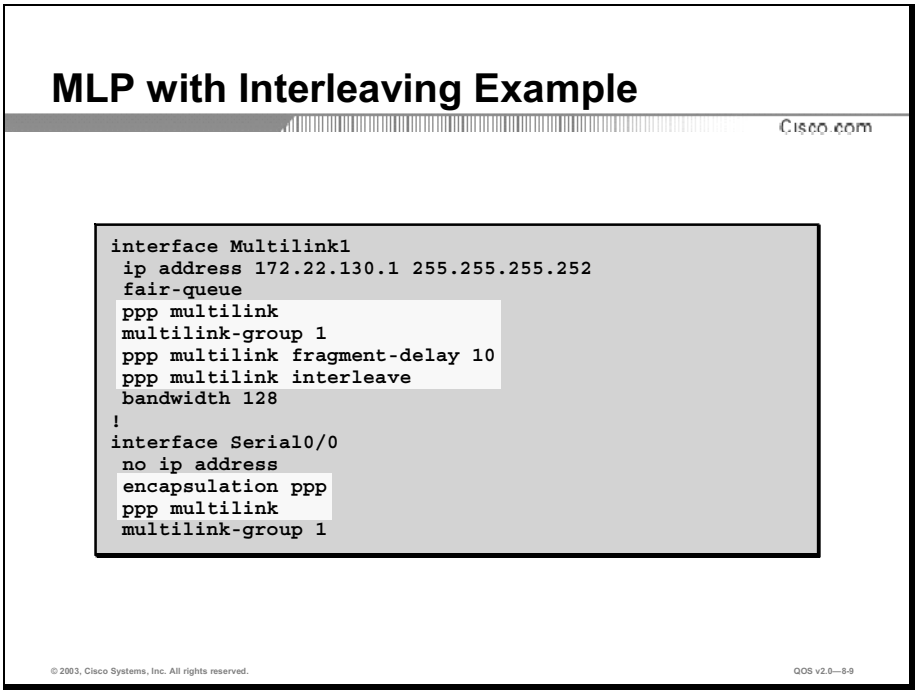

The figure shows an example configuration of MLP with interleaving on a multilink group interface. A non-default maximum desired delay of 10 ms is configured.

Monitoring MLP with Interleaving

This topic identifies the Cisco IOS commands that are used to monitor MLP with interleaving.

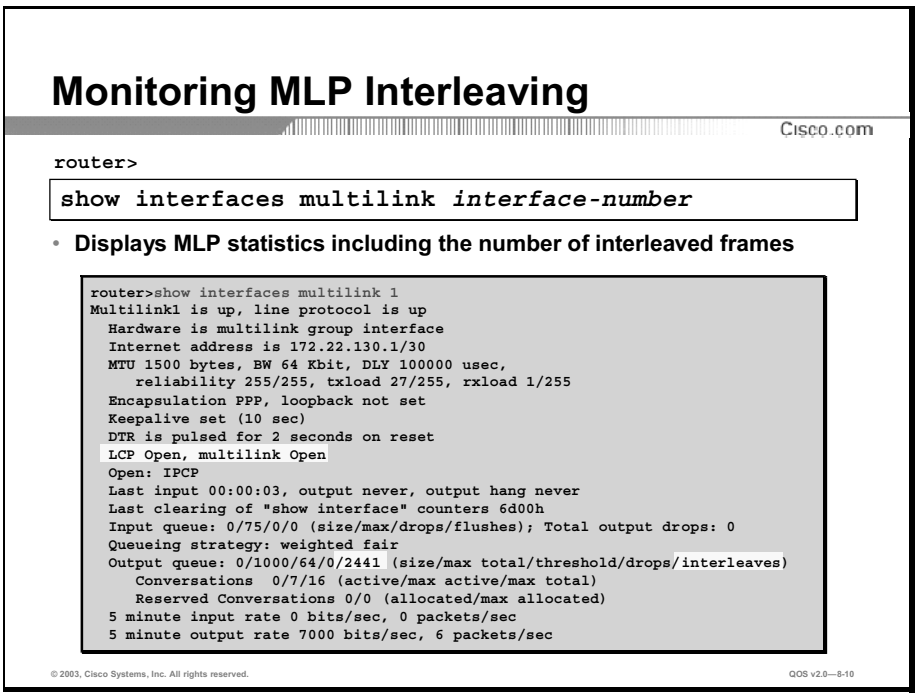

The **show interface multilink** command output includes MLP LFI statistics information and indicates whether MLP interleaving is enabled on the interface. Multilink should be in the open state along with link control protocol (LCP) and IP Control Protocol (IPCP).

Monitoring MLP Interleaving (Cont.)

Cisco.com

```
router#
debug ppp multilink fragments
```

- **Displays information about individual multilink fragments and interleaving events**

```
router#debug ppp multilink fragments
Multilink fragments debugging is on

Mar 17 20:03:08.995: Se0/0 MLP-FS: I seq C0004264 size 70
Mar 17 20:03:09.015: Se0/0 MLP-FS: I seq 80004265 size 160
Mar 17 20:03:09.035: Se0/0 MLP-FS: I seq 4266 size 160
Mar 17 20:03:09.075: Se0/0 MLP-FS: I seq 4267 size 160
Mar 17 20:03:09.079: Se0/0 MLP-FS: I seq 40004268 size 54
Mar 17 20:03:09.091: Se0/0 MLP-FS: I seq C0004269 size 70
Mar 17 20:03:09.099: Se0/0 MLP-FS: I seq C000426A size 70
Mar 17 20:03:09.103: Mul MLP: Packet interleaved from queue 24
Mar 17 20:03:09.107: Se0/0 MLP-FS: I seq C000426B size 70
Mar 17 20:03:09.119: Se0/0 MLP-FS: I seq C000426C size 70
Mar 17 20:03:09.123: Mul MLP: Packet interleaved from queue 24
Mar 17 20:03:09.131: Mul MLP: Packet interleaved from queue 24
Mar 17 20:03:09.135: Se0/0 MLP-FS: I seq C000426D size 70
Mar 17 20:03:09.155: Se0/0 MLP-FS: I seq C000426E size 70
```


The **debug ppp multilink fragments** command is a valuable troubleshooting tool when monitoring MLP LFI operations. This command outputs the result of every fragmentation operation, indicating whether the packets are fragmented into correct-sized fragments. This command should be used with extreme caution in a production environment, because of the amount of output that is created.

FRF.12 Frame Relay Fragmentation

This topic explains when FRF.12 can be used and how FRF.12 affects VoIP packets.

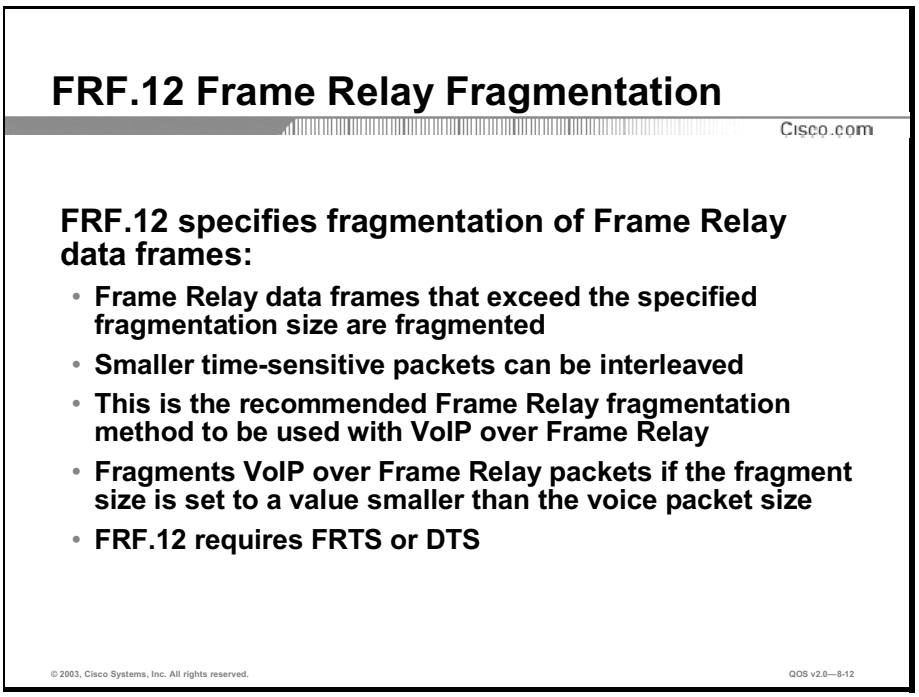

FRF.12 is the recommended fragmentation to be used with VoIP over Frame Relay while FRF.11c is the recommended fragmentation to be used with VoFR.

The FRF.12 Implementation Agreement defines FRF.12 fragmentation. The FRF.12 Implementation Agreement was developed to allow long data frames to be fragmented into smaller pieces and interleaved with real-time frames. In this way, real-time voice and non-real-time data frames can be carried together on lower-speed Frame Relay links without causing excessive delay and jitter to the real-time traffic like VoIP.

Because Frame Relay is a Layer 2 protocol, it has no way to tell which frame contains voice (VoIP) or data. Therefore, it will fragment all packets larger than the fragment size into smaller frames, including VoIP packets. In a VoIP over Frame Relay network, it is important to configure the fragment size on the data-link connection identifier (DLCI) so that VoIP frames will not get fragmented. For example, a G.711 VoIP packet without cRTP is 200 bytes long. For this DLCI, do not set the fragment size to less than 200 bytes.

Frame Relay permanent virtual circuits (PVCs) not configured for VoFR use normal Frame Relay (FRF.3.1) data encapsulation. If fragmentation is turned on for this DLCI, it uses FRF.12 for the fragmentation headers.

FRF.12 specifies three types of fragmentation, of which only one form is supported in Cisco IOS software:

- **End-to-end:** Supported in Cisco IOS software
  - Packets contain the FRF.12 fragmentation header
  - Occurs at the PVC level
  - LMI packets are not fragmented
- **UNI:** Not supported in Cisco IOS software releases
- **NNI:** Not supported in Cisco IOS software releases

FRF.12 is configured on a per-PVC basis. It should be noted that if one PVC or subinterface has fragmentation but another does not (on the same physical interface), large packets can and will delay smaller packets. For applications where fragmentation is needed, all PVCs or subinterfaces that carry large packets must have fragmentation turned on. FRF.12 fragmentation is configured within a Frame Relay map class using the **frame-relay fragment** command. The configured map class is then associated to the specific DLCIs. In addition, you must enable Frame Relay traffic shaping (FRTS) on the interface in order for FRF.12 fragmentation to work.

Cisco IOS release 12.1(5)T introduced a distributed version of FRF.12 Frame Relay fragmentation to use with DTS (distributed traffic shaping).

Configuring FRF.12 Frame Relay Fragmentation

This topic identifies the Cisco IOS commands that are required to configure FRF.12.

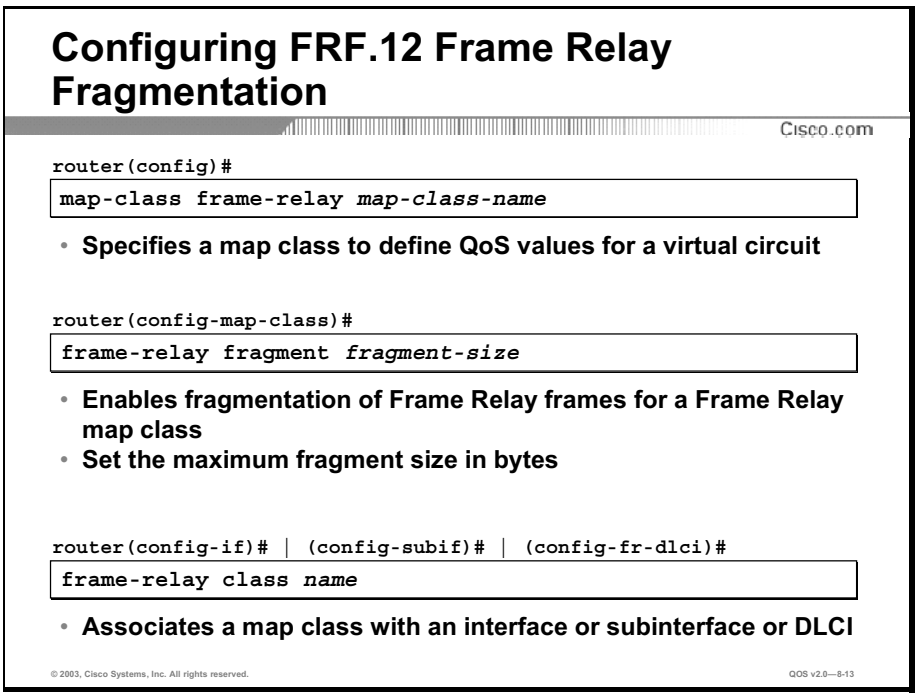

FRF.12 fragmentation is configured within the Frame Relay map class. The **frame-relay fragment** command sets the maximum fragment size in bytes. On an interface, the **frame-relay class** command applies the map class to the interface, subinterface, or a DLCI.

FRF.12 also requires FRTS to be enabled.

Example: FRF.12 Frame Relay Fragmentation

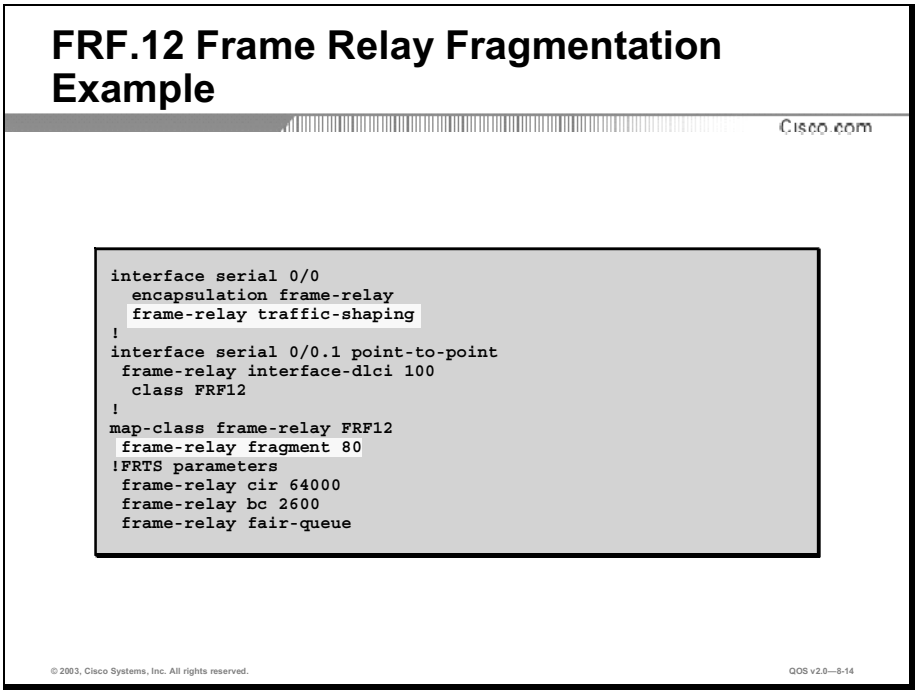

The figure shows a configuration example where FRF.12 fragmentation is applied to a data Frame Relay circuit configured on the serial 0/0.1 subinterface. The maximum fragment size is also set to 80 bytes. This would be used in a VoIP over Frame Relay environment.

FRF.12 also requires FRTS to be enabled. FRTS is enabled on the interface and the FRTS token bucket parameters are configured within the Frame Relay map class. In this figure, FRTS is enabled with a committed information rate (CIR) of 64 kbps, a committed burst (Bc) (normal burst size) of 2600 bits and uses weighted fair queuing (WFQ) as the shaping queue.

Monitoring FRF.12 Frame Relay Fragmentation

This topic identifies the Cisco IOS commands that are required to monitor FRF.12.

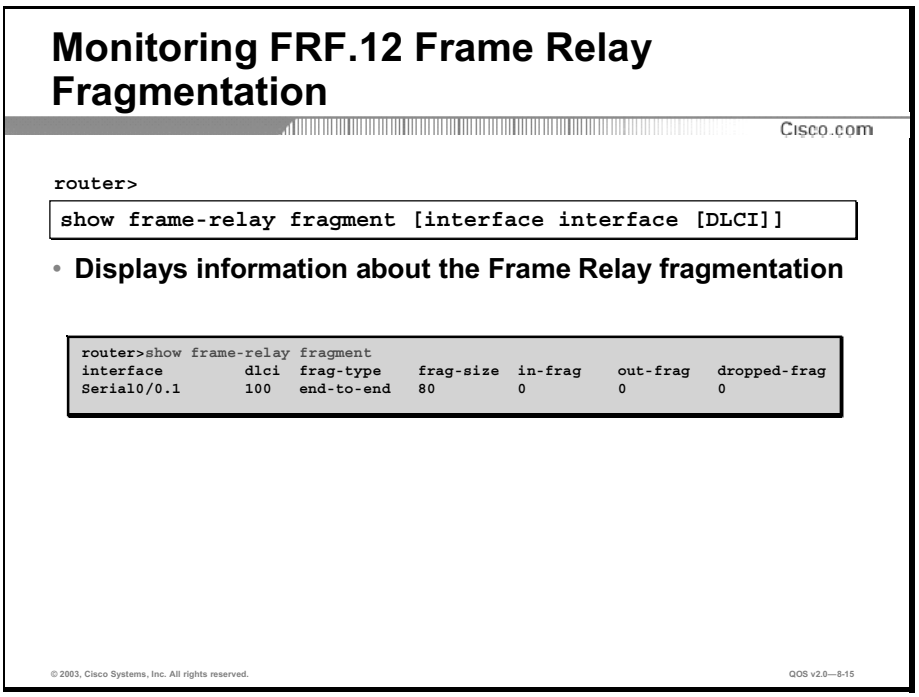

The **show frame-relay fragment** command displays information about the FRF.12 Frame Relay fragmentation process. The fragment type will always display end-to-end because this is the only type presently supported on Cisco IOS software. In addition to fragment type, the fragment size in bytes and associated DLCI is displayed.

Monitoring FRF.12 Frame Relay Fragmentation (Cont.)

Cisco.com

```
router>
show frame-relay pvc [interface interface] [dlci]
```

- **Displays statistics about PVCs for Frame Relay interfaces**

```
router>show frame-relay pvc 100

PVC Statistics for interface Serial0/0 (Frame Relay DTE)
DLCI = 100, DLCI USAGE = LOCAL, PVC STATUS = INACTIVE, INTERFACE = Serial0/0.1
<output omitted>
  Current fair queue configuration:
   Discard     Dynamic         Reserved
   threshold  queue count  queue count
    64            16                0
  Output queue size 0/max total 600/drops 0
  fragment type end-to-end fragment size 80
  cir 64000     bc   2600      be 0         limit 325    interval 40
  mincir 32000     byte increment 320   BECN response no  IF_CONG no
  frags 0         bytes 0         frags delayed 0         bytes delayed 0

  shaping inactive
  traffic shaping drops 0
```


The **show frame-relay pvc** command output includes settings related to the FRF.12 fragmentation process. This output shows the fragment size (80 bytes in this example), used on the Frame Relay PVC. The fragment type is end-to-end because Cisco IOS software currently only supports end-to-end FRF.12 fragmentation.

Summary

This topic summarizes the key points discussed in this lesson.

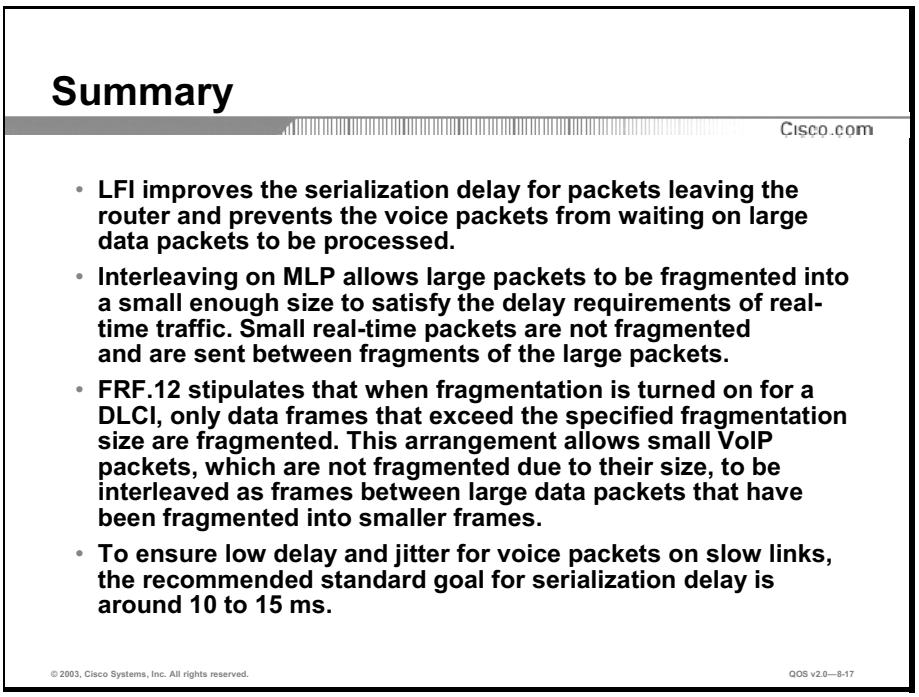

References

For additional information, refer to these resources:

- For more information on FRF.12 and DTS configuration, refer to “Frame Relay Traffic Shaping With Distributed QoS” at the following URL: http://www.cisco.com/warp/public/105/distribfrts.html#topic2%20URL
- For information on FRF.11c, refer to “VoFR Encapsulation and Fragmentation” at the following URLs:

  http://www.cisco.com/warp/public/788/voip/vofr_encap_frag_5733.html

  http://www.cisco.com/univercd/cc/td/doc/product/software/ios121/121newft/121limit/121e/121e2/vofvip/vofr3vip.htm

Next Steps

For the associated lab exercise, refer to the following section of the course Lab Guide:

- Lab Exercise 8-2: Configuring LFI

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) On a 64-kbps link, what is the proper fragment size (in bytes) to use?

A) 10 bytes

B) 20 bytes

C) 80 bytes

D) 160 bytes

Q2) How is the fragment size configured with multilink PPP and with FRF.12?

A) in bytes for FRF.12 and in milliseconds for multilink PPP

B) in milliseconds for FRF.12 and in bytes for multilink PPP

C) both in milliseconds

D) both in bytes

Q3) FRF.12 requires what else to be configured?

A) FRTS

B) FRF.11c

C) Multilink PPP

D) Frame Relay Inverse-Arp

E) LMI

Q4) With FRF.12, if the fragment size is set too small, what can happen to VoIP packets that are bigger than the fragment size?

A) The VoIP packets will be dropped.

B) The VoIP packets will not be fragmented because Cisco IOS software will automatically increase the fragment size.

C) The VoIP packets will not be fragmented and will be interleaved with the other fragments.

D) The VoIP packets will also be fragmented.

Q5) When enabling Multilink PPP, interleaving is also enabled by default.

A) true

B) false

Q6) Complete the following FRF.12 configuration:

```
interface serial 0/0
  encapsulation frame-relay
  _________________________
!
interface serial 0/0.1 point-to-point
 frame-relay interface-dlci 100
  ________________________
!
map-class frame-relay FRF12
 ________________________
!
!FRTS parameters
!
 frame-relay cir 64000
 frame-relay bc 2600
 frame-relay fair-queue
```

Quiz Answer Key

Q1) C

Relates to: Serialization Delay and Fragment Sizing

Q2) A

Relates to: Fragmentation Options

Q3) A

Relates to: Configuring FRF.12 Frame Relay Fragmentation

Q4) D

Relates to: FRF.12 Frame Relay Fragmentation

Q5) B

Relates to: Configuring MLP with Interleaving

Q6)

```
interface serial 0/0
  encapsulation frame-relay
  frame-relay traffic-shaping
!
interface serial 0/0.1 point-to-point
 frame-relay interface-dlci 100
  class FRF12
!
map-class frame-relay FRF12
 frame-relay fragment 80
!FRTS parameters
 frame-relay cir 64000
 frame-relay bc 2600
 frame-relay fair-queue
```

Relates to: Configuring FRF.12 Frame Relay Fragmentation

Module Assessment

Overview

Use this assessment to test what you learned in this module. The correct answers and solutions are found in the Module Assessment Answer Key.

Quiz: Link Efficiency Mechanisms

Complete the Quiz to assess what you have learned in the module.

Objectives

This activity tests your knowledge on how to meet these objectives:

- Explain how link efficiency mechanisms can be used to improve bandwidth efficiency and reduce delay
- Configure class-based TCP and class-based RTP header compression to improve bandwidth efficiency and reduce delay
- Configure LFI to improve bandwidth efficiency and reduce delay

Instructions

Complete these steps:

Step 1 Answer all questions in this quiz by selecting the best answer(s) to each question.

Step 2 Verify your results against the answer key located at the end of this section.

Step 3 Review the topics in this module that relate to the questions that you answered incorrectly.

Q1) With Layer 2 payload compression, what can be done to improve the compression/decompression delay of the router?

A) enable CEF switching
B) enable fast switching
C) use the stacker or predictor compression algorithm
D) use hardware-assisted compression

Q2) Which three of the following statements about FRF12 and FRF11.C are correct? (Choose three.)

A) FRF.12 fragmentation is used on DLCIs configured for VoIP.
B) FRF.12 fragmentation is used on DLCIs configured for VoFR.
C) FRF.11c fragmentation is used on DLCIs configured for VoIP.
D) FRF.11c fragmentation is used on DLCIs configured for VoFR.
E) FRF.12 fragmentation requires FRTS to be enabled.
F) Cisco IOS software supports UNI and NNI FRF.12 fragmentation and not end-to-end fragmentation.

Q3) When configuring multilink PPP with interleaving, where is the fragment size configured?

A) The fragment size in milliseconds is configured under the physical serial interfaces with PPP encapsulation and multilink PPP enabled.
B) The fragment size in milliseconds is configured under the logical multilink interface.
C) The fragment size in bytes is configured under the map-class.
D) The fragment size in bytes is configured within the policy-map.

Q4) Which two factors will influence the serialization delay? (Choose two.)

A) link speed

B) speed of the light in the media

C) router CPU processing power

D) packet size

Q5) To ensure good voice quality, what is the recommended fragment size?

A) 80 bytes per every 64 kbps of the clocking, which will result in a 10-ms serialization delay.

B) 40 bytes per every 64 kbps of the clocking, which will result in a 20-ms serialization delay.

C) 20 bytes per every 64 kbps of the clocking, which will result in a 10-ms serialization delay.

D) 120 bytes per every 64 kbps of the clocking, which will result in a 150-ms serialization delay.

Q6) Which two statements are true, based on the following show output? (Choose two.)

```
wg6r1#show policy-map interface s0/1
Serial0/1
Service-policy output: cust1
Class-map: voip (match-all)
300 packets, 202035 bytes
5 minute offered rate 2000 bps, drop rate 0 bps
Match: protocol rtp
Queueing
Strict Priority
Output Queue: Conversation 264
Bandwidth 384 (kbps) Burst 9600 (Bytes)
(pkts matched/bytes matched) 232/139975
(total drops/bytes drops) 15/15795
compress:
header ip rtp
UDP/RTP compression:
Sent: 285 total, 284 compressed,
9086 bytes saved, 176014 bytes sent
1.5 efficiency improvement factor
99% hit ratio, five minute miss rate 0 misses/sec, 0 max
rate 2000 bps
```

A) Both class-based TCP and RTP header compression are enabled for the VoIP traffic class.

B) IP payload compression is enabled for the VoIP traffic class.

C) LLQ is enabled for the VoIP traffic class.

D) Class-based RTP header compression is enabled for all RTP traffic.

Q7) From the **show interface multilink** output, which state(s) should be in the open state to indicate proper multilink PPP operation over an IP interface?

A) LCP, Multilink, and IPCP

B) LCP, IPCP, LFICP

C) LCP and MLP

D) LCP and IPCP

Q8) Which two of the following are correct regarding TCP and RTP header compression? (Choose two.)

A) Hardware assisted header compression is required to reduce the header compression delay.

B) TCP header compression compresses both the IP and TCP headers.

C) RTP header compression compresses the IP, UDP, and RTP headers.

D) TCP and RTP header compression can compress the respective protocol headers down to 10 bytes.

Q9) TCP header compression is most effective on which of the following applications?

A) VoIP

B) FTP

C) video streaming

D) Telnet

E) TFTP

Q10) RTP header compression is most effective on which of the following applications?

A) FTP

B) VoIP

C) TFTP

D) Telnet

E) HTTP

Scoring

You have successfully completed the quiz for this lesson when you earn a score of 80 percent or better.

Module Assessment Answer Key

Q1) D

Relates to: Link Efficiency Mechanisms Overview

Q2) A, D, E

Relates to: Link Fragmentation and Interleaving

Q3) B

Relates to: Link Fragmentation and Interleaving

Q4) A, D

Relates to: Link Fragmentation and Interleaving

Q5) A

Relates to: Link Fragmentation and Interleaving

Q6) C, D

Relates to: Class-Based Header Compression

Q7) A

Relates to: Link Fragmentation and Interleaving

Q8) B, C

Relates to: Class-Based Header Compression

Q9) D

Relates to: Class-Based Header Compression

Q10) B

Relates to: Class-Based Header Compression

Module Summary

This topic summarizes the key points discussed in this module.

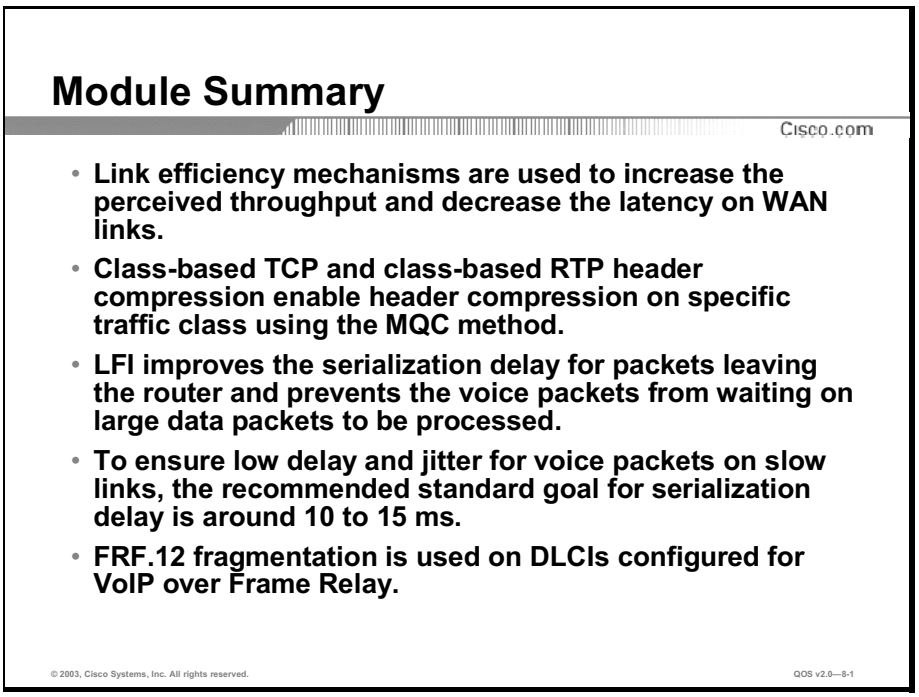

The different link efficiency mechanisms available in Cisco IOS software include class-based TCP and RTP header compression, Layer 2 payload compression, MLP with interleaving, and FRF.12 Frame Relay fragmentation. These QoS mechanisms are used on slower WAN links to improve the link efficiency by increasing the perceived throughput and decreasing the overall delay.

Class-based TCP and RTP header compressions are configured using the MQC method to reduce the protocol headers overhead. Header compression is most effective for packets with small payload size.

Layer 2 payload compression is not being implemented much because of the software compression delay overhead. Using a hardware compression module can help reduce the software compression delay.

For LFI, MLP with interleaving is supported on PPP interface while FRF.12 (for VoIP) and FRF.11.c (for VoFR) are supported on Frame Relay interface. LFI is used to reduce the serialization delay of large packets. Serialization delay is based on the link speed and packet size. To ensure high voice quality, the recommended serialization delay is 10 to 15 ms maximum. As a general rule, a fragment size of 80 bytes for every 64 kbps of link speed is recommended.

Module 9

QoS Best Practices

Overview

IP was designed to provide best-effort service for delivery of data packets and to run across virtually any network transmission media and system platform. The increasing popularity of IP has shifted the paradigm from “IP over everything,” to “everything over IP.” In order to manage the multitude of applications such as streaming video, voice over IP, e-commerce, enterprise resource planning (ERP), and others, a network requires quality of service (QoS) in addition to best-effort service. Different applications have varying needs for delay, delay variation (jitter), bandwidth, packet loss, and availability. These parameters form the basis of QoS. The IP network should be designed to provide the requisite QoS to applications.

To facilitate true end-to-end QoS on an IP network, the Internet Engineering Task Force (IETF) has defined two models: Integrated Services (IntServ) and Differentiated Services (DiffServ). IntServ follows the signaled QoS model, where the end-hosts signal their QoS need to the network. DiffServ works on the provisioned QoS model where network elements are set up to service multiple classes of traffic, with varying QoS requirements.

This module focuses on the implementation of the DiffServ model in service provider and enterprise networks. The first lesson provides examples of baseline traffic classifications. The second lesson provides a case study of a DiffServ implementation in a typical enterprise campus and service provider network.

Module Objectives

Upon completing this module, you will be able to correctly select the most appropriate QoS mechanisms for providing QoS using Cisco “best practices” in service provider and enterprise networks.

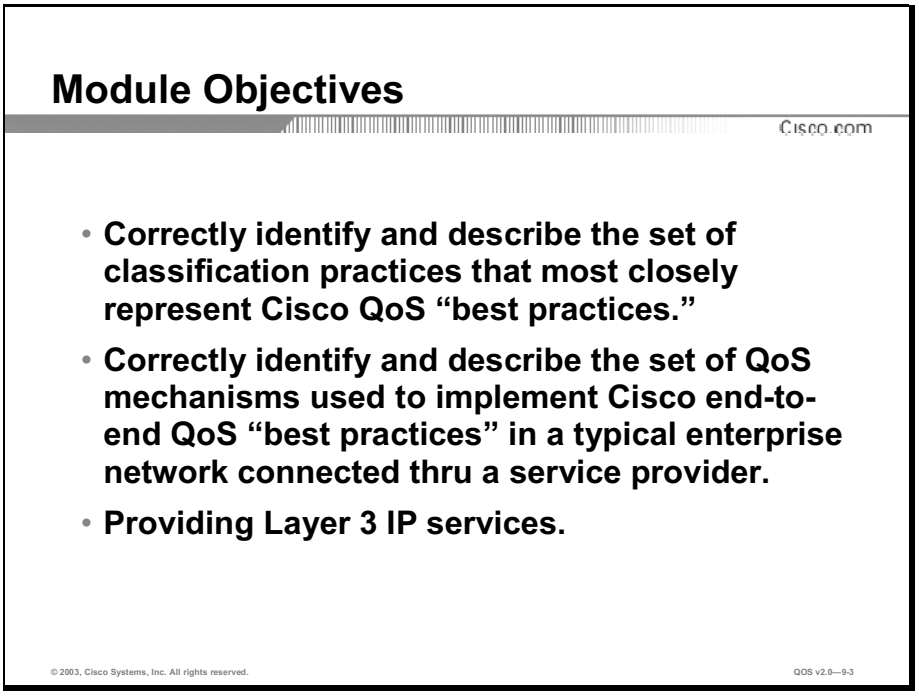

Module Outline

The outline lists the components of this module.

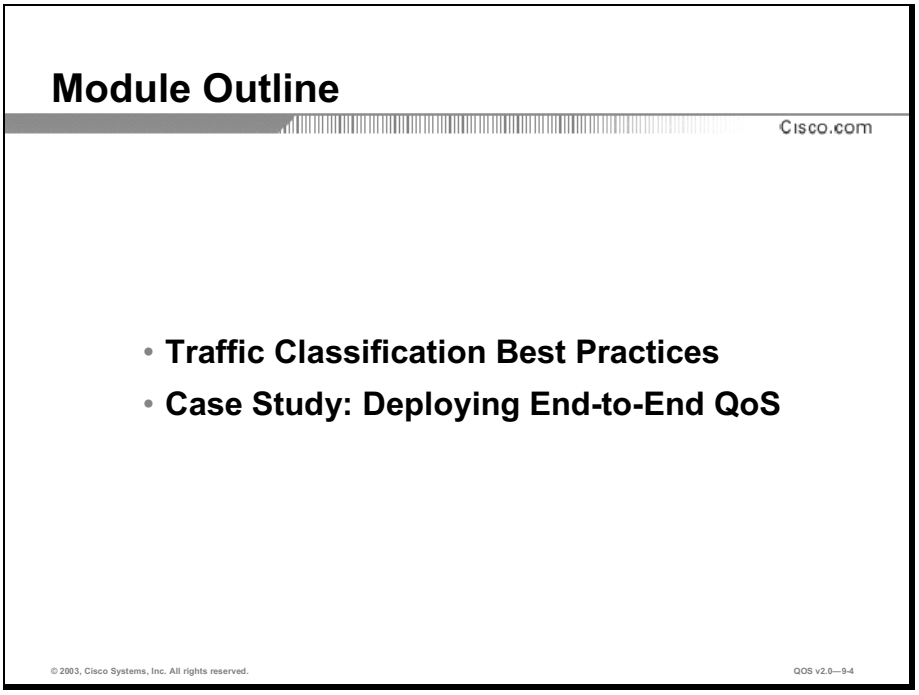

Traffic Classification Best Practices

Overview

Traffic classification entails using a traffic descriptor to categorize a packet within a specific group to define that packet and make it accessible for QoS handling on the network. Using proper traffic classification, the network traffic is partitioned into multiple priority levels or classes of service.

Relevance

Traffic classification is the pivotal key for proper QoS implementation. It is counterproductive to use too many discrete traffic classes. This is because the more classes that are defined, the less the distinction between the different service levels.

This lesson includes a baseline recommendation for traffic classification and marking in a typical enterprise and service provider environment.

Objectives

Upon completing this lesson, you will be able to correctly identify and describe the set of classification practices that most closely represent Cisco QoS “best practices.” This includes being able to meet these objectives:

- Define some of the key QoS best practices recommendations
- Explain the QoS requirements of the different application types
- List typical enterprise traffic classes then identify the delay, jitter, packet loss, and bandwidth requirements of each traffic class
- Map different enterprise application types into the appropriate PHB (DSCP)
- Describe how to map different enterprise traffic classes into appropriate service provider traffic classes

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- A good understanding of using the MQC configuration method to implement Cisco IOS QoS mechanisms
- A good understanding of how to configure low-latency queuing (LLQ) and class-based weighted fair queuing (CBWFQ)

Outline

The outline lists the topics included in this lesson.

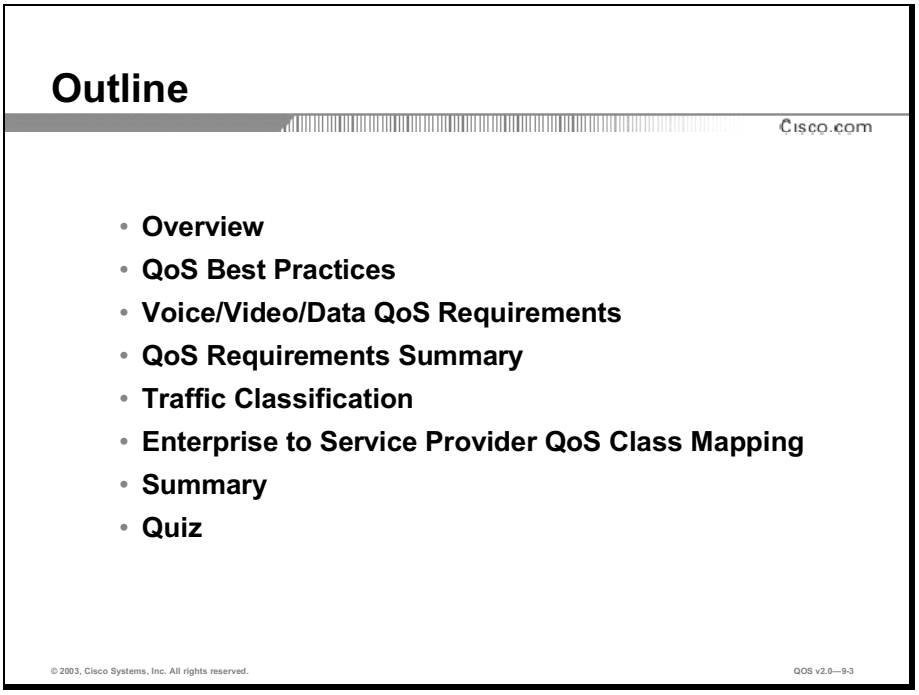

QoS Best Practices

This topic defines some of the QoS best practices recommendations.

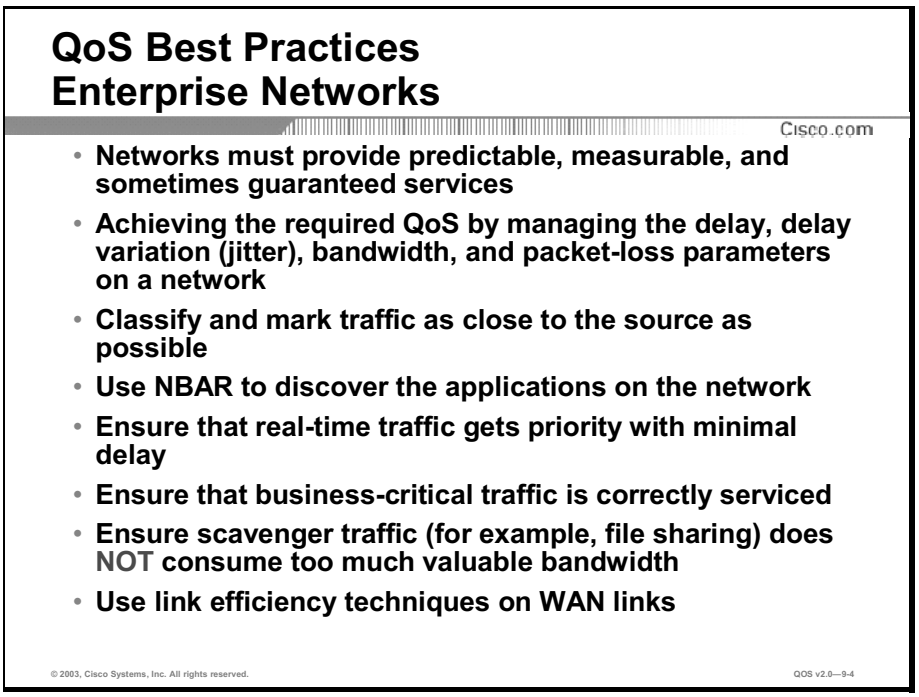

An IP network forms the backbone of any successful organization. These IP networks transport a multitude of applications and data, including high-quality video and delay-sensitive data (such as real-time voice). These applications stretch the IP network capabilities and resources, but also complement, add value, and enhance every business process. IP networks must provide secure, predictable, measurable, and sometimes guaranteed services. Achieving the required IP QoS by managing the delay, delay variation (jitter), bandwidth, and packet-loss parameters on an IP network becomes the secret to a successful end-to-end business solution.

IP QoS is the set of techniques and tools that are used to manage network resources. The figure lists some of the keys IP QoS best practices recommendations.

The first element of a QoS policy is to identify the traffic to be treated differently. By classifying the application traffic into different traffic classes, a baseline methodology is set to provide end-to-end QoS. DiffServ enables this classification by utilizing the differentiated services code point (DSCP) field. Using DiffServ, a properly designed network can deliver assured bandwidth, low latency, low jitter, and low packet loss for voice while simultaneously ensuring slices of available bandwidth to other traffic classes. Packets entering a DiffServ domain (collection of DiffServ routers) can be classified in a variety of ways:

- IP source and destination addresses
- Layer 4 protocol and port numbers
- Incoming interface
- MAC address
- IP precedence
- DSCP value
- Layer 2 information (such as Frame-Relay discard eligible [DE] bits, Ethernet 802.1p bits)
- The Cisco value-added mechanism network-based application recognition (NBAR)

It is best practice to classify and mark the traffic as close to the source of the traffic as possible.

In an enterprise environment, the IP QoS policies should allow critical business applications to receive requisite resources, while ensuring that other applications are not neglected. IP QoS policies should also ensure the quality of real-time traffic, like voice and video. It may also be important to prevent non-business-related network traffic (scavenger traffic) like file sharing traffic (that is, Kazaa, Napster, etc.) from taking up too much of the network bandwidth.

Network administrators often cannot justify continual upgrade of the link speeds in their networks. Cisco IOS software QoS features provide an alternative solution to link upgrade by managing the links efficiently to meet the application demands. Use QoS mechanisms such as Multilink PPP (MLP), link fragmentation and interleaving (LFI), compressed Real-Time Protocol (cRTP), CBWFQ, and LLQ to allow the most efficient distribution of the available bandwidth among the applications.

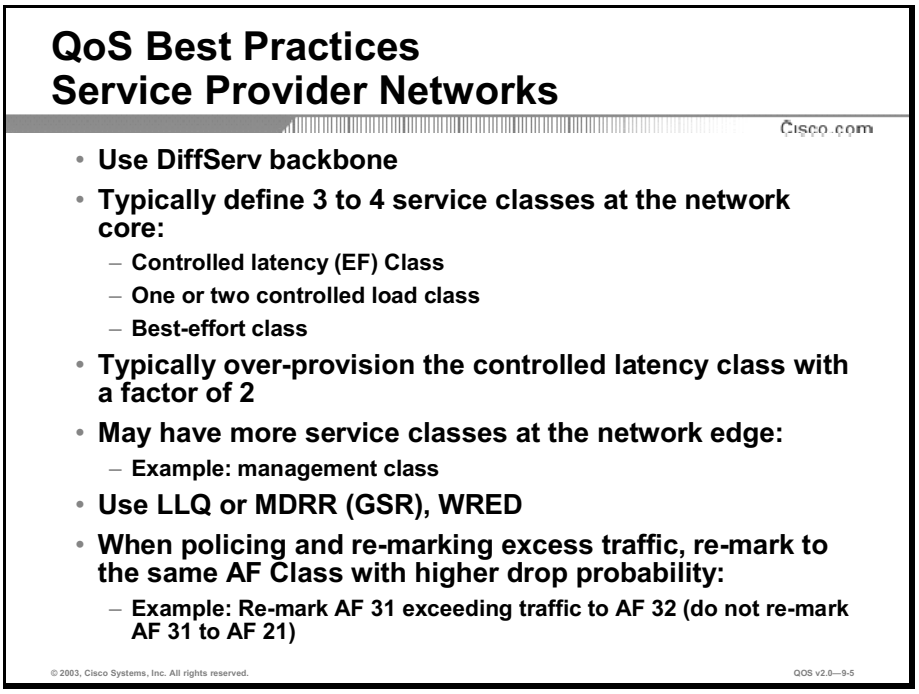

In an IP service provider environment, the recommended backbone architecture is to deploy a DiffServ backbone with a different over-provisioning ratio for the different traffic classes. Most service providers will typically classify their customer traffic into 3 or 4 traffic classes (controlled latency, controlled load, and best effort). At the network edge, the service provider may define additional traffic classes, for example, defining a management traffic class for all network management traffic (like Simple Network Management Protocol [SNMP] and Telnet traffic).

In the service provider core, where all the links are very high speed, the typical QoS mechanisms required are LLQ, modified deficit round robin (MDRR) (Gigabit Switch Router [GSR]), and weighted random early detection (WRED).

At the ingress to the service provider core, the customer traffic is typically policed by the service provider to ensure that the rate of customer traffic does not exceed the contractual rate. When re-marking the exceeding customer traffic, re-mark the DSCP value to the same Assured Forwarding (AF) class, but increase the drop probability. For example, re-mark AF 31 exceeding traffic to AF 32 or AF 33; do not re-mark AF 31 exceeding traffic to AF 21 or AF 11.

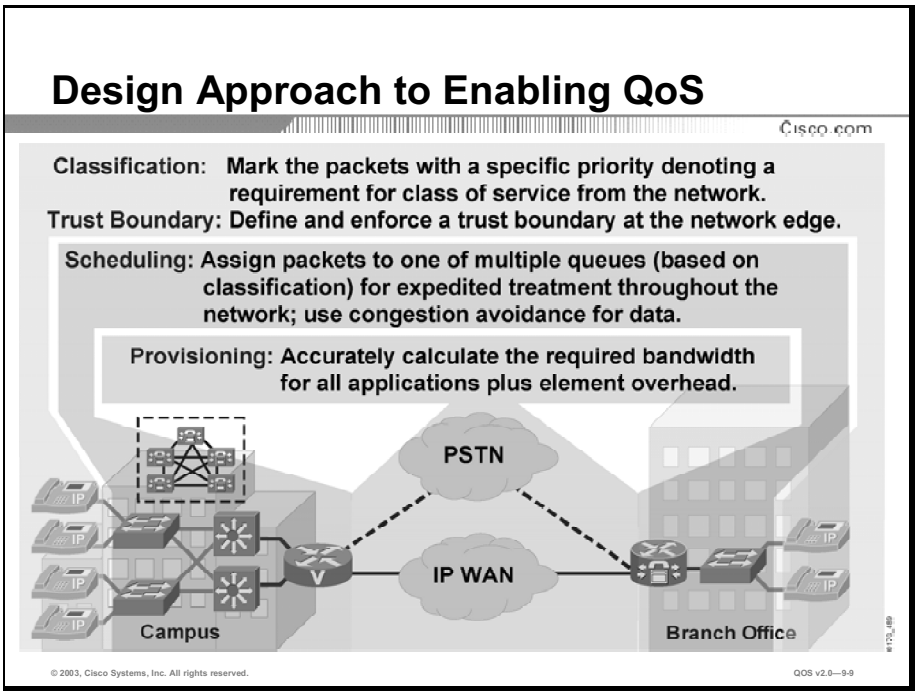

A set of QoS tools is needed to manage the delay, delay variation (jitter), bandwidth, and packet-loss parameters on a network. These QoS tools include the QoS classification, scheduling, and provisioning tools.

Only after traffic is identified can traffic policies be applied. The first element of a QoS policy is to identify the traffic to be treated differently. Classification must be consistent throughout the enterprise for it to be effective. QoS classification tools mark a packet with a specific priority. This marking is determined by examining:

- Layer 2 parameters like 802.1q class of service (CoS) bits, MAC address, Multiprotocol Label Switching (MPLS) experimental bits.
- Layer 3 parameters like IP precedence, DSCP, source/destination IP address.
- Layer 4 parameters like TCP or User Datagram Protocol (UDP) ports.

At the network edge, QoS markings may be accepted or rejected. This is referred to as the trust boundary. If packets are entering the network from a source that is trusted to set the packet marking properly, the packets are coming from a trusted device and can be left marked as-is. If the packets are entering the network from a device that cannot be trusted to properly mark packets, these packets must be re-marked. This is the location of the trust boundary.

The QoS scheduling tools are used to determine how a packet exits a node. Whenever packets enter a device faster than they can exit the device (as with speed mismatch), then a point of congestion can occur. Devices have buffers that allow for scheduling of higher-priority packets to exit sooner than lower-priority packets, which is commonly called queuing. Queuing is utilized only when a device is experiencing congestion and is bypassed when the congestion clears.

The QoS provisioning tools include the traffic policing and traffic shaping tools and the link efficiency mechanisms (LEM) tools. Traffic policers and shapers are the oldest form of QoS mechanisms. These tools have the same objectives: namely, to identify and respond to traffic violations. Policers and shapers usually identify traffic violations in a similar manner; however, their main difference is the manner in which they respond to violations:

- A policer drops excess traffic.
- A shaper delays excess traffic using a shaping queuing mechanism to hold packets and to shape the flow when the data rate is higher than expected.

LEMs are often used on WAN links to improve the perceived throughput and to reduce the delays on WAN links. Link efficiency mechanisms include LFI, payload compression, and header compression.

Implementing QoS is a means to use bandwidth efficiently, but not a blanket substitute for bandwidth itself. When the network is faced with ever increasing congestion, a certain point is reached where QoS alone will not solve bandwidth requirements. At such a point, nothing short of additional bandwidth will suffice.

Example: Cisco IOS QoS Tools Summary

In this course, you have learned about the various QoS tools available within Cisco IOS software.

The following list groups some of these tools into the proper categories.

- Classification tools:
  - Class-based marking
  - NBAR
  - QoS Policy Propagation on Border Gateway Protocol (QPPB)
  - QoS pre-classify
- Scheduling tools:
  - CBWFQ
  - LLQ
  - WRED
- Provisioning Tools:
  - Class-based shaping
  - Class-based policing
  - MLP interleaving
  - FRF.12
  - Class-based TCP and class-based Real-Time Transport Protocol (RTP) header compression

Voice/Video/Data QoS Requirements

This topic explains the QoS requirements of the different application types.

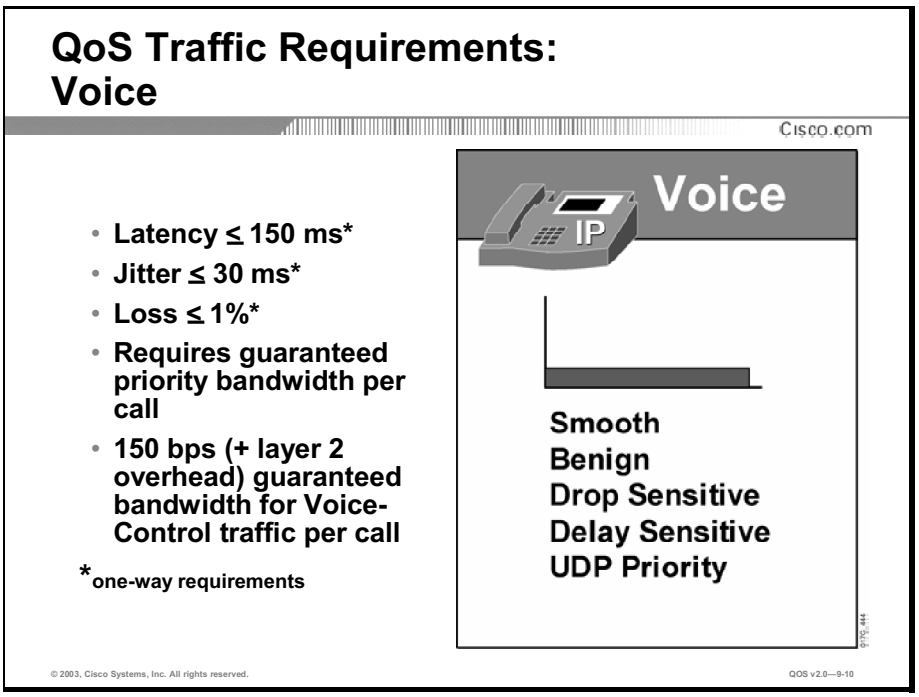

Voice quality is directly affected by all QoS quality factors: loss, delay, and jitter.

Loss causes voice clipping and skips. The industry standard coded algorithms that are used in Cisco digital signal processors (DSPs) can correct for up to 30 ms of lost voice. For example, if a 20 ms sample of voice payload is used per Voice over IP (VoIP) packet, only a single voice packet can be lost during any given time. If two successive voice packets are lost, the 30-ms correctable window is exceeded and voice quality begins to degrade.

Delay can cause voice quality degradation if it is above 200 ms. The International Telecommunication Union (ITU) standard for VoIP (G.114) states that a 150-ms one-way delay budget is acceptable for high voice quality.

With respect to jitter (delay variations), there are adaptive jitter buffers within Cisco IP telephony devices. However, these can usually only compensate for 20 to 50 ms of jitter.

To determine the bandwidth requirement per each VoIP call, several factors must be considered. These factors include the codec used, the sampling rate, and the Layer 2 protocol. Each VoIP call also requires 150 bps of bandwidth for the voice control (call-signaling) traffic.

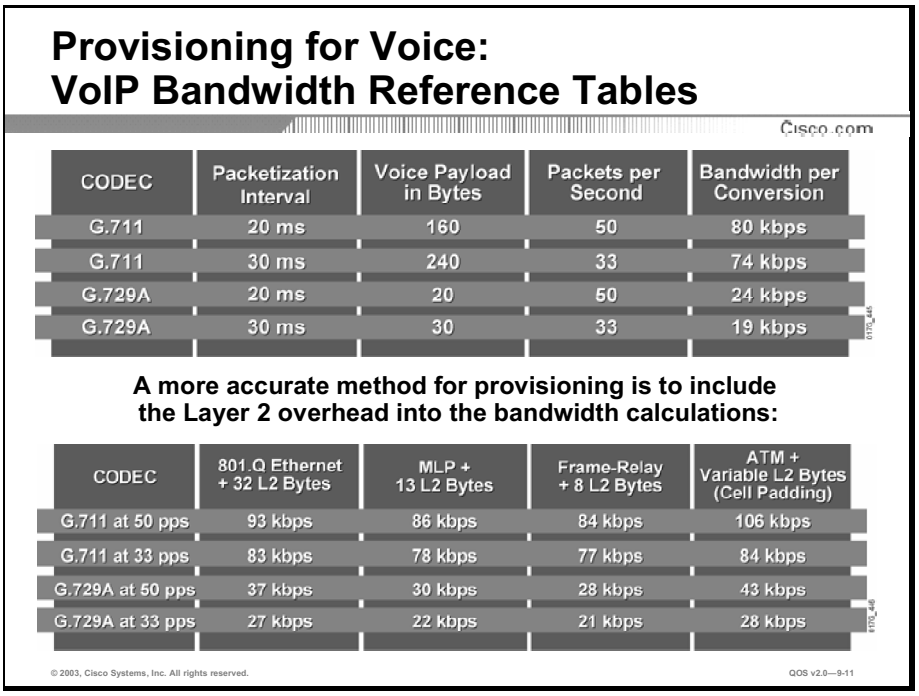

When addressing the bandwidth requirement for voice traffic, keep the following in mind:

- The required bandwidth per call depends on the codec, the packetization interval, and the Layer 2 protocol.
- The bandwidth required for the voice control (signaling) traffic, which is about 150 bps per call.

The bandwidth consumed by VoIP streams is calculated by adding the packet payload (in bits) and all the headers (in bits), then multiplying the total number of bits per packet by the per-second packet rate. This does not take into account the effect of any link efficiency tools such as RTP header compression and LFI. This graphic illustrates the VoIP bearer traffic bandwidth requirements for different codecs with different sampling rates. To have a more accurate method for VoIP bandwidth provisioning, also include the Layer 2 overhead (that is, preambles, headers, flags, cyclic redundancy checks (CRC)s, ATM cell-padding, and others) in bandwidth calculations.

Example: G.711 Voice Bearer Bandwidth Requirement Calculation

The following example shows how to calculate the VoIP bearer bandwidth requirement for a single VoIP call using a G.711 codec (Layer 2 overhead not included):

```
G.711             = 160 bytes payload size

Packet size       = payload size + IP/UDP/RTP headers
            = 160 bytes + 20 bytes + 8 bytes + 12 bytes
            = 200 bytes

Sampling Rate = 20 msec per sample = 50 samples per second

Bandwidth (bytes/sec) without Layer 2 overhead
= 200 bytes/packet x 50 packets/second
= 10000 bytes/second

Bandwidth (bits/sec) without Layer 2 overhead
= 10000 bytes/second  * 8 bits/byte
= 80000 bytes/second (80 kbps)
```


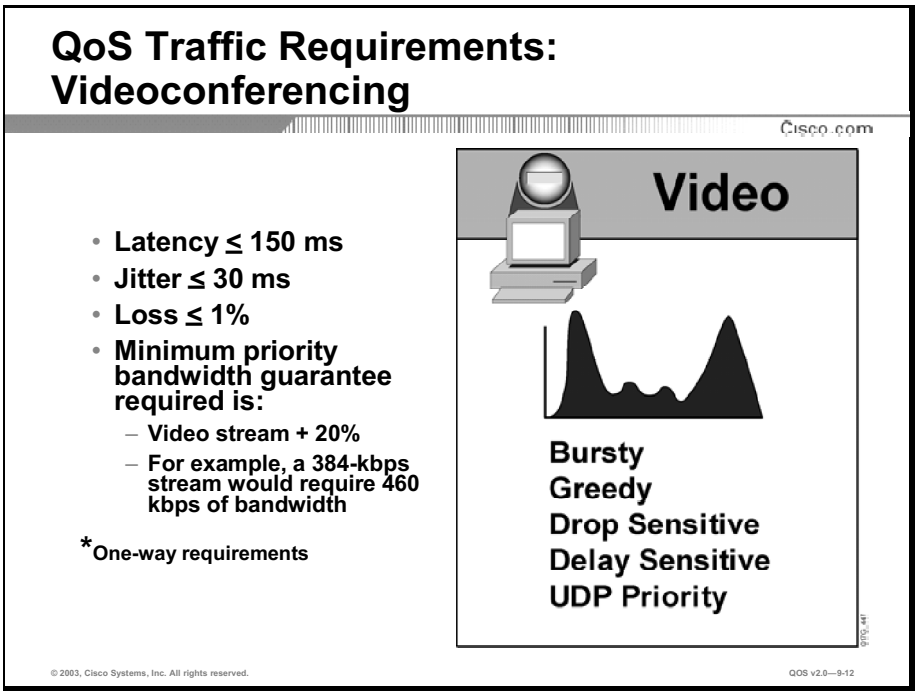

Videoconferencing has the same latency, jitter, and packet-loss requirements as voice, but the traffic patterns of videoconferencing are radically different from voice. For example, videoconferencing traffic has varying packet sizes and extremely variable packet rates. Because of its bursty nature, videoconferencing has two unique best-practice requirements in provisioning for strict-priority bandwidth:

- The LLQ should be provisioned to the stream rate plus 20 percent.
- The LLQ burst parameter should be provisioned to 30 KB per 384-kbps stream.

Compared to videoconferencing, streaming video applications have more lenient QoS requirements. Video streaming applications are more delay insensitive and are largely jitter insensitive due to application buffering. When provisioning for streaming video applications, take into account the video content distribution requirements. Video file distribution is similar to FTP traffic in nature and can have a major impact on network performance because of the large file size. Content distribution traffic should be managed to avoid impacting the network.

Example: Calculating the Bandwidth Requirement for a 384-kbps Videoconference Stream

The following example shows how to calculate the bandwidth requirement for a 384-kbps videoconferencing stream.

```
384 kbps + (20% x 384 kbps)    = 384 kbps + 76.8 kbps
                               = 460.8 kbps
```

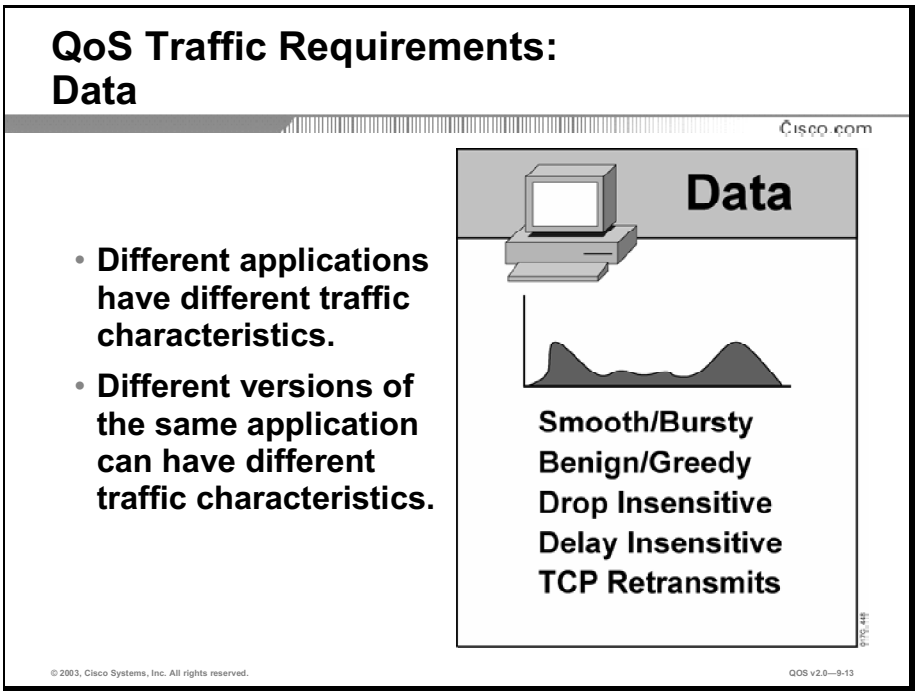

Because bandwidth requirements vary greatly from application to application (and even between versions of the same applications) it is not possible to provide a blanket rule for provisioning data bandwidth.

Traffic analysis and lab testing are required to ascertain bandwidth requirements for data applications.

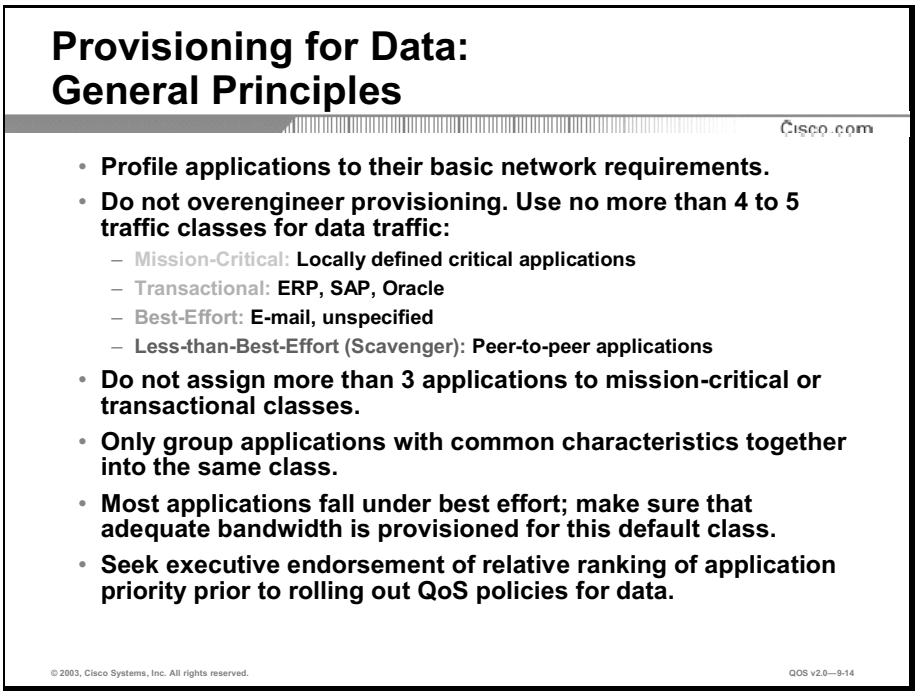

For data classification, a relative-priority model of no more than four to five traffic classes has been proven to work well in most enterprise environments: for example, a mission-critical, transactional, best-effort (default), and less-than-best-effort (scavenger) class. Each traffic class can be assigned a minimum bandwidth guarantee using CBWFQ. An example of using a relative-priority model between the different data traffic classes is:

- The mission-critical traffic class gets a minimum of 50 percent of the remaining bandwidth after the LLQ is serviced.
- The transactional traffic class gets a minimum of 20 percent of the remaining bandwidth after the LLQ is serviced.
- The best-effort (default/best effort) traffic class gets a minimum of 25 percent of the remaining bandwidth after the LLQ is serviced.
- The less-than-best-effort traffic class gets a minimum of 5 percent of the remaining bandwidth after the LLQ is serviced.

It is counterproductive to use too many discrete classes for data traffic. This is because the more classes that are defined, the less the distinction between service-levels. Additionally, if too many applications are assigned to the mission-critical class, then the overall effectiveness of QoS provisioning is dampened. Taken to an extreme, if all applications are assigned to the mission-critical class, then the end result is the same as when no QoS is provisioned at all. Mission-critical applications should be limited to those that directly contribute to the core operations of the business. These are usually highly interactive and are therefore sensitive to delay and loss. Examples of mission-critical applications include ERP applications, such as Oracle, SAP and PeopleSoft, as well as proprietary applications that were designed in-house.

When classifying the different network applications into traffic classes, it is very important to try to only group applications with the common characteristics and QoS requirements together into the same traffic class. For example, because interactive applications and bulk transfer applications have very different bandwidth and delay requirements, interactive and bulk traffic should not be classified into the same traffic class.

Example: RCL Enterprise

In this example, the RCL enterprise network administrator used NBAR and determined the major sources of network traffic are from the following applications:

- Structured Query Language (SQL), VoIP, FTP, HTTP, Domain Name System (DNS), Exchange, Napster

The RCL senior management team has determined that only SQL and HTTP traffic are mission-critical to the company.

RCL has decided to implement four traffic classes as follows:

- A real-time traffic class for VoIP traffic—all RTP traffic
- A mission-critical data traffic class for mission-critical traffic—SQL and HTTP traffic
- A bulk traffic class for business traffic that is not mission-critical—FTP, DNS, Exchange
- A best-effort traffic class for all best-effort traffic—Napster and all other traffic

The RCL network administrator has decided not to implement a scavenger traffic class to limit the Napster traffic at this time, but is constantly monitoring the Napster traffic. In the future, if more bandwidth is required for the business-related applications, a scavenger traffic class will be added and the Napster traffic will be classified into this scavenger traffic class.

QoS Requirements Summary

This topic lists typical enterprise traffic classes and identifies the delay, jitter, packet-loss, and bandwidth requirements of each enterprise traffic class.

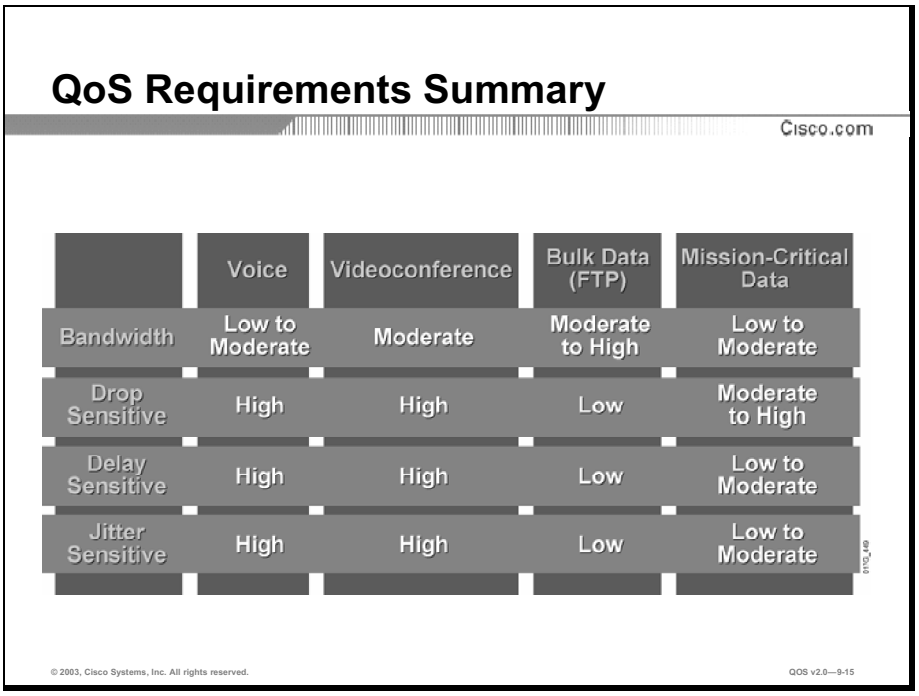

This table summarizes the key QoS requirements (bandwidth, packet loss, delay and jitter) for some of the major different applications category.

Example: QoS Requirements of the Major Applications Category

VoIP and videoconferencing applications share the same low drop, low delay, and low jitter requirements. For example, a latency of less than 150 ms, a jitter of less than 30 ms, and a packet loss of less than 1 percent are typically required.

Mission-critical applications like ERP applications have different bandwidth requirements, depending on the application and even the particular version of the application. In general, mission-critical applications are more drop, delay, and jitter sensitive than bulk traffic.

Bulk applications like FTP are less sensitive to drop, delay, and jitter, but generally require more bandwidth than real-time traffic, such as voice traffic.

Traffic Classification

This topic maps different enterprise applications types into the appropriate per-hop behavior (PHB) (DSCP).

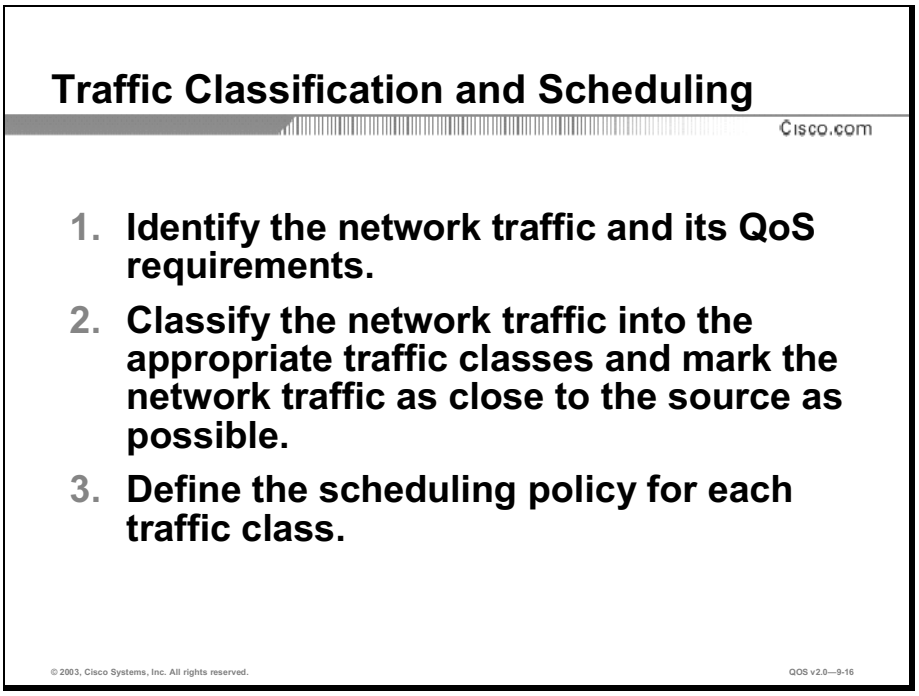

After the QoS requirements for the different applications running on the network have been determined, the next step is to determine how to classify the traffic into multiple service classes. It is best practice to classify and mark traffic as close to the source of the traffic as possible. At the enterprise campus LAN, classification and marking is typically done at the access or distribution layer. Both Cisco IOS routers and IOS Catalyst switches have class-based marking capabilities. NBAR is also supported on Cisco IOS routers and some IOS Catalyst switches.

After the traffic classification requirements have been determined, the next step is to establish the proper scheduling policy for each traffic class. For example, at the WAN edge router where the traffic exits the enterprise campus LAN onto the WAN, it is best practice to implement LLQ/CBWFQ to provided congestion management.

Enterprise WAN Edge Bandwidth Allocation Example

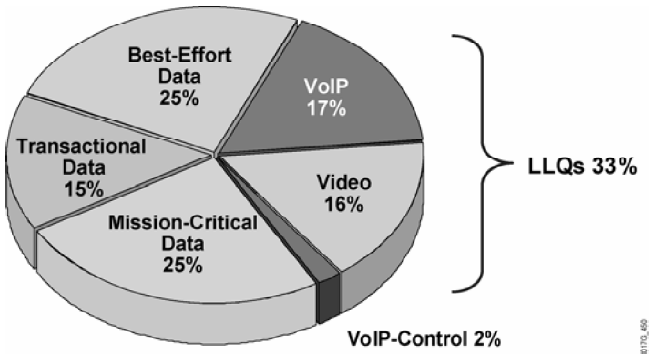

The amount of maximum bandwidth to be allocated for the LLQ class depends on many factors such as the router platform, the traffic profiles, and QoS features enabled.

The following design principles apply when determining bandwidth provisioning requirements for combining voice, video, and data over a single WAN link:

- Video requires 460 kbps of LLQ for each 384 kbps stream (stream-rate plus 20 percent).
- VoIP also requires LLQ bandwidth. The actual amount of bandwidth depends on the number of IP Phones, codecs used, sampling-rates, and Layer 2 protocol.
- The ratios of voice and video within the LLQ will vary, depending on the voice and video traffic requirements.
- As a general rule, the total LLQ provisioning is recommended to be ≤ 33 percent of the link capacity. This maximum number varies depending on many factors, such as the router platform and the traffic profiles.
- In a relative-priority model, mission-critical traffic should have greater bandwidth guarantees than the other data classes.
- Because most traffic may end up in the default (best-effort) class, ensure there is enough bandwidth left over for the best-effort traffic class.
- By default, the sum of all bandwidth guarantees should be ≤ 75 percent of the link bandwidth. For example, if the link bandwidth is 10 Mbps, the maximum available bandwidth will be 7.5 Mbps.

Example: LLQ Bandwidth Allocation

The example in the graphic shows a LLQ bandwidth allocation example between six different traffic classes:

- **VoIP bearer traffic class:** The VoIP bearer traffic class has a maximum bandwidth guarantee of 17 percent of the link bandwidth.
- **Video traffic class:** The video traffic class has a maximum bandwidth guarantee of 16 percent of the link bandwidth.
- **VoIP control traffic class:** The VoIP control traffic class has a minimum bandwidth guarantee of 2 percent of the link bandwidth.
- **Mission-critical data class:** The mission-critical traffic class has a minimum bandwidth guarantee of 25 percent of the link bandwidth.
- **Transactional data class:** The transactional traffic class has a minimum bandwidth guarantee of 15 percent of the link bandwidth.
- **Best-effort (default) class:** The best-effort (default) traffic class has a minimum bandwidth guarantee of 25 percent of the link bandwidth.

Note: In this example, the max reservable bandwidth on the link is set to 0 percent so up to 100 percent of the link bandwidth can be guaranteed between the different classes.

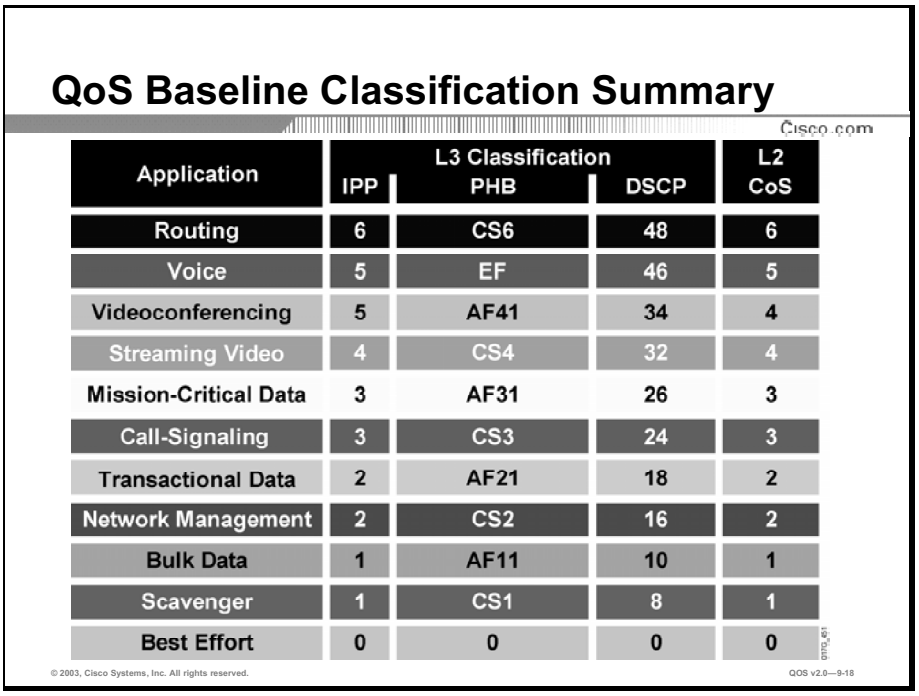

This figure illustrates a QoS baseline recommendation for traffic classification and markings. In this example, the enterprise traffic is classified into 11 traffic classes including:

- Five data classes (mission critical, transactional, bulk, scavenger, best effort)
- One class for network management traffic
- One class for routing protocol traffic
- One class for voice traffic
- One separate class for voice call-signaling traffic
- One class for videoconferencing traffic
- One class for streaming video traffic

In the figure, the recommended IP precedence value, the PHB, the DSCP value, and the Layer 2 CoS markings are shown for each traffic class. The markings of some of the traffic classes are as follows:

- Voice is marked with DSCP 46 and CoS 5.
- Voice call-signaling traffic is marked with Class Selector (CS) 3 and CoS 3.
- Videoconferencing (AF 41) is marked differently than streaming video (CS 4) because the QoS requirements are different for videoconferencing than for streaming video.
- Mission-critical data is marked with AF 31 and CoS 3.
- A scavenger class is used and is marked with CS 1 and CoS 1.
- Routing traffic is marked with CS 6 and CoS 6. Cisco IOS software assigns an IP precedence of 6 to routing protocol packets on the control plane.

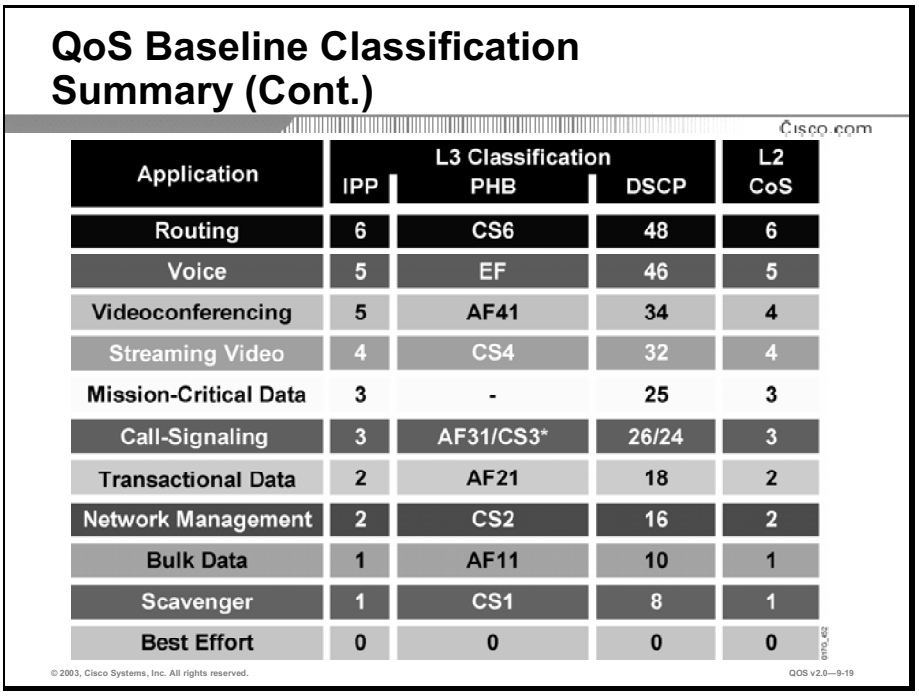

A marking of CS 3 is recommended for voice call-signaling traffic. However, Cisco IP Phones and other Cisco IP telephony devices currently mark voice call-signaling traffic to AF 31. Until migration to CS 3 is complete, both AF 31 and CS 3 should be reserved for the voice call-signaling traffic class.

Because voice call-signaling traffic is currently marked to AF 31, the mission-critical data should be marked as DSCP 25 instead of AF 31, as shown in the previous figure, until migration to CS 3 is complete.

Example: LLQ Example on the Enterprise WAN Edge Router

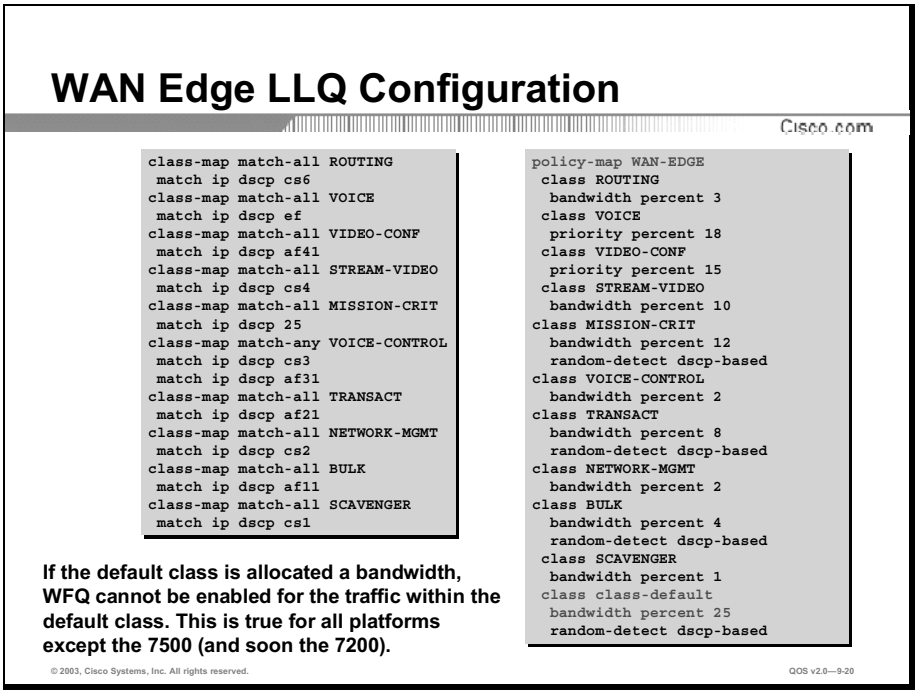

This figure illustrates an example of an enterprise WAN-edge router configuration using LLQ with class-based WRED on certain data traffic classes. This assumes the markings for the different traffic classes are already done at the access or at the distribution layer within the campus network.

Currently (except for the 7500 router platform) all traffic classes except for the default traffic class support only FIFO queuing within the class. Check Cisco.com for the latest information on weighted fair queuing (WFQ) support within each traffic class.

On all platforms, the default traffic class can support either FIFO or WFQ within the class. But if the default traffic class is allocated a minimum bandwidth as shown in this graphic, WFQ will not be supported in the default traffic class. The only current exception is for the 7500 series platforms. In this case, the default traffic class will support only FIFO queuing.

Note: In this example, the max reservable bandwidth on the link is set to 0 percent so up to 100 percent of the link bandwidth can be guaranteed between the different classes.

Enterprise to Service Provider QoS Class Mapping

This topic describes how to map different enterprise traffic classes into traffic classes that are appropriate for service provider use.

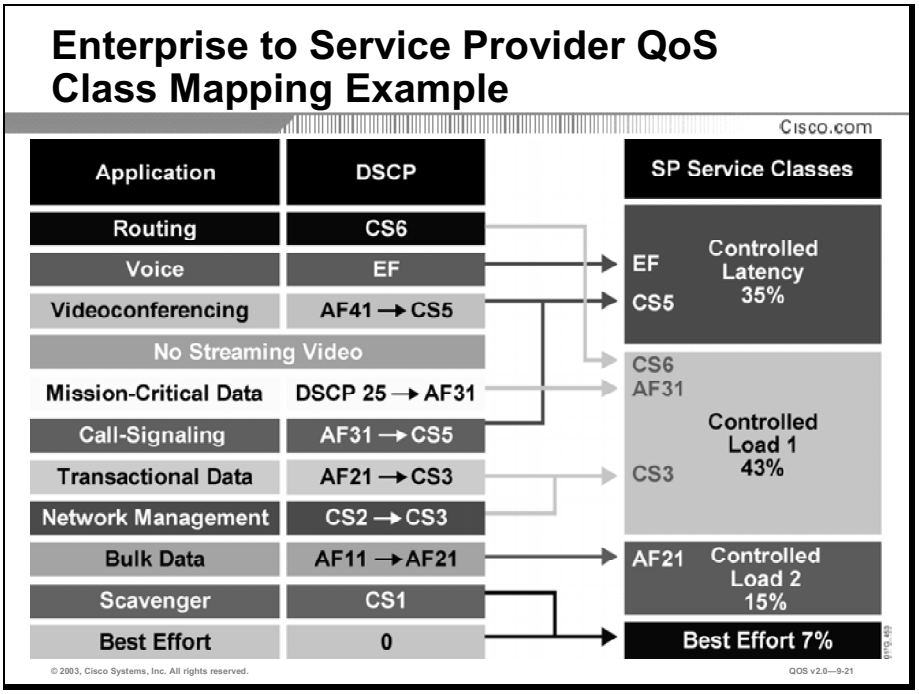

This figure illustrates how the different enterprise traffic classes can be mapped into the four traffic classes offered by a service provider. In the example, there is no streaming video traffic. The four traffic classes offered by the service provider are:

- **Controlled Latency:** A traffic class for all real-time traffic. The controlled latency class admits packets marked with CS 5 and Expedited Forwarding (EF).
- **Controlled Load 1:** A traffic class for all highly bursty traffic. The controlled load 1 class admits packets marked with CS 3, AF 31, and CS 6.
- **Controlled Load 2:** A traffic class for the less bursty traffic. The controlled load 2 class admits packets marked with AF 21.
- **Best Effort:** A traffic class for the best-effort traffic. All other traffic is assigned to the best-effort class.

The figure also shows an example of bandwidth allocation between the four traffic classes where 35 percent of the link bandwidth is allocated for the controlled latency traffic class, 43 percent for controlled load 1 traffic class, 15 percent for controlled load 2 traffic class, and 7 percent for best-effort traffic class.

Example: Remapping the Enterprise Managed CE Traffic Classes into the Traffic Classes Offered by the Service Provider

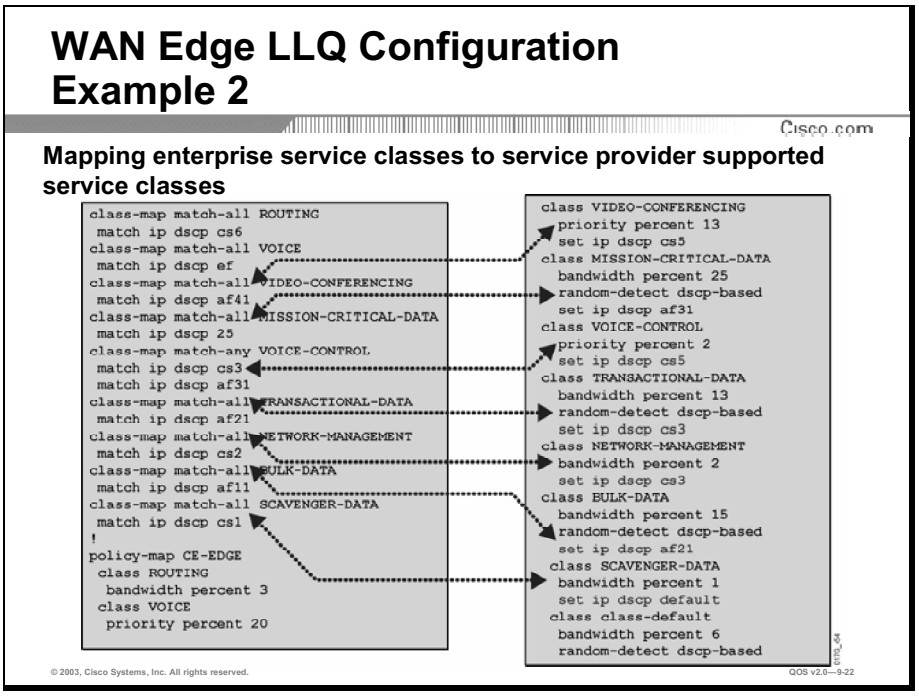

This figure illustrates a sample content engine (CE) provider edge (PE) router configuration using LLQ with class-based WRED on certain data traffic classes. Class-based markings are also used to re-mark the enterprise markings into the markings that are expected by the service provider where:

- Videoconferencing traffic is re-marked from AF 41 to CS 5 and has a maximum bandwidth guarantee of 13 percent of the link bandwidth.
- Mission-critical traffic is re-marked from DSCP 25 to AF 31 and has a minimum bandwidth guarantee of 25 percent of the link bandwidth.
- Voice control (signaling) traffic is re-marked from AF 31 and CS 3 to CS 5 and has a maximum bandwidth guarantee of 2 percent of the link bandwidth.
- Transactional data traffic is re-marked from AF 21 to CS 3 and has a minimum bandwidth guarantee of 13 percent of the link bandwidth.
- Network management traffic is re-marked from CS 2 to CS 3 and has a minimum bandwidth guarantee of 2 percent of the link bandwidth.
- Bulk data traffic is re-marked from AF 11 to AF 21 and has a minimum bandwidth guarantee of 15 percent of the link bandwidth.
- Scavenger data traffic is re-marked from CS 1 to DSCP 0 and has a minimum bandwidth guarantee of 1 percent of the link bandwidth.
- Routing traffic is not re-marked and has a minimum bandwidth guarantee of 3 percent of the link bandwidth.

- Voice bearer traffic is not re-marked and has a maximum bandwidth guarantee of 20 percent of the link bandwidth.
- The default class is not re-marked and has a minimum bandwidth guarantee of 6 percent of the link bandwidth.

Note: In this example, the max reservable bandwidth on the link is set to 0 percent so up to 100 percent of the link bandwidth can be guaranteed between the different classes.

Summary

This topic summarizes the key points discussed in this lesson.

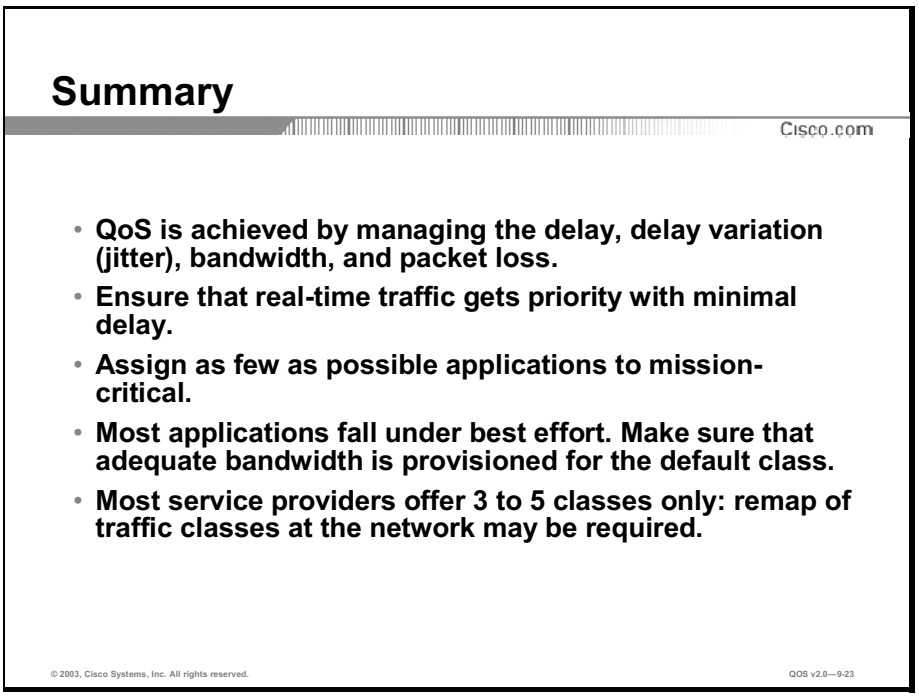

References

For additional information, refer to this resource:

- For more information on QoS Classification best practices, refer to “Classification & Marking” at the following URL: http://www.cisco.com/pcgi-bin/Support/browse/psp_view.pl?p=Internetworking:Classification_and_Marking&viewall=true

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) According to the ITU standard, to ensure high voice quality, what is the recommended maximum one-way delay budget for VoIP traffic?

A) 10 ms

B) 50 ms

C) 150 ms

D) 300 ms

E) 500 ms

Q2) On a typical converged network supporting voice, video, and data, what type of traffic should be serviced in the LLQ?

A) mission-critical data

B) voice and videoconferencing

C) videoconferencing and streaming video

D) voice and mission-critical data

E) voice, videoconferencing, streaming video, and mission-critical data

Q3) Cisco IP Phones and other Cisco IP telephony devices currently mark call-signaling traffic to what DSCP value?

A) EF

B) CS1

C) CS 6

D) AF 11

E) AF 31

F) AF 41

Q4) What are two proper Layer 3 and Layer 2 markings for voice bearer traffic? (Choose two.)

A) AF 31

B) AF 41

C) EF

D) CoS 4

E) CoS 5

F) CoS 6

Q5) The adaptive jitter buffer within Cisco IP telephony devices can usually only compensate for how much jitter?

A) 1 to 10 ms of jitter

B) 20 to 50 ms of jitter

C) 150 to 200 ms of jitter

D) 250 to 300 ms of jitter

Q6) What three factors determine the amount of bandwidth needed per VoIP call? (Choose three.)

A) codec type

B) voice sampling rate

C) Layer 2 Media

D) jitter buffer size

E) router CPU speed

Q7) Mission-critical data should be classified with what CoS value?

A) 1

B) 2

C) 3

D) 4

E) 5

F) 6

G) 7

Q8) Review the following QoS configuration:

```
class-map match-all ROUTING
 match ip dscp cs6
class-map match-all VOICE
 match ip dscp ef
class-map match-all VIDEO-CONFERENCING
 match ip dscp af41
class-map match-all MISSION-CRITICAL-DATA
 match ip dscp 25
class-map match-any VOICE-CONTROL
 match ip dscp cs3
 match ip dscp af31
class-map match-all TRANSACTIONAL-DATA
 match ip dscp af21
class-map match-all NETWORK-MANAGEMENT
 match ip dscp cs2
class-map match-all BULK-DATA
 match ip dscp af11
class-map match-all SCAVENGER-DATA
 match ip dscp cs1
!
policy-map CE-EDGE
 class ROUTING
  bandwidth percent 3
 class VOICE
  priority percent 20
class VIDEO-CONFERENCING
  priority percent 13
  set ip dscp cs5
class MISSION-CRITICAL-DATA
  bandwidth percent 25
  random-detect dscp-based
  set ip dscp af31
class VOICE-CONTROL
  priority percent 2
  set ip dscp cs5
class TRANSACTIONAL-DATA
  bandwidth percent 13
  random-detect dscp-based
  set ip dscp cs3
```



```
class NETWORK-MANAGEMENT
  bandwidth percent 2
  set ip dscp cs3
class BULK-DATA
  bandwidth percent 15
  random-detect dscp-based
  set ip dscp af21
 class SCAVENGER-DATA
  bandwidth percent 1
  set ip dscp default
 class class-default
  bandwidth percent 6
  random-detect dscp-based
```

Based on the above configuration, which two of the following statements are correct? (Choose two.)

A) Only voice bearer and voice control (signaling) traffic will be serviced by the LLQ.

B) Voice control traffic will be remapped to CS 5.

C) Scavenger traffic is guaranteed less bandwidth than the best-effort class.

D) Routing protocol traffic will be serviced by the LLQ.

E) WRED is enabled for the voice traffic class.

F) Scavenger traffic will be remapped to CS 1.

Q9) In a typical converged network, which two of the following is considered a QoS best practice? (Choose two.)

A) Traffic classification and marking should be performed at the high speed core layer and close to the destination.

B) A tool like NBAR should be used to discover all the applications running on the network.

C) Ensure real-time traffic gets the highest priority and minimum delay.

D) All traffic that is not real-time should be classified into the scavenger traffic class.

E) Ensure mission-critical applications have higher priority over all other traffic.

Q10) Which of the following applications are typically less delay and jitter-sensitive than VoIP? (Choose all that apply.)

A) videoconferencing

B) ERP

C) streaming video

D) call-signaling

E) FTP

Quiz Answer Key

Q1) C

Relates to: QoS Requirements Summary

Q2) B

Relates to: Traffic Classification

Q3) E

Relates to: Traffic Classification

Q4) C, E

Relates to: Traffic Classification

Q5) B

Relates to: QoS Requirements Summary

Q6) A, B, C

Relates to: QoS Requirements Summary

Q7) C

Relates to: Traffic Classification

Q8) B, C

Relates to: Enterprise to Service Provider QoS Class Mapping

Q9) B, C

Relates to: QoS Best Practices

Q10) B, C, D, E

Relates to: Voice/Video/Data QoS Requirements

Case Study: Deploying End-to-End QoS

Overview

When using public transport, a traveler may benefit from contractual commitments from the public transport provider; for example, a guarantee from an airline that 95 percent of their flights will arrive within 5 minutes of the scheduled time. The commitments may include other parameters or metrics such as number of stops en route. The more competitive the market for the particular service, the more comprehensive and the tighter the commitments or service level agreements (SLA) that are offered.

In the same way, the increase of competition between IP service providers, together with the heightened importance of IP to business operations, has led to an increased demand and consequent supply of IP services with tighter SLAs for IP performance.

The DiffServ architecture enables IP networks to be engineered to support tight SLA commitments.

Relevance

For a service provider IP service, the SLA commitments are based on delay, jitter, packet-loss rate, throughput, and availability. This lesson gives examples of these IP QoS SLA parameters, and describes how the different Cisco IOS QoS mechanisms should be used so that SLA parameters can be met.

In service provider IP backbone networks, it is noted that the mechanisms employed at the edge of the network to deliver tight SLAs are more involved than those used in the core. In the core, where traffic is aggregated, SLA requirements for a traffic class can be translated into the appropriate bandwidth requirements, and the problem of SLA assurance can effectively be reduced to that of bandwidth provisioning.

Objectives

Upon completing this lesson, you will be able to correctly identify and describe the set of QoS mechanisms that are used to implement Cisco end-to-end QoS “best practices” in a typical enterprise network connected through a service provider that is providing Layer 3 IP services. This includes being able to meet these objectives:

- Explain IP QoS SLA and provide some SLA examples
- Explain the typical network requirements within each functional block (campus LAN, WAN edge, service provider backbone, branch), which makes up the end-to-end network
- Explain the best practice QoS implementations and configurations within the campus LAN
- Explain the best practice QoS implementations and configurations on the WAN CE and PE routers
- Explain the best practice QoS implementations and configurations on the service provider backbone (P) and PE routers

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- A good understanding of using the MQC configuration method to implement Cisco IOS QoS mechanisms
- A good understanding of Cisco IOS QoS classification/marking, scheduling, and provisioning mechanisms

Outline

The outline lists the topics included in this lesson.

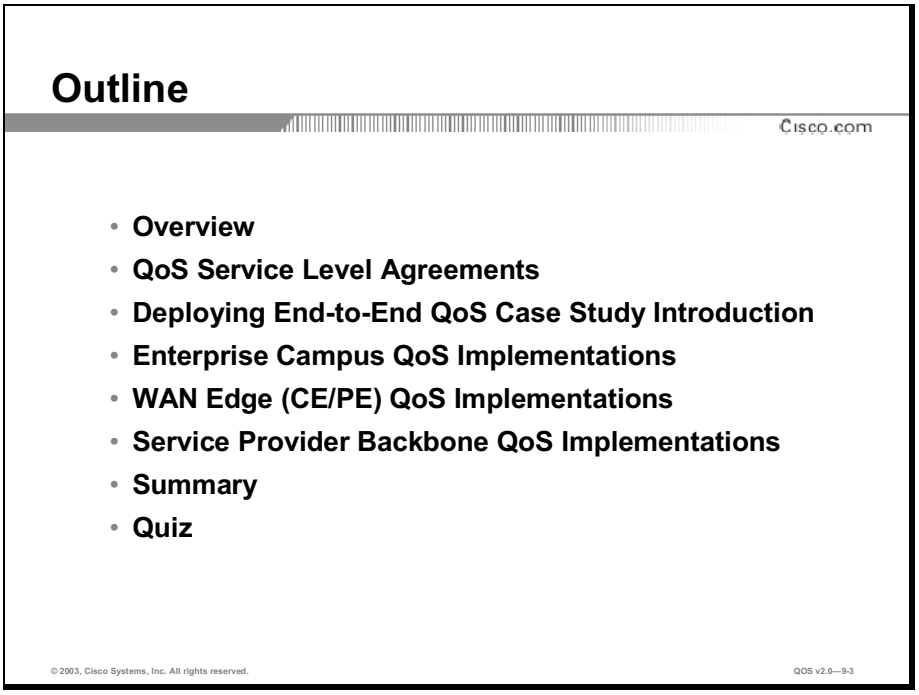

QoS Service Level Agreements

This topic explains IP QoS SLA and provides some SLA examples.

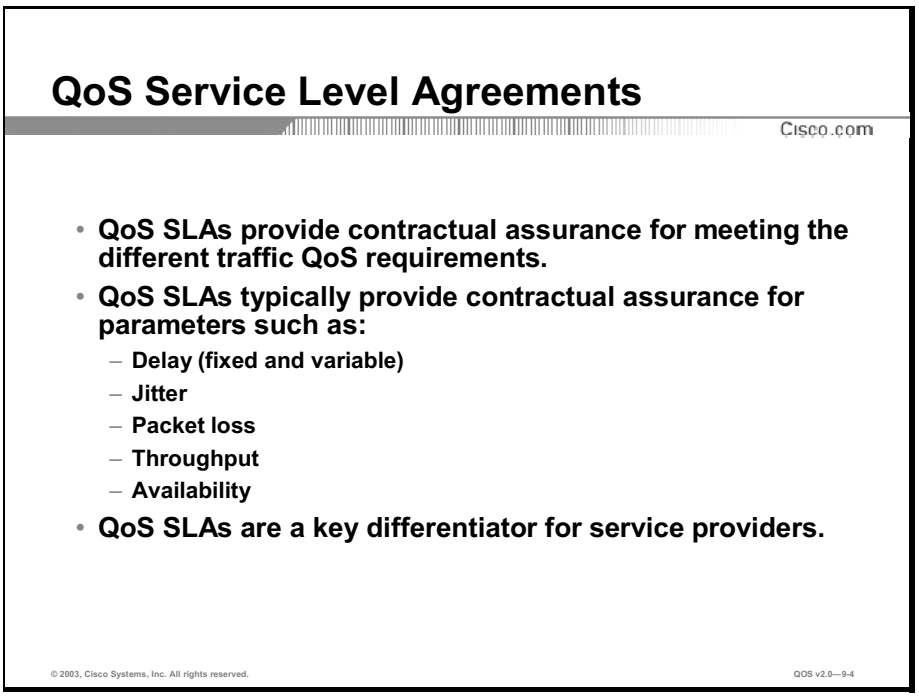

An SLA stipulates the delivery and pricing of numerous service levels and spells out penalties for shortfalls. SLAs can cover an assortment of data services like frame relay, leased lines, Internet access, Web hosting, etc. The best way to understand SLAs is to break them into two pieces: negotiating the technology agreement and verifying that promises made are promises kept.

To support integrated voice, video, and data services, service providers are under increasing pressure to offer differentiated service levels to their customers, often in conjunction with SLAs that provide contractual assurance for meeting the different traffic QoS requirements. A QoS SLA typically provides contractual assurance for parameters such as delay, jitter, packet loss, throughput, and availability.

With the rapid growth of new multimedia real-time applications like IP telephony, web-conferencing and e-learning, offering IP QoS SLA is becoming a key service differentiator for service providers.

Enterprise Network with Traditional Layer 2 Service

- **Provider sells the customer a Layer 2 service**
- **PPP SLA from the provider**
- **Enterprise WAN likely to get congested**
- **IP QoS required for voice, video, data integration**
- **Service provider is not involved in IP QoS**

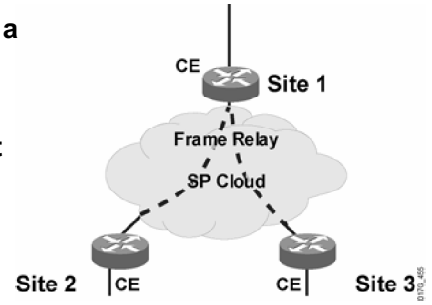

This figure illustrates a service provider providing only Layer 2 services to the enterprise customer. The customer edge (CE) routers at the various customer sites are inter-connected by Frame Relay virtual circuits (VCs). These VCs can be fully meshed, partially meshed, or set up as hub-and-spokes, depending on the customer requirements.

In this environment, the service provider is only responsible for the end-to-end Layer 2 VC connections. The service provider only provides a PPP SLA guarantee for each VC connection. The service provider is not involved with providing IP QoS to the customer.

To provide IP QoS for voice, video, and data integration over the Frame Relay VCs, the customer must configure the proper QoS mechanisms like traffic shaping, LLQ, FRF.12, and cRTP at the WAN CE routers because the Frame Relay WAN link is likely to get congested.

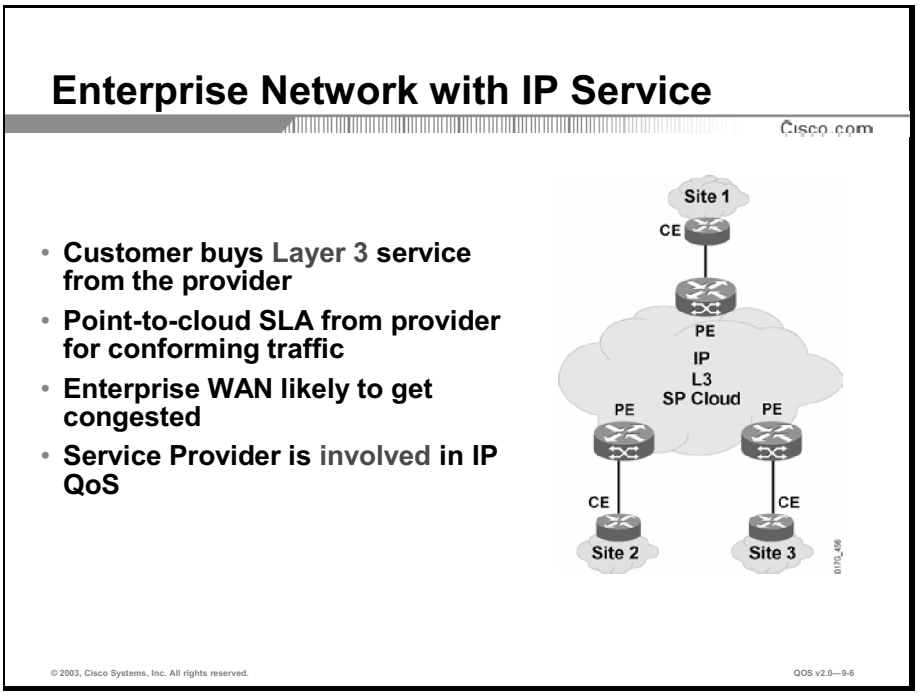

This figure illustrates a service provider that is providing Layer 3 services to the enterprise customer. The CE routers at the various customer sites connect to the provider edge (PE) of the service provider router. From a particular customer site perspective, every IP address that is not located on-site is reachable via the service provider IP backbone network.

In this environment, because the service provider is responsible for providing IP services, the service provider can provide value-added IP services to the customer by providing point-to-cloud SLAs for the conforming traffic from the customer. For example, divide customer traffic at the network edge into controlled latency, controlled load 1, and controlled load 2 classes, and then provide IP QoS assurances to each traffic class that is conforming to the contractual rate over a DiffServ IP backbone. For all nonconforming (exceeding) traffic, the service provider can re-mark and delivery all nonconforming traffic with best-effort service.

Know the SLA Offered by Your SP

Cisco.com

- **SLA typically includes between 3 and 5 classes.**
- **Real-time traffic gets fixed bandwidth allocation.**
- **Data traffic gets variable bandwidth allocation with minimum guarantee.**
- **Additional classes not visible to customer may exist at the edge (for example, management/control traffic).**

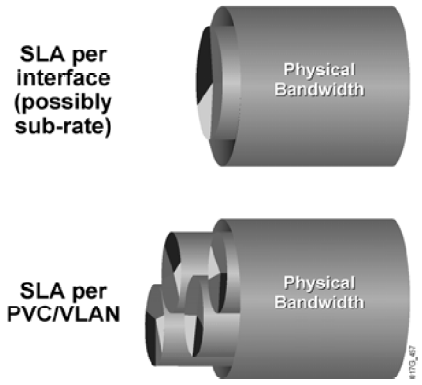

A typical IP QoS SLA offered by most service providers often includes 3 to 5 traffic classes; for example, a real-time traffic class, a mission-critical data traffic class, one or two other data traffic classes, and a best-effort traffic class. The IP QoS SLA for the real-time traffic class should be guaranteed a fixed maximum bandwidth while the data traffic classes should be guaranteed a minimum bandwidth. Typically, the bandwidth allocation is configured as a percentage of the interface bandwidth. Each traffic class can also have a latency, delay, jitter, and packet-loss guarantee.

At the CE PE edge, there may be additional traffic classes that are used by the service providers only. For example, a management traffic class for traffic like Telnet or SNMP, from the service provider to the service provider-managed CE routers.

If a single physical interface is only serving one customer, the SLA is typically set up per interface. To provide easy bandwidth upgrades, service providers often install a high-speed link to the customer and then offer a sub-rate access.

If a single physical interface is serving many different customers, the SLA is typically set up per PVC or per virtual LAN (VLAN). To provide easy bandwidth upgrades, the service provider often installs a high-speed link to the customer and then offers a sub-rate access.

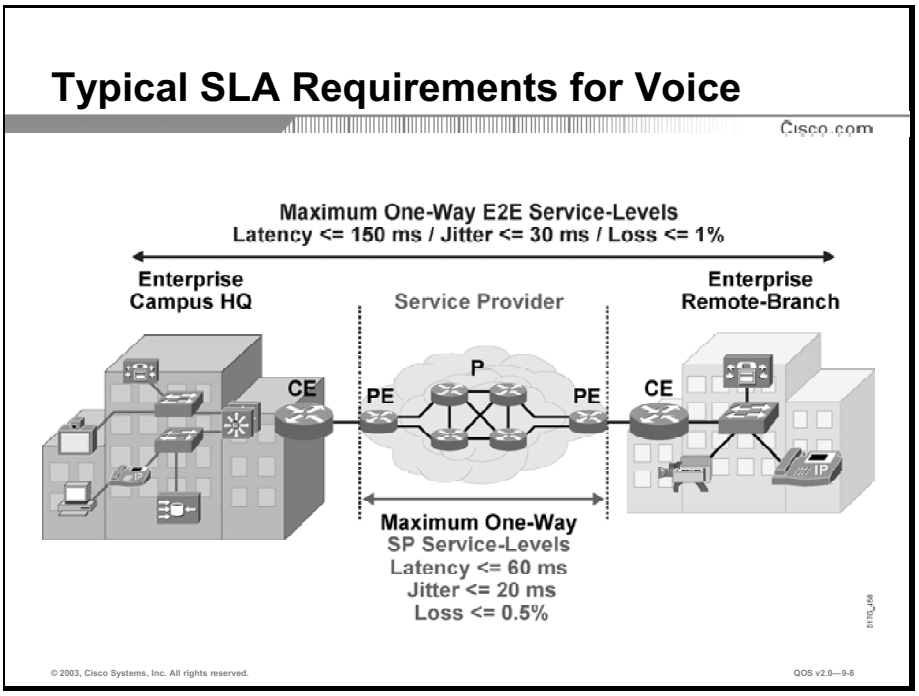

To meet the QoS requirements for the different traffic types, both the enterprise and the service provider must implement the proper QoS mechanisms to provide end-to-end QoS for the packets traversing across a service provider IP network. In the figure, the enterprise headquarters and the enterprise branch office are connected to a service provider that is providing Layer 3 services.

In this example, the service provider is providing an SLA for voice traffic with a latency of 60 ms or less, a jitter of 20 ms or less, and a packet loss of 0.5 percent or less. To meet the end-to-end QoS requirements for voice packets within headquarters, within the branch office, and on the two WAN edge links, the accumulative total delay should be 90 ms (150 to 60) or less, the jitter should be 10 ms (30 to 20) or less, and the packet loss should be 0.5 percent (1 to 0.5) or less.

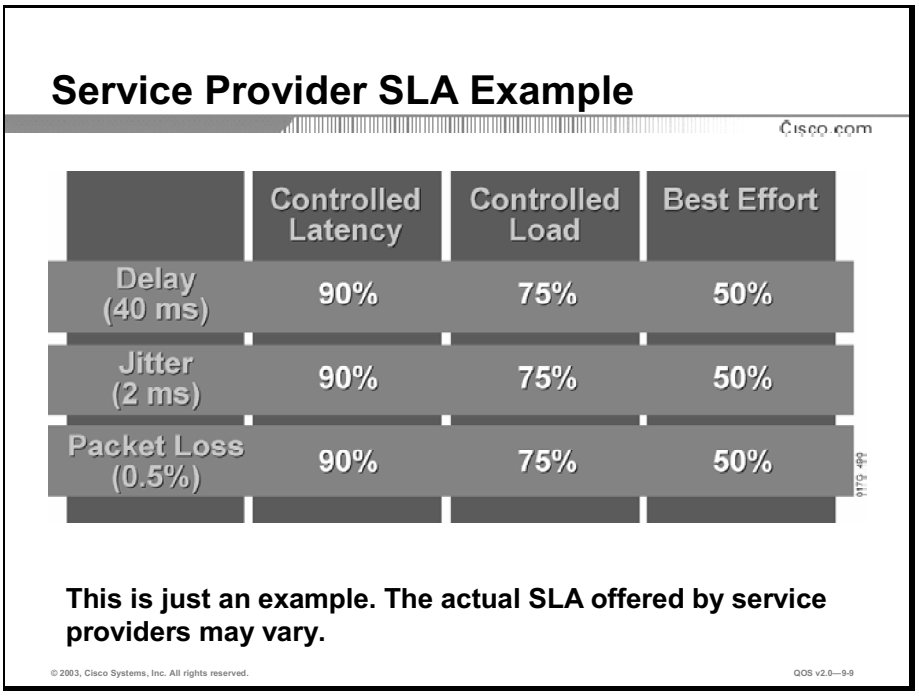

This figure shows an example of a typical IP QoS SLA from an IP service provider. In this example, the service provider is offering three service classes to the customer: controlled latency, controlled load, best effort. This example SLA guarantees include:

- For the controlled latency class:
  - A one-way delay of 40 ms that is guaranteed 90 percent of the time.
  - A jitter of 2 ms that is guaranteed 90 percent of the time.
  - A packet loss of 0.5 percent that is guaranteed 90 percent of the time.
- For the controlled load class:
  - A one-way delay of 40 ms that is guaranteed 75 percent of the time.
  - A jitter of 2 ms that is guaranteed 75 percent of the time.
  - A packet loss of 0.5 percent that is guaranteed 75 percent of the time.
- For the best-effort class:
  - A one-way delay of 40 ms that is guaranteed 50 percent of the time.
  - A jitter of 2 ms that is guaranteed 50 percent of the time.
  - A packet loss of 0.5 percent that is guaranteed 50 percent of the time.

This is just an example. Actual SLA offered by service providers may vary.

In the United States tier 1 Internet backbone, a typical round-trip time (RTT) delay between two service provider POPs (points of presence) is about 40 ms. It is typical for a service provider to offer a monthly average loss on its network of less than 1 percent.

Deploying End-to-End QoS Case Study Introduction

This topic explains the typical network requirements within each functional block (headquarter campus LAN, WAN edge, service provider backbone, branch office), which makes up the end-to-end network used in the case study.

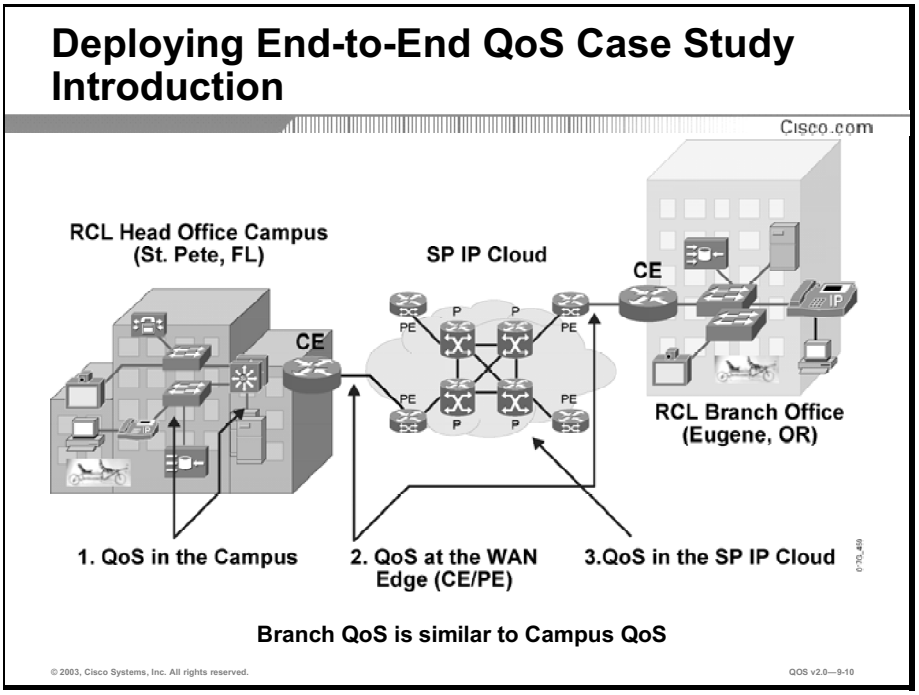

To meet the QoS requirements for different traffic types, both the enterprise and the service provider must implement the proper IP QoS mechanisms to provide end-to-end QoS for the packets traversing across a service provider network.

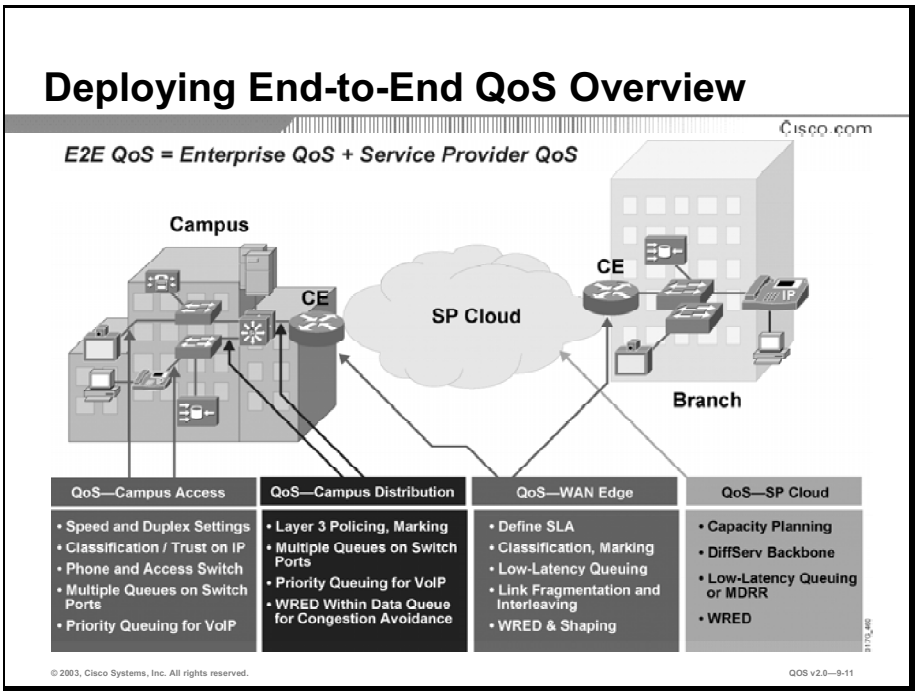

To provide end-to-end QoS, both the enterprise and service provider must implement the proper QoS mechanisms to ensure the proper PHBs for each traffic class across the whole network. Until recently, IP QoS was not an issue in an enterprise campus network where bandwidth is plentiful. But as more applications like IP telephony, videoconferencing, e-learning, and mission-critical data applications are being implemented in the campus, it has become evident that buffer management, not just bandwidth, is an issue that must be addressed. IP QoS functions such as classification, scheduling, and provisioning are also now required within the campus to manage bandwidth and buffers to minimize loss, delay, and jitter.

This figure lists some of the requirements within the different building blocks that make up the end-to-end network. Most of the more complex QoS configurations occur at the WAN edge. In the IP core (service provider cloud), only queuing like LLQ or MDRR and WRED should be required.

Enterprise Campus QoS Implementations

This first part of the case study explains some of the best practice QoS implementations and configurations within the campus LAN.

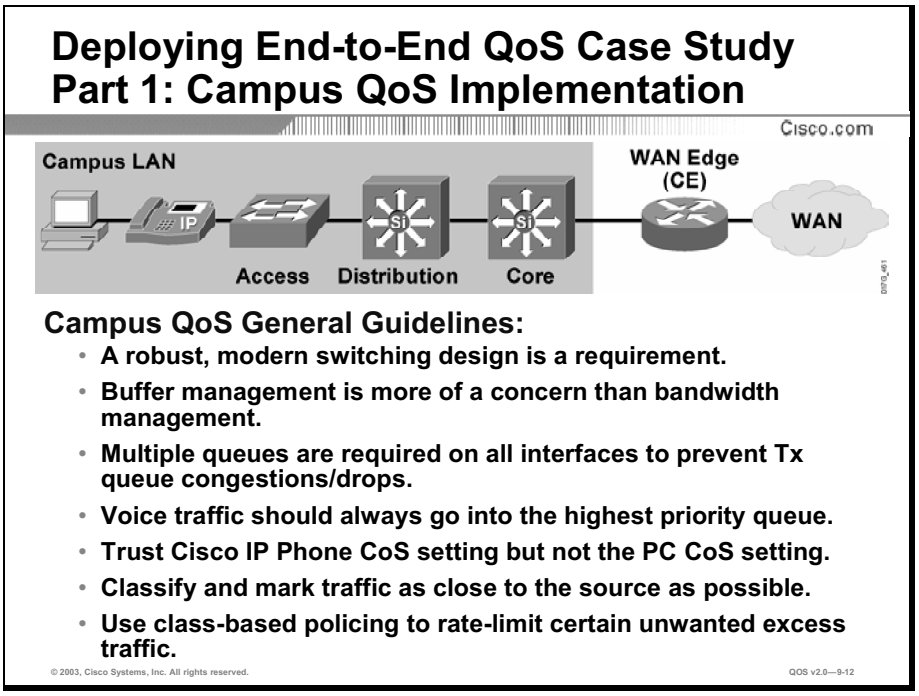

Applications like IP telephony, videoconferencing, e-learning, and mission-critical data applications are becoming more commonly implemented in enterprise networks. IP QoS functions such as classification, scheduling, and provisioning are required within the campus to manage the LAN bandwidth and the LAN switches output buffers to minimize loss, delay, and jitter within these campus networks. There are many points in a campus network where QoS mechanisms are required. Some of the general guidelines when implementing campus QoS include:

- The campus network design should use a hierarchical design model, which has 3 layers: access layer, distribution layer, and core layer.
- Use multiple queues on the transmit interfaces to minimize the potential for dropped or delayed traffic caused by transmit buffer congestion.
- Real-time voice and video traffic should have priority over the data traffic.
- Establish proper trust boundaries. For example, at the access layer switches, only trust the IP Phone CoS marking, not the PC CoS marking.
- Classify and mark the traffic as soon as possible. For example, use class-based markings at the access layer switches to classify and mark the traffic coming into the network.
- NBAR can be used on the access or distribution switches to classify traffic if the switch supports NBAR.
- At the access or distribution layers, use class-based policing to limit any undesired non-business-related excess traffic: for example, rate-limiting peer-to-peer file-sharing traffic to a minimal rate.

- No classification or marking should be done at the core layer, as this slows down traffic. The core layer provides high-speed transport between the distribution layers and generally requires no QoS mechanisms.
- Most of the more complex QoS configurations like traffic shaping, header compression, LFI, LLQ, WRED, and re-marking of traffic occur at the WAN CE router.

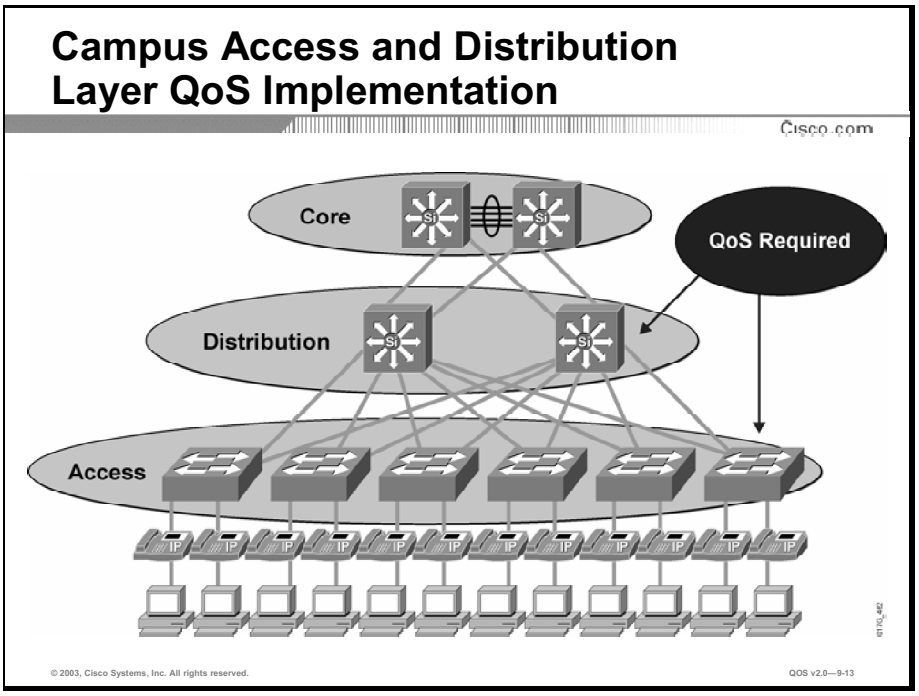

The graphics will go through the typical QoS configurations that are required at the access and distribution layer switches. In this example, core layer configurations will not be illustrated. Typically, the high-speed core layer does not require any QoS mechanisms because all the links are very high speed.

Catalyst 3550 Access Switch

- 1P3Q2T (recommended) or 4Q2T
- WRR default
  - Each queue is given a relative weight with one priority queue
- Default values are as follows:

CoS Value	CoS Priority Queues
0, 1	1
2, 3	2
4, 5	3
6, 7	4 (priority queue)

- Move CoS 5 traffic to queue 4
- Trust Cisco IP Phone on access links
- Trust DSCP or CoS on uplinks
- CoS-to-DSCP mappings
- Classify and mark mission-critical traffic using ACLs

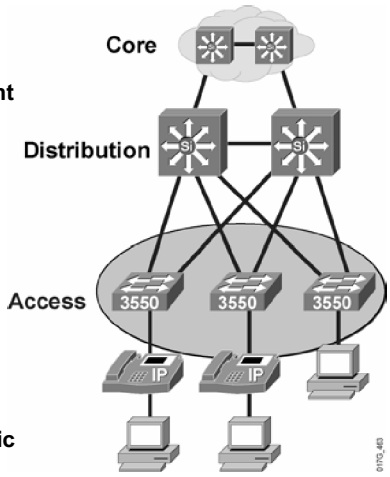

In this case study, the access layer switches are the Catalyst 3550 switches.

The recommended configuration for the transmit interface for the Catalyst 3550 is one priority queue and three queues, each with a single drop threshold.

Scheduling is performed using weighted round robin (WRR), where each queue is given a relative weight while the priority queue is serviced exhaustively. The default CoS to queue mapping is as follows:

- Frames marked with a CoS 6 and 7 will go into queue 4 (priority queue for the Catalyst 3550).
- Frames marked with a CoS 4 and 5 will go into queue 3.
- Frames marked with a CoS 2 and 3 will go into queue 2.
- Frames marked with a CoS 0 and 1 will go into queue 1.

When servicing voice traffic, CoS 5 frames should be configured to go to the priority (expedite) queue. By default, Cisco IP Phones forward the voice traffic with an 802.1Q priority of 5 (CoS 5).

At the access layer switch, configure the trust boundary to trust only the CoS marking on the Cisco IP Phone, not the CoS marking on the PC. Also configure the trust boundary to trust the DSCP or CoS marking from the distribution layer switches, depending on the capabilities of the distribution layer switch that is attached.

For Layer 2 to Layer 3 QoS mappings, the Catalyst 3550 has default CoS-to-DSCP mappings. However, these default mappings can be manually configured to override the default mappings. The default CoS-to-DSCP map is as follows:

CoS-to-DSCP Default Mapping

Marker	Value							
CoS Values	0	1	2	3	4	5	6	7
DSCP Values	0	8	16	24	32	40	48	56

As a general rule, it is best to classify and mark the traffic as close to the source as possible. Therefore, at the access layer switches, class-based markings can be configured to classify and mark the traffic.

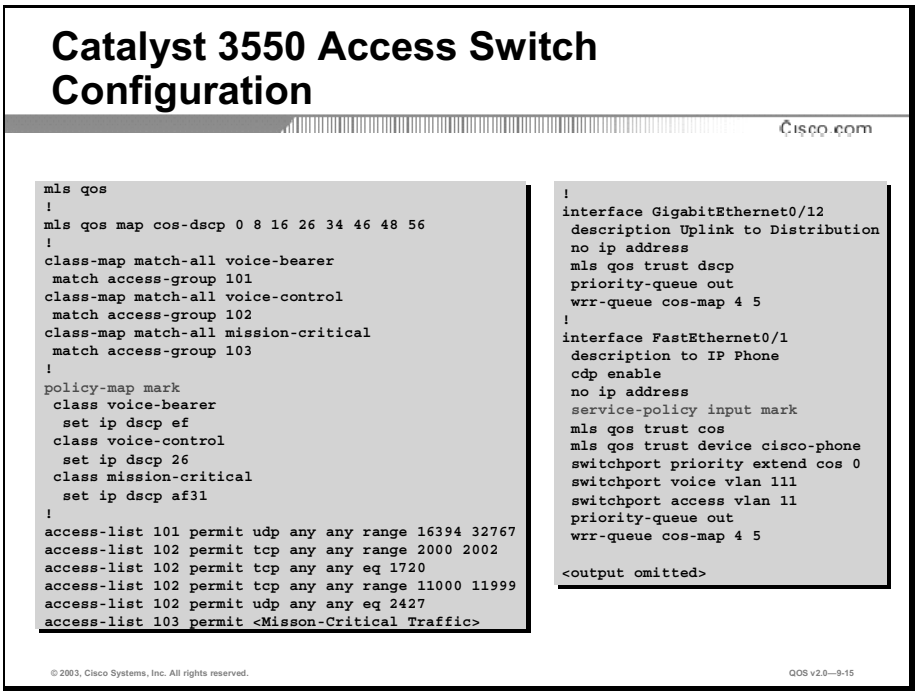

This figure shows a sample QoS configuration on a Catalyst 3550 access switch.

When classification of traffic by the access-layer switch is required, the Catalyst 3550 provides a powerful set of features for classifying traffic as it enters the network. This configuration example uses three class maps to classify traffic into three classes (voice-bearer, voice-control and mission-critical) using three extended IP access control lists (ACLs). The access ports are configured to trust only the CoS markings from the Cisco IP Phone; all other traffic will be classified and marked using the traffic policy called “mark”. This traffic policy is used to mark the voice-bearer class with a DSCP of EF, the voice-control (call-signaling) class with a DSCP of 26, and the mission-critical class with a DSCP of AF 31. This traffic policy is then applied to the switch access ports.

Note: If using the Catalyst 2950, the ACLs are limited to match a single TCP or UDP port. For example, the Catalyst 2950 ACLs cannot be used to match against a UDP or TCP port range for matching VoIP RTP bearer and control traffic.

The uplink switch ports are configured to trust the DSCP markings from the distribution layer switches. Layer 3 capable distribution layer switches are being used in this case study.

The Catalyst 3550 supports multiple VLANs on the access port to which the Cisco IP Phone is attached. The **switchport voice vlan** interface configuration command is used to configure voice VLAN on the port. In this example, all the voice traffic from the access port will belong to VLAN 111 while the data traffic will belong to VLAN 11.

To ensure EF for the voice-bearer traffic, all CoS 5 traffic is configured to go into queue 4 (the priority queue [PQ]) along with the CoS 6 and CoS 7 traffic.

The default CoS-to-DSCP mapping is also changed so that CoS 3 (voice control) maps to AF 31, CoS 4 (videoconferencing) maps to AF 41, and CoS 5 (voice bearer) maps to DSCP 46 (EF).

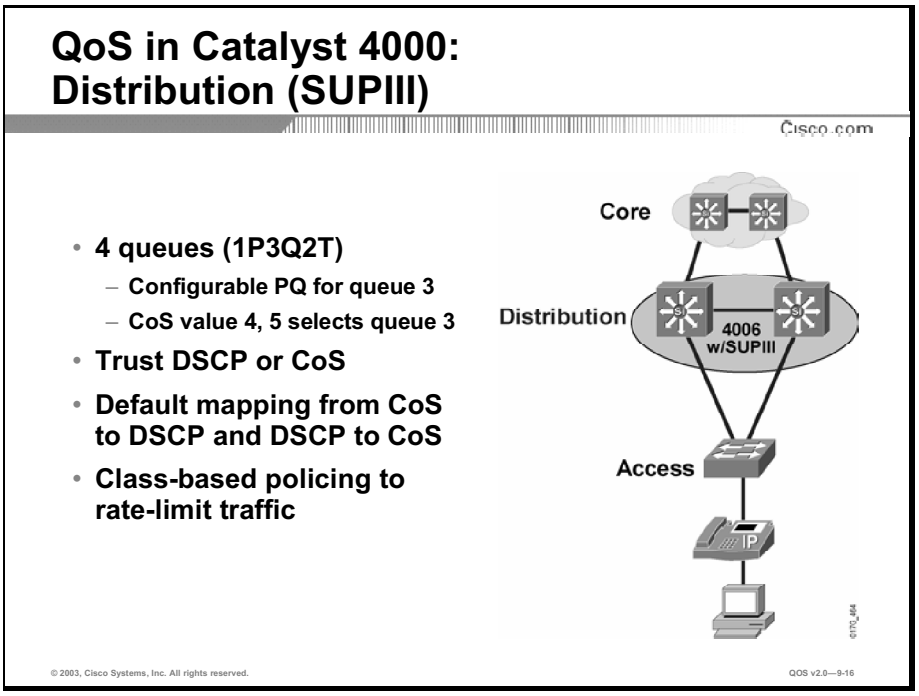

In this case study, the distribution layer switches are the Catalyst 4000 switches.

The Catalyst 4000 has a total of four queues, each with two different drop thresholds. The link bandwidth can be configured to be shared differently among the four transmit queues by assigning the minimum guaranteed bandwidth for each transmit queue. By default, all four queues are scheduled in a round robin manner with each queue having 25 percent of the link bandwidth. The transmit queue 3 on each port can be configured with higher priority. When transmit queue 3 is configured with higher priority, packets in transmit queue 3 are scheduled ahead of packets in other queues. The default DCSP-to-transmit queue map (which can be manually configured to override the default) is as follows:

- DSCP 0-15 to queue 1
- DSCP 16-31 to queue 2
- DSCP 32-47 to queue 3
- DSCP 48-63 to queue 4

At the distribution switch, configure the trust boundary to trust the DSCP or CoS marking from the core and access layer switches. For Layer 2 to Layer 3 and Layer 3 to Layer 2 QoS mappings, the Catalyst 4000 has default CoS-to-DSCP and DSCP-to-CoS mappings. However, these default mappings can be manually configured to override the default mappings. The Catalyst 4000 has the following CoS-to-DSCP and DSCP-to-CoS mappings:

CoS-to-DSCP Default Mapping

Marker	Value							
CoS Values	0	1	2	3	4	5	6	7
DSCP Values	0	8	16	24	32	40	48	56

DSCP-to-CoS Default Mapping

Marker	Value							
DSCP Values	0 - 7	8 - 15	16 - 23	24 - 31	32 - 39	40 - 47	48 - 55	56 - 63
CoS Values	0	1	2	3	4	5	6	7

Class-based policing may also be implemented to rate-limit certain traffic classes.

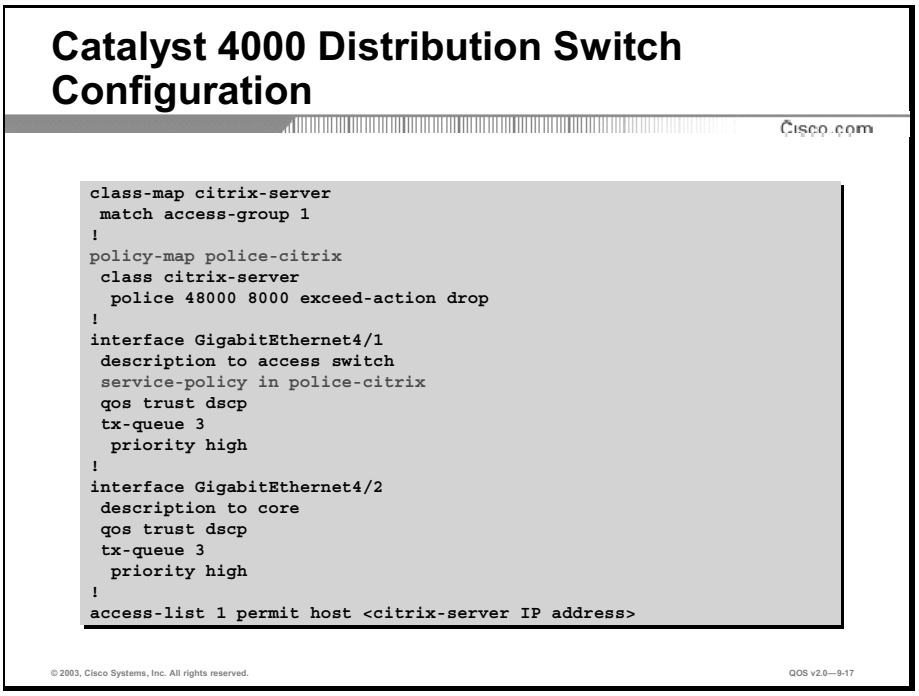

This figure shows a sample QoS configuration on a Catalyst 4000 distribution switch.

This configuration uses a class map to classify the Citrix server traffic into the citrix-server traffic class.

A traffic policy called “police-citrix” is used to police the Citrix traffic to a CIR of 48000 bps with a committed burst (Bc) of 8000 bits. All excess traffic is dropped.

The downlink and uplink ports are configured to trust the DSCP markings from the access and core layer switches.

The transmit queue 3 is configured as the high-priority queue.

WAN Edge (CE/PE) QoS Implementations

This second part of the case study explains some of the best practice QoS implementations and configurations on the WAN CE and PE routers.

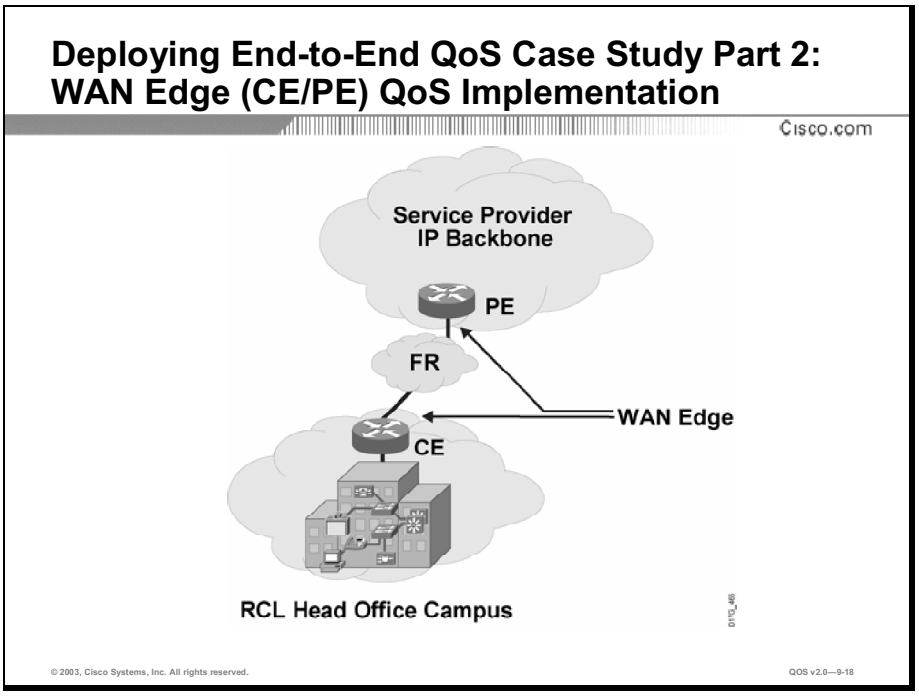

The following slides will go through the typical QoS configurations that are required at the CE and PE WAN routers.

On the CE PE WAN link, LLQ/CBWFQ, traffic shaping, cRTP, and LFI are typically required.

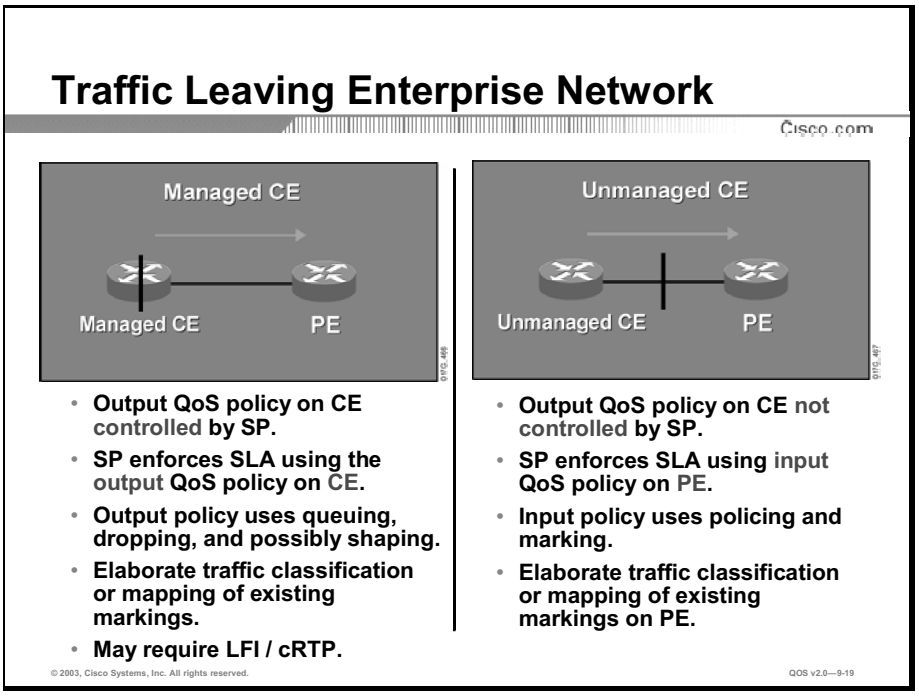

The QoS requirements on the CE and PE router will differ, depending on whether or not the CE is managed by the service provider.

For traffic leaving the enterprise CE router towards the service provider PE router, the figure illustrates the general QoS requirements on the CE and PE routers.

For managed CE service, the WAN edge output QoS policy on the CE will be managed and configured by the service provider.

For unmanaged CE service, the WAN edge output QoS policy on the CE will be managed and configured by the enterprise customer.

For managed CE service, the service provider can enforce the SLA for each traffic class using the output QoS policy on the CE; for example, use LLQ/CBWFQ to give a maximum bandwidth guarantee to the real-time voice and video traffic class and give a minimum bandwidth guarantee to the data traffic classes and use class-based shaping to provide a maximum rate limit to each data traffic class.

For unmanaged CE service, because the service provider has no control of the CE, the service provider can only enforce the SLA for each traffic class at the input of the PE router. For example, use class-based policing to rate-limit the input traffic rate of the different traffic classes and to re-mark the exceeding traffic.

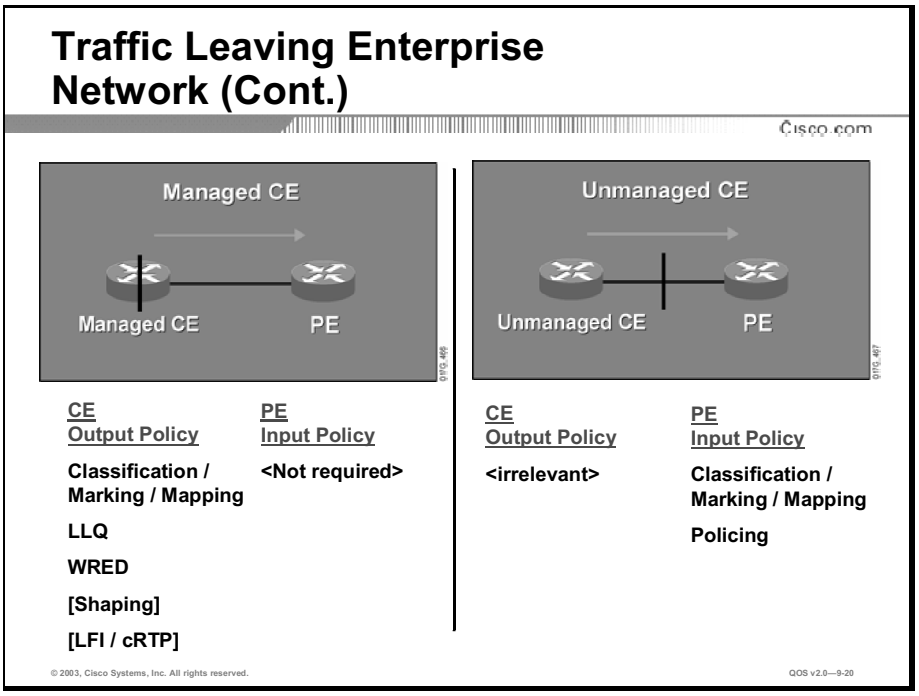

For the traffic leaving the CE router, this figure illustrates the different QoS mechanisms that are commonly implemented at the CE and PE routers, depending on whether or not the CE is managed by the service provider.

For example, for unmanaged CE, the CE output policy is managed and configured by the enterprise customer. Therefore, it is irrelevant to the service provider. At the PE input interface, the service provider will have a policy to classify, mark, or remaps the traffic. The service provider also typically implements traffic policing to rate-limit the input traffic rate from the enterprise customer so the traffic rate does not exceed the contractual rate as specified in the SLA.

For managed CE, the CE output policy is managed and configured by the service provider. The service provider typically has an output policy on the CE router to classify and mark the traffic exiting the CE router. LLQ/CBWFQ and WRED are also used to for congestion management and congestion avoidance. To compensate for speed mismatch or oversubscription, traffic shaping may be required. To improve the link efficiency, LFI and cRTP are used.

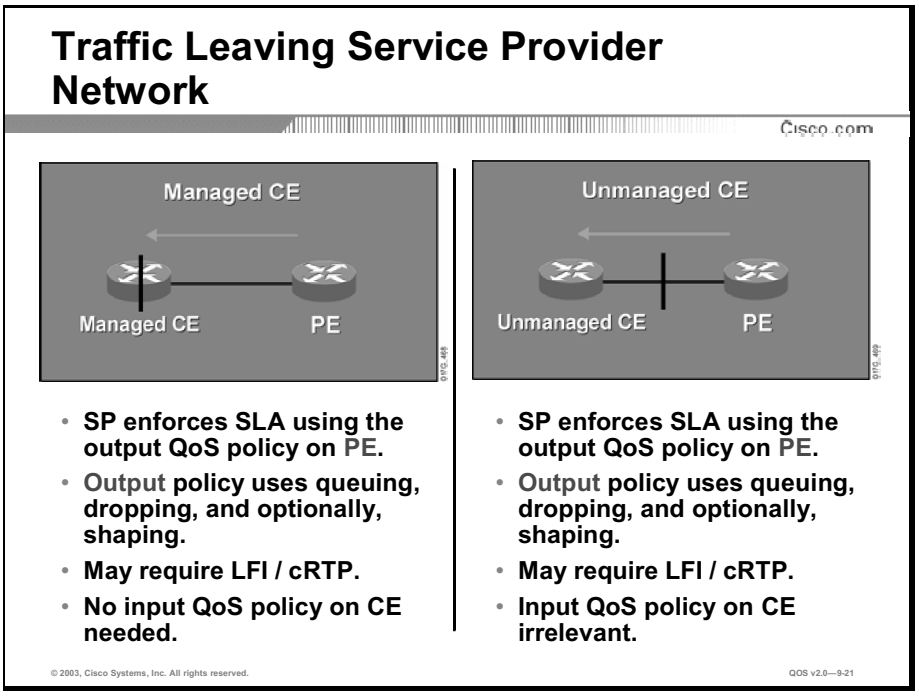

For traffic leaving the service provider PE router towards the enterprise CE router, this figure illustrates the general QoS requirements on the CE and PE routers.

For both managed and unmanaged CE service, the service provider can enforce the SLA for each traffic class using the output QoS policy on the PE. For example, use LLQ/CBWFQ to give a maximum bandwidth guarantee to the real-time voice and video traffic class and give a minimum bandwidth guarantee to the data traffic classes and use class-based shaping to provide a maximum rate limit to each data traffic class.

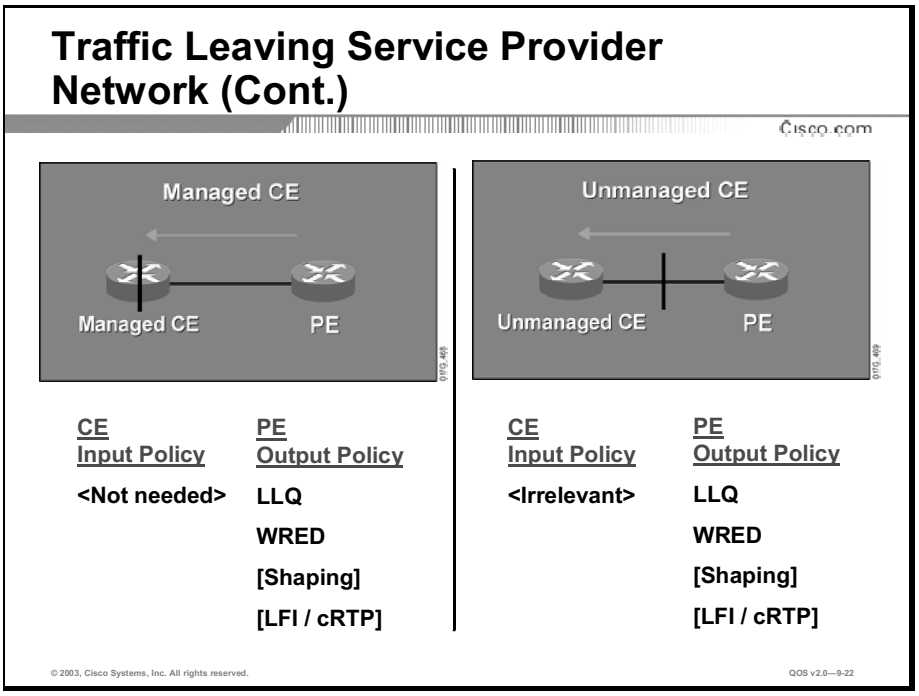

For traffic leaving the service provider PE router towards the enterprise CE router, this figure illustrates the different QoS mechanisms that are commonly implemented at the PE router.

For both managed and unmanaged CE service, the service provider typically has an output policy on the PE router using LLQ/CBWFQ and WRED for congestion management and congestion avoidance. To compensate for speed mismatch or over-subscription, traffic shaping may be required. To improve the link efficiency, LFI and cRTP are used.

A customer edge input policy is not required for managed and unmanaged CE services.

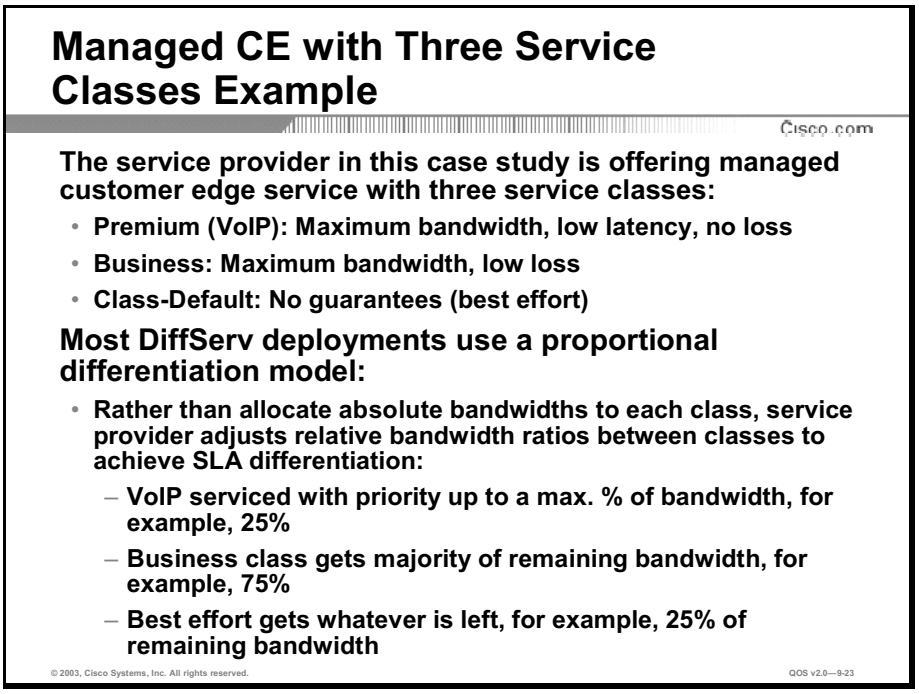

In this case study, the service provider is implementing an IP DiffServ backbone and is offering three traffic classes with different SLAs for each:

- The premium traffic class is intended for voice traffic. This class has a maximum bandwidth limit, a low latency, and no loss guarantee.
- The business traffic class is intended for mission-critical traffic. This class has a minimum bandwidth and a low loss guarantee.
- The default (best effort) traffic class is intended for all other traffic. This class has no guarantee.

When implementing LLQ/CBWFQ, the bandwidth guarantees can be specified in kbps or in percent of the available bandwidth or in percentage of the remaining available bandwidth. Most DiffServ deployments today use a proportional differentiation model where the bandwidth guarantees are configured as a percentage instead of a fixed kbps.

For example, with three traffic classes, the bandwidth allocation can be divided as such:

- The premium (VoIP) class LLQ can have a maximum bandwidth guarantee of 25 percent of the link bandwidth.
- The business (mission-critical) class can have a minimum bandwidth guarantee of 75 percent of the remaining bandwidth after the LLQ is serviced.
- The default (best effort) class can have a minimum bandwidth guarantee of whatever is left, in this case, 25 percent of the remaining bandwidth after the LLQ is serviced.

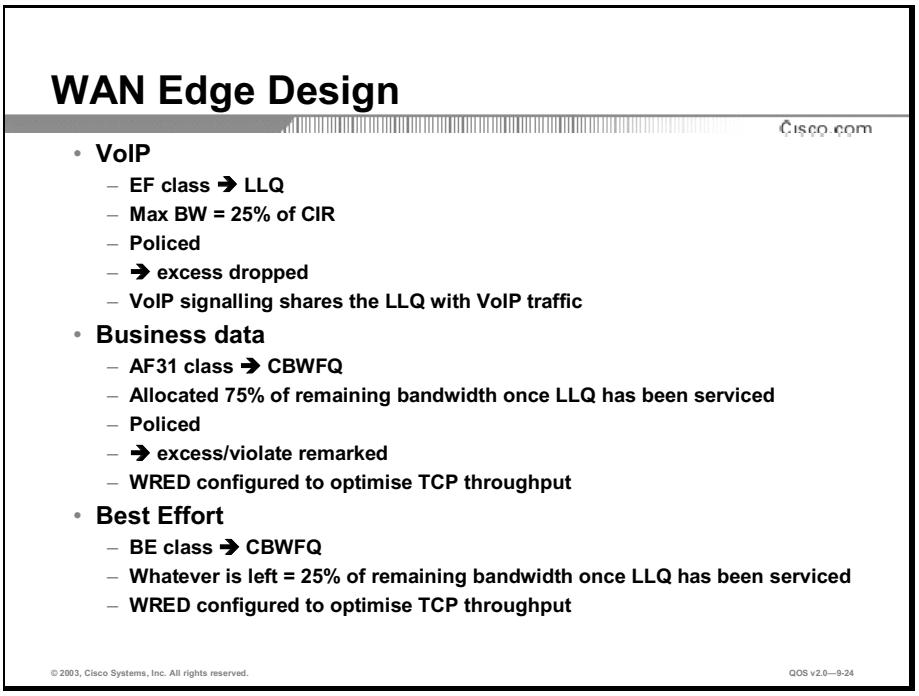

This slide highlights the required PHB for each of the three traffic classes supported by the service provider.

For the premium (VoIP) traffic class, the VoIP packets will be marked with EF and go into the LLQ with the following parameters:

- The LLQ will be policed and have a maximum bandwidth of 25 percent of the CIR.
- All excess traffic will be dropped.
- The call-signaling traffic will share the LLQ with the VoIP bearer traffic.

For the business (mission-critical) traffic class, the business class packets will be marked with AF 31 and go into the CBWFQ:

- This class will be policed and have a minimum bandwidth guarantee of 75 percent of the remaining available bandwidth.
- All exceeding and violating traffic will be re-marked then sent.
- WRED will be used on this traffic class to optimize TCP throughput.

For the default (best effort) traffic class, the best-effort class packets will be marked with DSCP of 0:

- This class is not policed and has a minimum bandwidth guarantee of 25 percent of the remaining available bandwidth.
- WRED will be used on this traffic class to optimize TCP throughput.

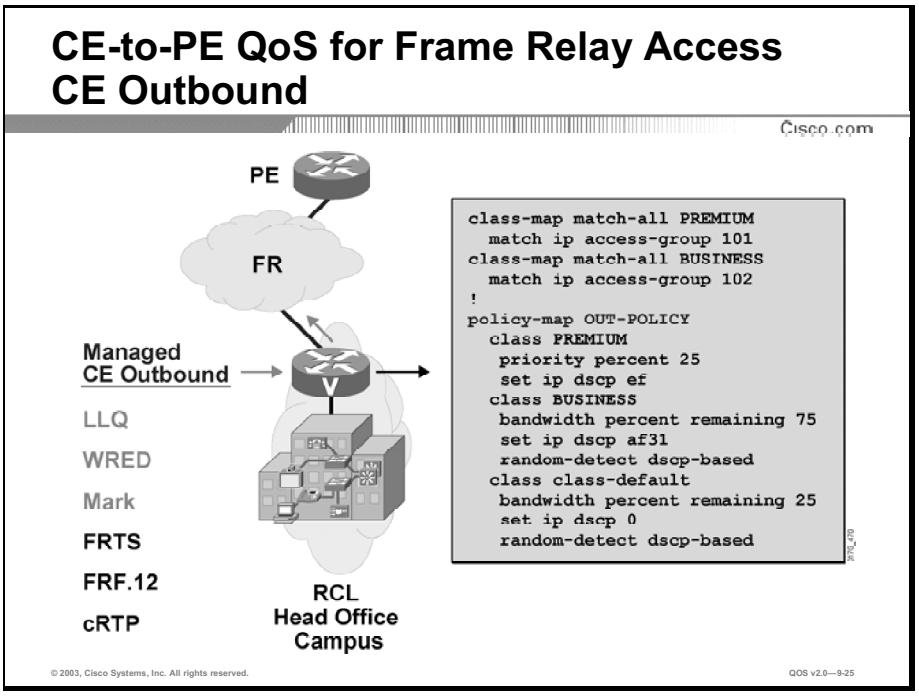

This figure shows the QoS configurations on the managed CE router outbound interface to implement the required QoS policy required for each of the three service provider traffic classes. The traffic policy called “OUT-POLICY” is configured to provide the LLQ/CBWFQ and WRED. Each traffic class bandwidth guarantee is configured using a percentage rather than a fixed bandwidth in kbps.

Both the enterprise voice bearer and voice control traffic will be matched using access-list 101 and will be serviced by the LLQ with a maximum bandwidth guarantee of 25 percent of the link bandwidth.

The enterprise mission-critical traffic will be matched using access-list 102 and has a minimum guarantee bandwidth of 75 percent of the remaining available bandwidth after the LLQ has been serviced.

All other traffic from the enterprise will be classified into the default class and has a minimum guarantee bandwidth of 25 percent of the remaining available bandwidth after the LLQ has been serviced.

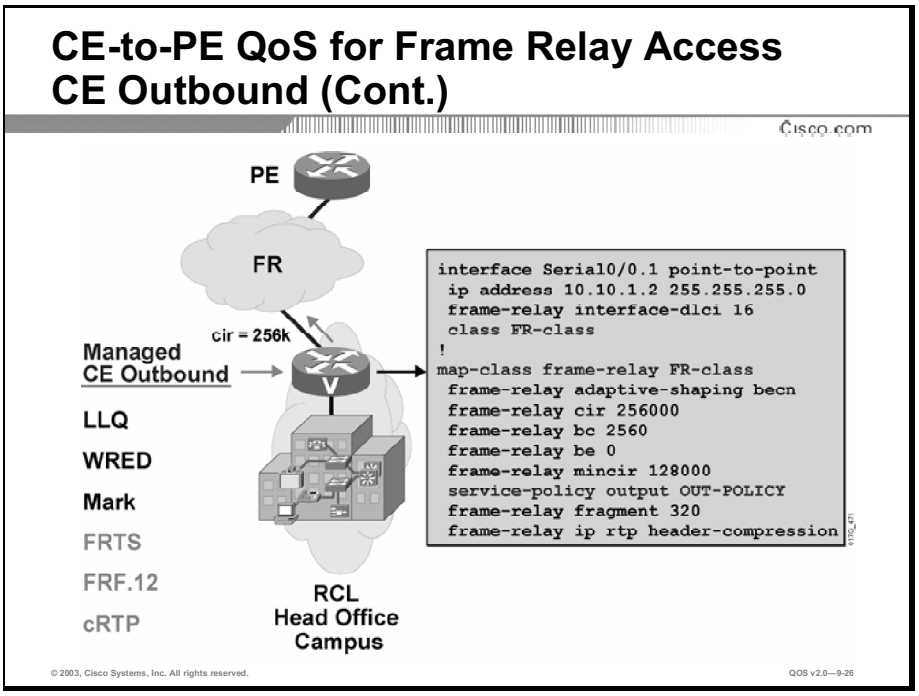

In this example, the CE and PE link is a Frame Relay link and traffic shaping is implemented on the PVC using FRTS.

FRTS is configured using a Frame Relay map class with a committed information rate (CIR) of 256 kbps, a Bc of 2560 bits, an excess burst (Be) of 0 (no bursting), and a minimum CIR (mincir) of 128 kbps. The CIR is the rate you want to normally send at when there is no congestion. The CIR needs to be remote end link speed or the actual CIR of the virtual circuit. The Bc is the amount you will send per time interval. The CIR and Bc will be used to compute a committed time window (Tc), where Tc = Bc/CIR. For FRTS, CLI will only allow Bc values that would result in a Tc that is between 10 ms and 125 ms. A recommended value for Tc is 10 ms. To get a Tc of 10 ms, the Bc should set to 1/100 of the CIR. The mincir is the lower bound value that you will slow down to during congestion. If adaptive shaping is enabled, the mincir must be larger than required voice bandwidth.

FRF.12 fragmentation and interleaving and cRTP are also enabled within the Frame Relay map class. The fragment size in bytes is set to derive 10 ms to 15 ms delay maximum. The fragment size should be same on both ends.

The “OUT-POLICY” traffic policy defined on the previous slide is applied within the Frame Relay map class.

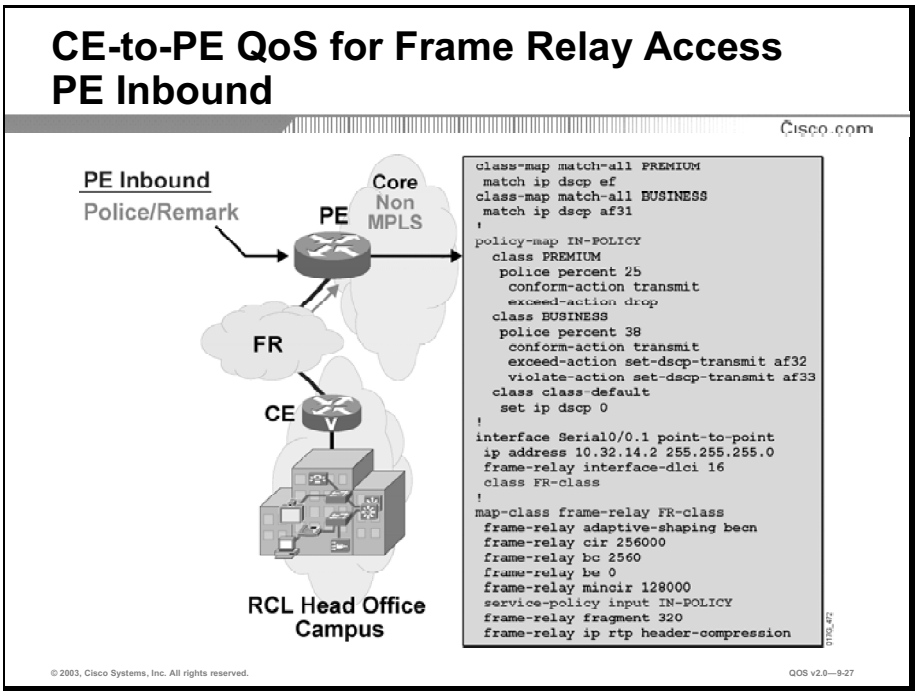

This figure shows the QoS configurations on the ingress PE router inbound interface to implement the required QoS policy that is required for each of the three service provider traffic classes.

In this case, a traffic policy called “IN-POLICY” is configured to provide the required class-based policing. For the premium class, the rate limit is set to 25 percent of the link bandwidth. All exceeding premium class traffic is dropped. For the business class, the rate limit is set to 38 percent of the link bandwidth. All exceeding and violating business class traffic are re-marked with a higher drop probability then sent. The default class is not policed.

The “IN-POLICY” traffic policy is applied within the Frame Relay map class.

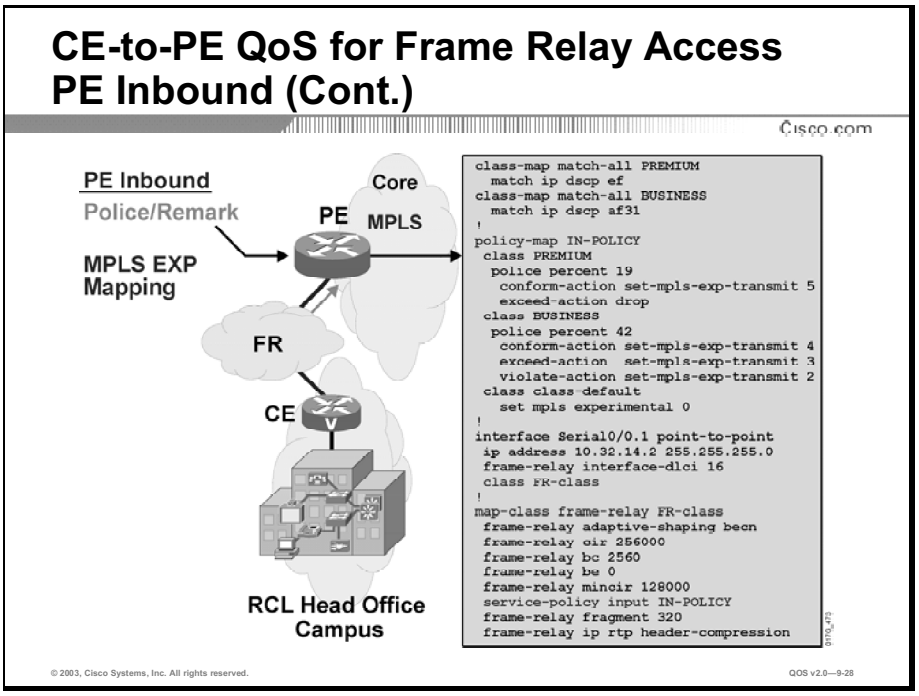

This figure is similar to the previous figure except that the service provider core is now Multiprotocol Label Switching (MPLS) instead of IP. The differences here are the markings within the traffic policy. In this case, a traffic policy called “IN-POLICY” is configured to provide the required class-based policing. The premium and business class traffic is policed and then marked with the proper MPLS experimental bits:

- All conforming premium traffic is marked with MPLS experimental 5 then sent.
- All exceeding premium traffic is dropped.
- All conforming business traffic is marked with MPLS experimental 4 then sent.
- All exceeding business traffic is marked with MPLS experimental 3 then sent.
- All violating business traffic is marked with MPLS experimental 2 then sent.
- The best-effort class traffic is not policed but just marked with MPLS experimental 0.

The “IN-POLICY” traffic policy is applied within the Frame Relay map class.

Service Provider Backbone QoS Implementations

This third and last part of the case study explains some of the best practice QoS implementations and configurations on the service provider IP core PE and P routers.

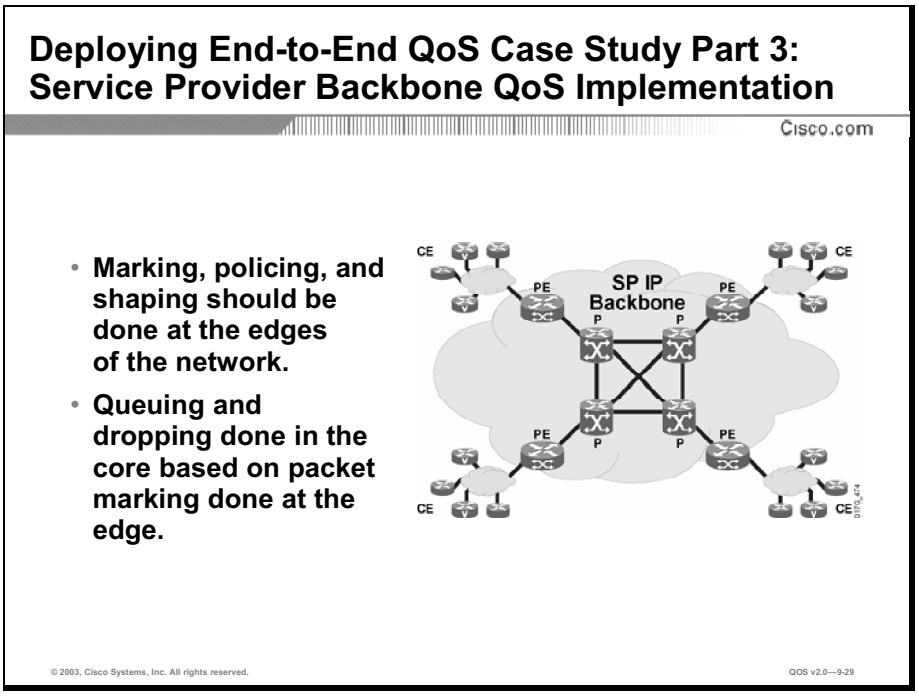

The figures will go through the typical QoS configurations required at the PE and P routers within the service provider IP core network.

The service provider IP core is used to provide high-speed packet transport. Therefore, all the markings, policing, and shaping should be performed only at the PE router on the PE-to-CE link and not at the core.

Using the DiffServ model, only the edge requires complex QoS policy. At the core, only queuing and WRED are required. The operation of queuing and WRED will be based on the markings done at the edge (PE).

It is common to implement GSRs within the service provider IP core. With the GSRs, the queuing mechanism is MDRR. If a router like the 7500 is used in the core, then LLQ will be the queuing mechanism. (In this case study, only the 7500 LLQ/WRED configurations will be discussed.)

Service Provider Backbone

- **Overprovisioning best-effort backbone is an alternative but, has these drawbacks:**
  - **Expensive**
  - **Fate sharing**
  - **Planning mistakes**
  - **Failure conditions**
  - **Unexpected traffic demand**
- **DiffServ backbone is better**

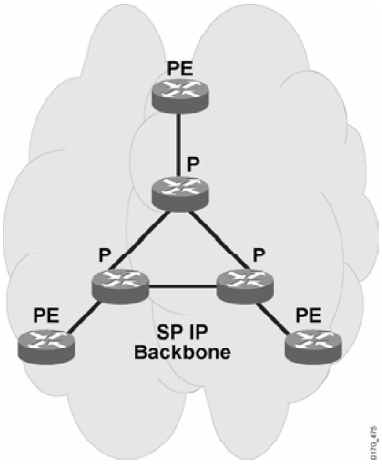

Two of the IP backbone design methods include a best-effort backbone with overprovisioning and a DiffServ backbone.

The more traditional approach is to use a best-effort backbone with overprovisioning. However, to meet the application needs of today (VoIP, videoconferencing, e-learning, and so on), deploying a DiffServ backbone and offering different SLAs for the different traffic classes can greatly reduce the cost and improve the delay, jitter, and packet loss and meet network QoS requirements.

With overprovisioning, service provider typically uses an overprovisioning factor of 2. For example, if the aggregate traffic load on the network is 10 Gbps, then the network is provisioned for 20 Gbps of maximum capacity. Some of the problems with a best-effort backbone with overprovisioning include:

- If the capacity planning is not accurate and congestion occurs, because the traffic types are not differentiated, the VoIP packets will not be treated with higher priority than other data traffic resulting in suboptimal treatment for VoIP packets.
- Overprovisioning for all traffic is very expensive to implement.
- During capacity planning, maybe not all the failure scenarios are analyzed. Unplanned failures can cause unexpected congestions in the network. A link or node failure leading to traffic re-routing can take up all the excess capacity.
- The network can experience unexpected traffic demands, which can cause congestions in the network.
- Denial of Service (DOS) attacks on one service will affect all other services.

Using DiffServ, the traffic is isolated into different classes, and each traffic class is provisioned with a different traffic policy that is based on the QoS requirements of the traffic class. This will reduce the cost and provide better overall latency, delay, and jitter.

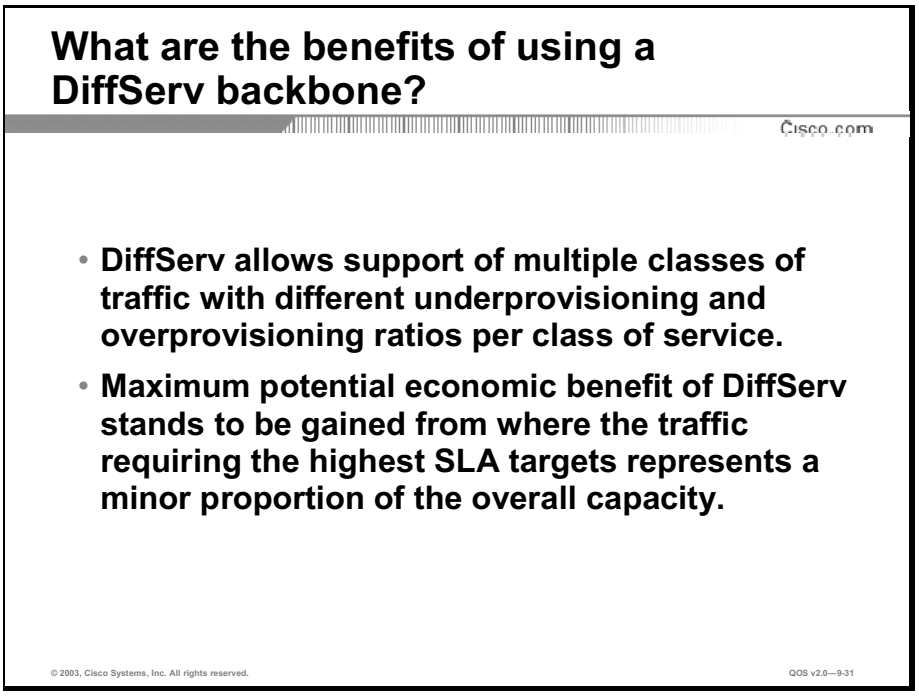

This figure lists two of the main benefits of using DiffServ over a best-effort with overprovisioning backbone.

Instead of overprovisioning based on the aggregate bandwidth of all traffic, with DiffServ, each traffic class can be designed with different provisioning ratios. If the premium traffic class is only 20 percent of the total capacity, then overprovisioning for 100 percent of the traffic load is expensive and not necessary in order to guarantee the QoS requirements for the premium traffic class.

Let us examine an overprovisioning example with and without DiffServ:

- If the total aggregate bandwidth is 10 Gbps, where the premium class = 2 Gbps, the business class = 3 Gbps, and the default class = 5 Gbps:
  - With best effort and an overprovisioning ratio of 2:1, the provisioned bandwidth = 10G * 2= 20 Gbps.
  - With DiffServ, and the premium class having an overprovisioning ratio of 2:1, the business class having a lower overprovisioning ratio of 1.5:1, and the default class not having any overprovisioning, the provisioned bandwidth = (2 G * 2) + (3 G * 1.5) + 5 G = 13.5Gbps.

By isolating the traffic into different traffic classes then treating the different traffic classes with different PHBs, DiffServ can reduce the bandwidth requirement on the network while achieving the same SLA when compared to the non-DiffServ case.

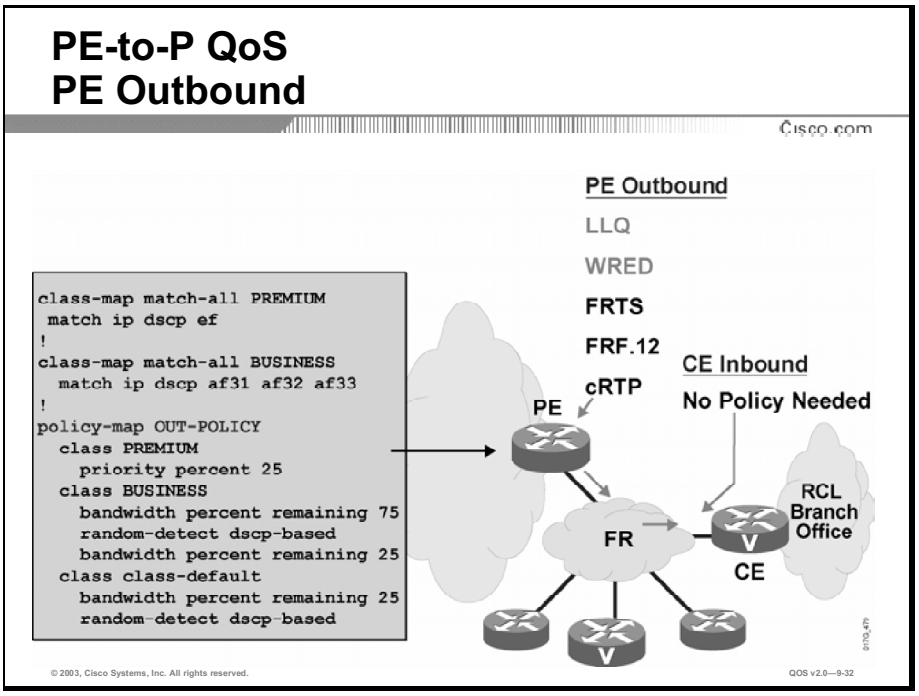

The figures will go through the typical QoS configurations required at the service provider core PE and P routers.

The complex QoS policies with classification and markings, policing, shaping, LFI, cRTP are only required at the edge. In the core, only LLQ (or MDRR for GSR) and WRED are needed.

This figure shows the QoS configurations on the ingress PE router outbound interface to the P router.

In this case, a traffic policy called “OUT-POLICY” is configured to provide the LLQ/CBWFQ and WRED. Each traffic class bandwidth guarantee is configured using a percentage rather than a fixed bandwidth in kbps.

No inbound policy is required on the P router.

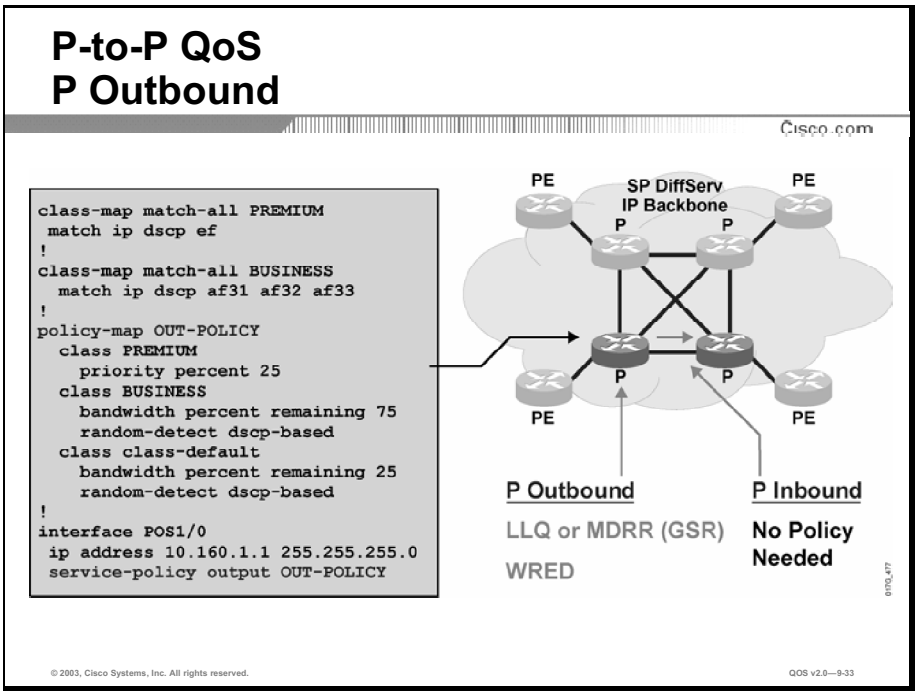

This figure shows the QoS configurations on the P router outbound interface to another P router.

In this case, a traffic policy called “OUT-POLICY” is configured to provide the LLQ/CBWFQ and WRED. Each traffic class bandwidth guarantee is configured using a percentage rather than a fixed bandwidth in kbps.

No inbound policy is required on the receiving P router.

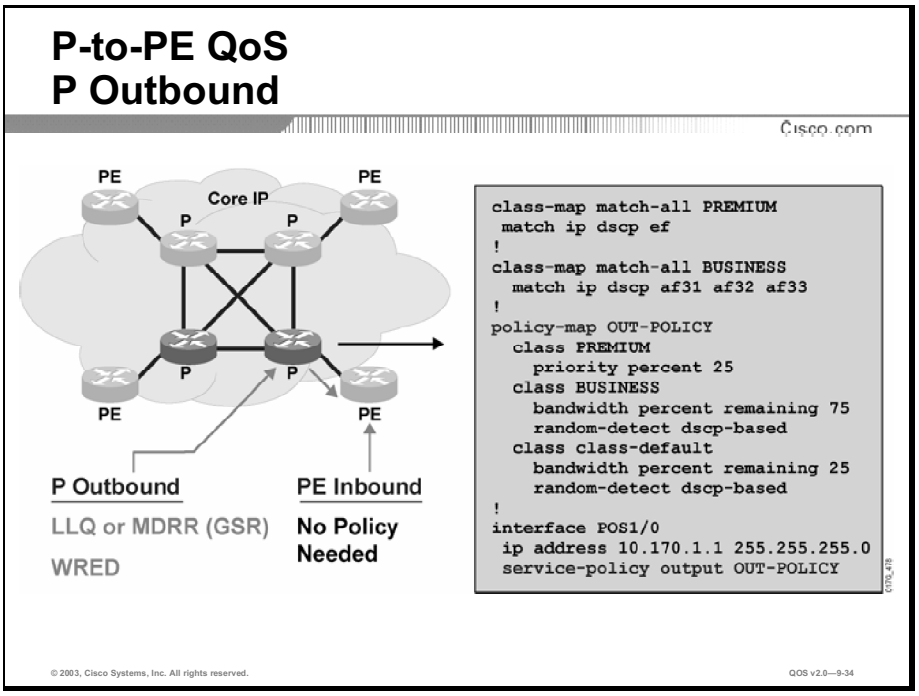

This figure shows the QoS configurations on the P router outbound interface to the egress PE router.

In this case, a traffic policy called “OUT-POLICY” is configured to provide the LLQ/CBWFQ and WRED. Each traffic class bandwidth guarantee is configured using a percentage rather than a fixed bandwidth in kbps.

No inbound policy is required on the egress PE router.

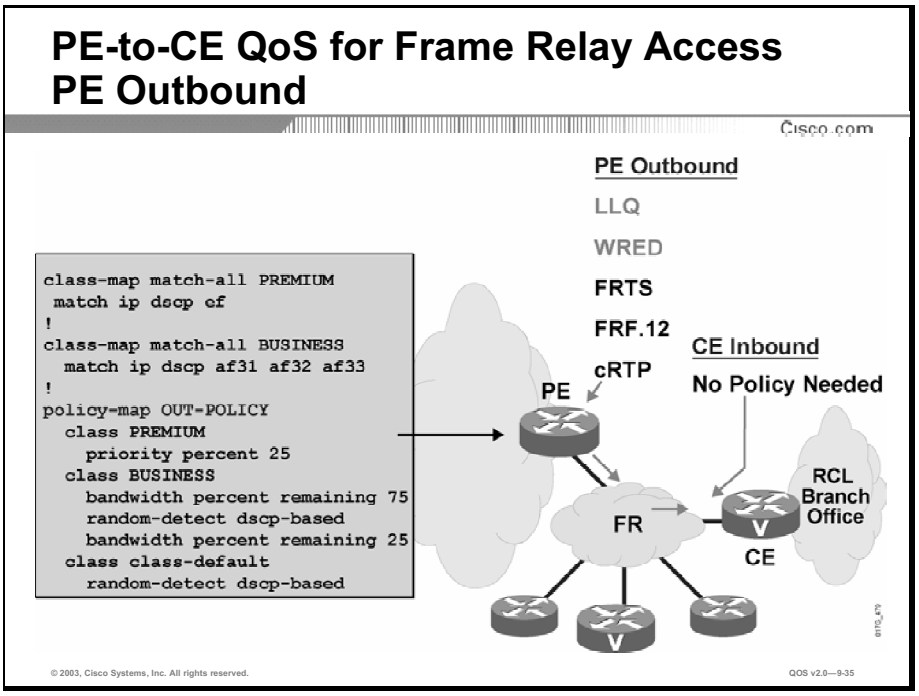

This is the continual discussion of the PE CE WAN edge QoS requirements.

This figure shows the QoS configurations on the egress PE router outbound interface to implement the required QoS policy required for each of the three service provider traffic classes (as previously discussed).

In this case, a traffic policy called “OUT-POLICY” is configured to provide the LLQ/CBWFQ and WRED. Each traffic class bandwidth guarantee is configured using a percentage rather than a fixed bandwidth in kbps.

No inbound policy is required on the CE router.

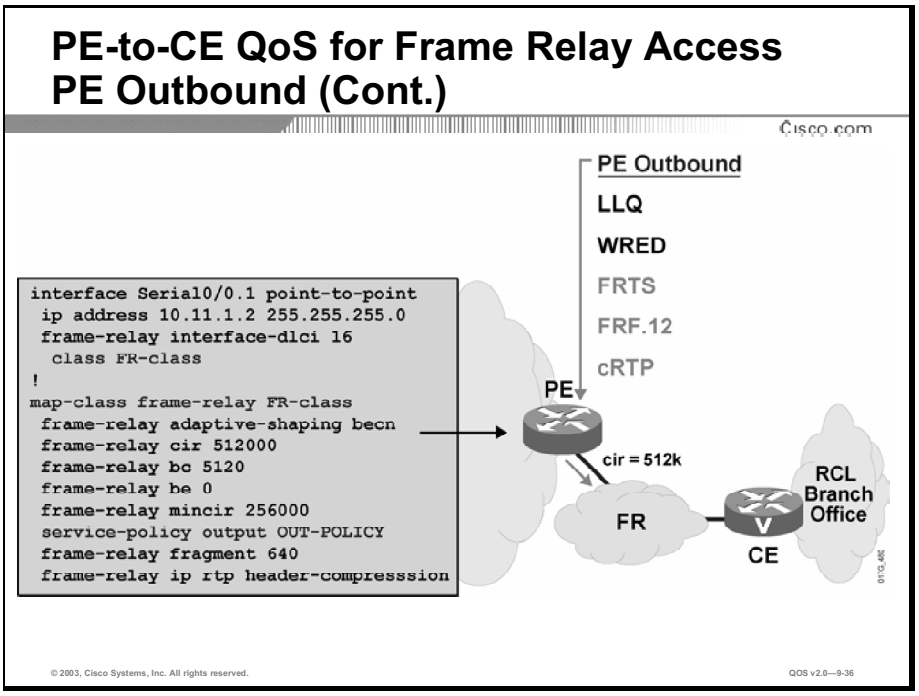

In this example, the PE-CE link is a Frame Relay link and Frame Relay traffic shaping (FRTS) is enabled on the permanent virtual circuit (PVC).

FRTS is configured using a Frame Relay map class with a CIR of 512 kbps, a Bc of 5120 bits, a Be of 0 (no bursting), and a mincir of 256 kbps. The CIR is the rate you want to normally send at when there is no congestion. The CIR needs to be remote end link speed or the actual CIR on virtual circuit. The Bc is the amount you will send per time interval. The CIR and Bc will be used to compute a Tc, where Tc = Bc/CIR. For FRTS, CLI will only allow Bc values that would result in a 125 ms < Tc <10 ms. A recommended value for Tc is 10 ms. To get a Tc of 10 ms, the Bc should set to 1/100 of the CIR. The mincir is the lower bound value you will slow down to during congestion. If adaptive shaping is enabled, the mincir must be larger than required voice bandwidth.

FRF.12 fragmentation and interleaving and cRTP are also enabled within the Frame Relay map class. The fragment size in bytes is set to derive 10 ms to 15 ms delay maximum. The fragment size should be same on both ends.

The “OUT-POLICY” traffic policy defined on the previous slide is applied within the Frame Relay map class.

Summary

This topic summarizes the key points discussed in this lesson.

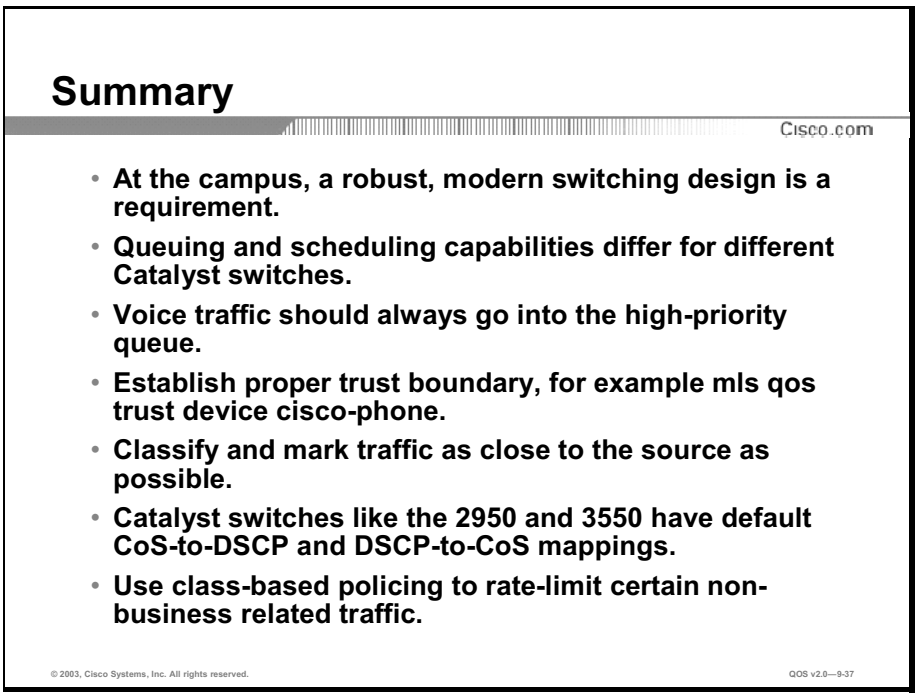

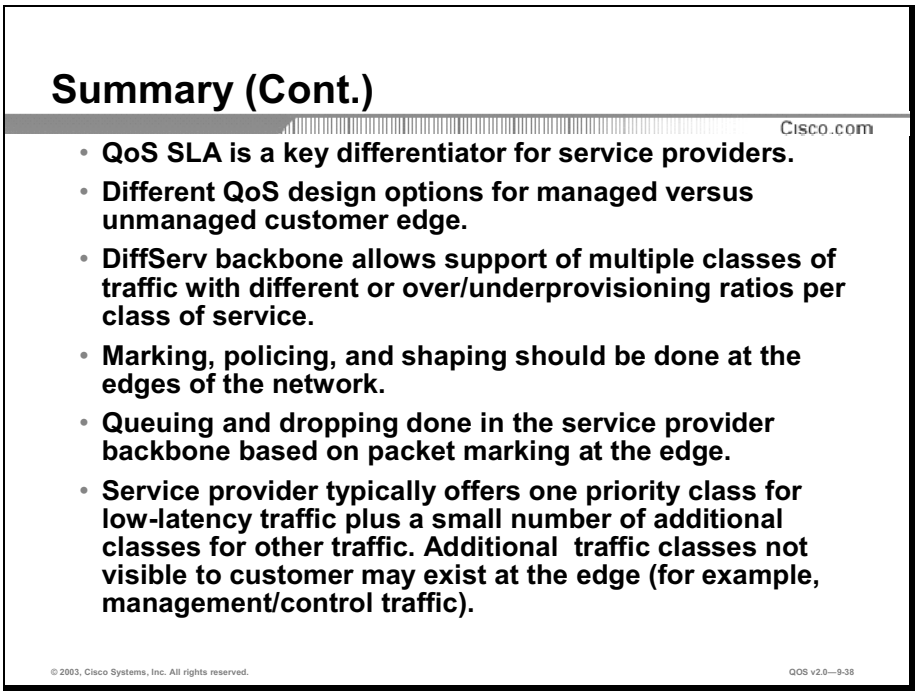

References

For additional information, refer to these resources:

- For more information on Enterprise QoS Design, refer to “Cisco AVVID Network Infrastructure Enterprise Quality of Service Design” at the following URL: http://www.cisco.com/application/pdf/en/us/guest/netsol/ns17/c649/ccmigration_09186a00800d67ed.pdf
- For more information on end-to-end QoS implementation, refer to “Implementing DiffServ for End-to-End Quality of Service Overview” at the following URL: http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122cgcr/fqos_c/fqcprt7/qcfdfsrv.pdf

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) What is an advantage of using a DiffServ backbone over a best-effort backbone?

A) provide better guarantee for the best-effort traffic class

B) simplify configurations on the core backbone routers

C) reduce the cost by using different overprovisioning ratio for different traffic class

D) provide more deterministic network response by allowing the use of a single overprovisioning ratio for all aggregated traffic

Q2) At which two layers of the campus network is traffic policing typically implemented? (Choose two.)

A) core layer

B) access layer

C) distribution layer

D) data center layer

Q3) Which one of the following is correct about traffic classification?

A) It is best practice to classify and mark traffic at the core layer switches.

B) It is best practice to classify and mark traffic as close to the destination as possible.

C) It is best practice to classify and mark traffic as close to the source as possible.

D) It is best practice to classify and mark traffic using an external server running NBAR.

Q4) Which two QoS mechanisms are typically deployed at the service provider core (non-GSR) routers? (Choose two.)

A) LFI

B) LLQ

C) WRED

D) class-based shaping

E) class-based RTP header compression

Q5) What two queue types can the Catalyst 2950 queues be configured as? (Choose two.)

A) WFQ

B) FIFO

C) PQ

D) CQ

E) WRR

Q6) With managed CE service, what inbound QoS policy is typically required at the PE router?

A) none

B) LFI and cRTP

C) traffic shaping

D) traffic policing

E) classification/marking

Q7) When would an SLA be set up per PVC or per VLAN?

A) if using managed CE services

B) if the service provider is using a DiffServ backbone

C) if a single physical interface is serving one customer

D) if a single physical interface is serving many different customers

E) if the service provider is using a best-effort backbone with overprovisioning

Q8) What is used by service providers to provide a contractual assurance for parameters such as delay, jitter, packet loss, throughput, and availability?

_____

Q9) What are two possible causes of congestion on LAN switch ports? (Choose two.)

F) port speed mismatch

G) ISL or 802.1Q trunking overhead

H) full duplex communications on trunk ports

I) excessive delays caused by CSMA/CD operations

J) many access ports being aggregated into a single uplink

Quiz Answer Key

Q1) C

Relates to: Service Provider Backbone QoS Implementations

Q2) B, C

Relates to: Enterprise Campus QoS Implementations

Q3) C

Relates to: Enterprise Campus QoS Implementations

Q4) B, C

Relates to: Service Provider Backbone QoS Implementations

Q5) C, E

Relates to: Enterprise Campus QoS Implementations

Q6) A

Relates to: WAN Edge (CE/PE) QoS Implementations

Q7) D

Relates to: QoS Service Level Agreements

Q8) An IP QoS Service Level Agreement

Relates to: QoS Service Level Agreements

Q9) A, E

Relates to: Enterprise Campus QoS Implementations

Module Assessment

Overview

Use this assessment to test what you learned in this module. The correct answers and solutions are found in the Module Assessment Answer Key.

Quiz: QoS Best Practices

Complete the Quiz to assess what you have learned in the module.

Objectives

This activity tests your knowledge on how to meet these objectives:

- Correctly identify and describe the set of classification practices that most closely represent Cisco QoS "best practices"
- Correctly identify and describe the set of QoS mechanisms used to implement Cisco end-to end QoS "best practices" in a typical enterprise network connected through a service provider providing Layer 3 IP services

Instructions

Complete these steps:

Step 1 Answer all questions in this quiz by selecting the best answer(s) to each question.

Step 2 Verify your results against the answer key located at the end of this section.

Step 3 Review the topics in this module that relate to the questions that you answered incorrectly.

Q1) Which three of the following QoS tools is used to manage the delay, delay variation (jitter), bandwidth, and packet-loss parameters on a network? (Choose three.)

A) coding/decoding (codec) tools
B) flow-control (windowing) tools
C) scheduling tools
D) provisioning tools
E) classification tools

Q2) Which of the following statements is true regarding the bandwidth requirement of data applications?

A) It is not possible to provide a blanket rule for provisioning data bandwidth because the bandwidth requirements vary greatly from application to application (and even between versions of the same applications).
B) All ERP applications have similar bandwidth requirements.
C) All data applications should be provisioned a minimum bandwidth guarantee matching their data rate + 20 percent.
D) ERP applications should use the LLQ and be provisioned a maximum bandwidth guarantee matching their data rate.

Q3) List three QoS requirements for voice. (Choose three.)

A) latency $\leq$ 150 ms
B) jitter $\leq$ 30 ms
C) loss $\leq$ 1 percent
D) bandwidth = 80 kbps per call minimum for all codecs
E) retransmit interval $\leq$ 10 ms

Q4) Which three of the following statements is a general guideline for implementing campus QoS? (Choose three.)

A) Use multiple queues on the transmit interfaces to minimize the potential for dropped or delayed traffic caused by transmit buffer congestion.

B) Classification or marking should be done at the high-speed core layer.

C) Most of the more complex QoS configurations, like traffic shaping, header compression, LFI, LLQ, and WRED, occur at the (CE) WAN edge router.

D) Establish proper trust boundary. For example, at the access layer switches, only trust the IP Phone CoS marking, but not the PC CoS marking.

E) Always use NBAR to classify traffic because all Cisco Catalyst switches support NBAR.

F) QoS in the campus is generally not a buffer management issue as much as it is a bandwidth management issue. Therefore, link efficient mechanisms should be implemented between the access and distribution layer links.

Q5) The Cisco Catalyst 2950 switch has four transmit queues with a single drop threshold. By default, the transmit queues are serviced with which type of scheduling?

A) Priority

B) WRR

C) MDRR

D) FIFO

Q6) What are two typical PHB behaviors required for the EF traffic class? (Choose two.)

A) policed to a maximum bandwidth by LLQ

B) guaranteed a minimum bandwidth by CBWFQ

C) the dropping method should be tail-drop

D) the dropping method should be WRED to prevent congestion

Q7) Based on the following configuration, which two statements are true? (Choose two.)

```
class-map match-all PREMIUM
  match ip access-group 101
class-map match-all BUSINESS
  match ip access-group 102
!
policy-map OUT-POLICY
  class PREMIUM
   priority percent 25
   set ip dscp ef
  class BUSINESS
   bandwidth percent remaining 75
   set ip dscp af31
   random-detect dscp-based
  class class-default
   bandwidth percent remaining 25
   set ip dscp 0
   random-detect dscp-based
```

A) The premium traffic class has a maximum bandwidth guarantee equal to 25 percent of the available interface bandwidth.

B) The business traffic class has a maximum bandwidth guarantee equal to 75 percent of the interface bandwidth.

C) The default traffic class has a maximum bandwidth guarantee equal to 25 percent of the interface bandwidth.

D) Tail drop is used on the business traffic class.

E) WRED is used on the premium traffic class.

F) WRED is used on the default traffic class.

Q8) What is wrong with the following configuration?

```
class-map match-all PREMIUM
  match ip dscp ef
class-map match-all BUSINESS
  match ip dscp af31
!
policy-map IN-POLICY
!
class PREMIUM
  police percent 25
   exceed-action drop
!
class BUSINESS
  police percent 56
   exceed-action set-mpls-exp-transmit 3
   violate-action set-mpls-exp-transmit 2
!
 class class-default
   set mpls experimental 0
!
interface Serial0/0.1 point-to-point
 ip address 10.32.14.2 255.255.255.0
 frame-relay interface-dlci 16
 service-policy input IN-POLICY
```

Q9) If the service provider core network is using Cisco GSR as the P routers, what will be the queuing mechanism used on the Cisco GSR P routers?

A) MDRR
B) LLQ
C) CBWFQ
D) D-WFQ
E) IP RTP Priority

Q10) The QoS requirements on the CE and PE router differ depending on what factor?

A) whether or not the PE router is managed by the service provider or not
B) whether or not the CE router is managed by the service provider or not
C) whether or not the service provider is using an MPLS core or not
D) the number of traffic classes supported by the service provider
E) the SLAs offered by the service provider

Scoring

You have successfully completed the quiz for this lesson when you earn a score of 80 percent or better.

Module Assessment Answer Key

Q1) C, D, E

Relates to: Traffic Classification Best Practices

Q2) A

Relates to: Traffic Classification Best Practices

Q3) A, B, C

Relates to: Traffic Classification Best Practices

Q4) A, C, D

Relates to: Case Study: End-to-End QoS

Q5) A

Relates to: Case Study: End-to-End QoS

Q6) A, C

Relates to: Case Study: End-to-End QoS

Q7) A, F

Relates to: Case Study: End-to-End QoS

Q8) The conform actions are missing from the configuration.

Relates to: End-to-End QoS

Q9) A

Relates to: Case Study: End-to-End QoS

Q10) B

Relates to: Case Study: End-to-End QoS

Module Summary

This topic summarizes the key points discussed in this module.

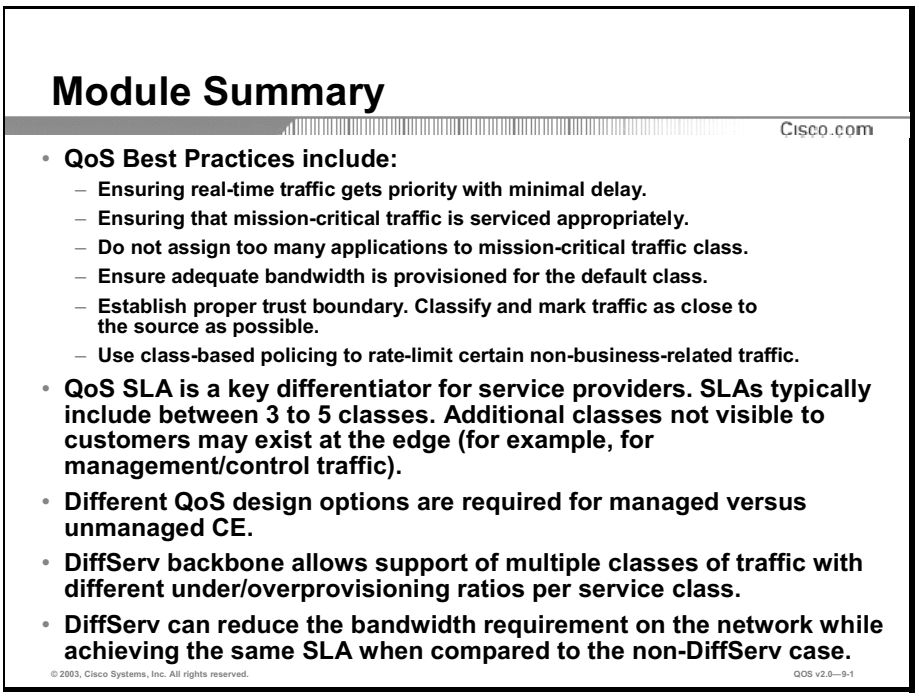

When an enterprise network is connected by a service provider that provides Layer 3 IP services, both the enterprise and the service provider must implement the proper IP QoS mechanisms. This is necessary to satisfy the end-to-end QoS requirements of the different applications running at the enterprise. A key differentiator for service providers when offering Layer 3 IP services is to offer SLAs for each of the traffic classes supported by the service provider. SLAs provide the needed assurance to the enterprise customers that the service provider network can meet the QoS requirements of the enterprise applications traversing across the service provider network.

This module discusses some of the key best practices for deploying IP QoS at the enterprise and at the service provider network. The first lesson of the module included a baseline traffic classification recommendation detailing the PHB, DSCP, IP precedence, and Layer 2 CoS values for some of the common enterprise traffic types. The recommendation can be referenced by network engineers and designers as a guide for overall system design and QoS feature implementation.

The second lesson of this module provides a long case study illustrating the typical IP QoS implementations at the enterprise campus LAN, at the WAN edge (CE and PE routers), and at the service provider core.

Within the enterprise campus LAN, a hierarchical design is required along with proper buffer management at the LAN switches. Classification and marking should be performed as close to the source as possible. Class-based policing can be implemented to rate-limit certain non-business-related applications (like peer-to-peer file sharing applications, Napster, and so on).

More complex QoS configurations will be required at the WAN edge. The QoS configuration required at the WAN edge CE and PE routers is different depending on whether or not the CE is managed by the service provider or by the enterprise customer. LLQ/CBWFQ will be required for congestion management. Traffic shaping will be required on Frame Relay PVCs. Link efficient mechanisms like LFI and cRTP will be required to improve the WAN link efficiency. Class-based marking is also required if the enterprise customer traffic classes need to be remapped into the service provider traffic classes.

At the service provider high-speed core, typically only WRED or LLQ/MDRR are required, based on the previous packet QoS marking. In the service provider core (backbone), two of the design options include implementing a best-effort backbone with overprovisioning or implementing a DiffServ backbone with a different provisioning ratio for each traffic class. By deploying a DiffServ backbone, service providers can significantly reduce their cost while providing the required SLAs for their customers.

Course Glossary

The Course Glossary for *Implementing Cisco Quality of Service* (QOS) v2.0 highlights and defines key terms and acronyms used throughout this course. Many of these terms are also described in the Cisco Internetworking Terms and Acronyms resource, available via http://www.cisco.com.

Acronym or Term	Short Definition for Mouseover	Expansion and Definition of Acronym or Term	Source for Definition	Additional Explanation/ Description for this Course Context
802.1p	IEEE specification for defining priority.	Defined in the IEEE 802.1Q specification as the priority field. Within the Tag Control Information field, the priority field is 3-bits and hence represents up to 8 priority levels. The IEEE 802.1p specification defines the use of these 3 bits.	http://www.cisco.com/univercd/cc/td/doc/product/software/ios121/121newft/121t/121t3/dtbridge.htm	
802.1Q	IEEE specification for tagging Ethernet frames with VLAN information.	IEEE specification for inserting VLAN membership information into Ethernet frames. The VLAN information is inserted into the frame in a field called the Tag Control Information that contains VLAN and QoS priority identification.	http://www.cisco.com/univercd/cc/td/doc/product/software/ios121/121newft/121t/121t3/dtbridge.htm	
ACK	acknowledgment. Notification indicating successful receipt of a message.	acknowledgment. Notification sent from one network device to another to acknowledge that some event occurred (for example, the receipt of a message). Sometimes abbreviated ACK.	ITA/Jan 2001	
ACL	access control list. Filter list used for services such as security, QoS, and routing.	access control list. A list kept by routers to control access to or from the router for a number of services. (For example, to prevent packets with a certain IP address from leaving a particular interface on the router).	ITA/Jan 2001	
AF	Assured Forwarding	Assured Forwarding. A DSCP PHB providing for the delivery of IP packets in four independently forwarded AF classes. Within each AF class, an IP packet can be assigned one of three different levels of drop precedence: low, medium and high.	IETF RFC 2597	
AH	authentication header. Security protocol providing data authentication.	authentication header. A security protocol that provides data authentication and optional anti-replay services. AH is embedded in the data to be protected (a full IP datagram).	ITA/Jan 2001	
ATM	Asynchronous Transfer Mode. International standard for cell relay.	Asynchronous Transfer Mode. The international standard for cell relay in which multiple service types (such as voice, video, or data) are conveyed in fixed-length (53-byte) cells. Fixed-length cells allow cell processing to occur in hardware, thereby reducing transit delays. ATM is designed to take advantage of high-speed transmission media, such as E3, SONET, and T3.	ITA/Jan 2001	

Acronym or Term	Short Definition for Mouseover	Expansion and Definition of Acronym or Term	Source for Definition	Additional Explanation/ Description for this Course Context
BA	behavior aggregate	behavior aggregate. A collection of packets with the same DiffServ code point crossing a link in a particular direction.	IETF RFC 2475	
backbone		Part of a network that acts as the primary path for traffic that is most often sourced from, and destined for, other networks.	ITA/Jan 2001	
bandwidth	The rated throughput capacity of a given network medium.	The difference between the highest and lowest frequencies available for network signals. The term also is used to describe the rated throughput capacity of a given network medium or protocol. The frequency range necessary to convey a signal measured in units of hertz (Hz). For example, voice signals typically require approximately 7 kHz of bandwidth and data traffic typically requires approximately 50 kHz of bandwidth.	ITA/Jan 2001	
Bc	committed burst	committed burst. Negotiated tariff metric in Frame Relay internetworks. The maximum amount of data (in bits) that a Frame Relay internetwork is committed to accept and transmit at the committed information rate (CIR).	ITA/Jan 2001	
Be	excess burst	excess burst. Negotiated tariff metric in Frame Relay internetworks. The number of bits that a Frame Relay internetwork attempts to transmit after Bc is accommodated. Be data, in general, is delivered with a lower probability than Bc data because Be data can be marked as DE by the network.	ITA/Jan 2001	
BECN	backward explicit congestion notification. Frame Relay congestion signaling mechanism.	backward explicit congestion notification. Bit set by a Frame Relay network in frames traveling in the opposite direction of frames encountering a congested path. DTE receiving frames with the BECN bit set can request that higher-level protocols take flow control action as appropriate.	ITA/Jan 2001	
best-effort delivery	No guarantee made on delivery of packets.	Describes a network system that does not use a sophisticated acknowledgment system to guarantee reliable delivery of information.	ITA/Jan 2001	
BGP	Border Gateway Protocol.	Border Gateway Protocol. Interdomain routing protocol that replaces EGP. BGP exchanges reachability information with other BGP systems. It is defined by RFC 1163.	ITA/Jan 2001	

Acronym or Term	Short Definition for Mouseover	Expansion and Definition of Acronym or Term	Source for Definition	Additional Explanation/ Description for this Course Context
buffer	A storage area used for handling data in transit.	A storage area used for handling data in transit. Buffers are used in internetworking to compensate for differences in processing speed between network devices. Bursts of data can be stored in buffers until they can be handled by slower processing devices. Sometimes referred to as a packet buffer.	ITA/Jan 2001	
burst	A sequence of signals counted as one unit in accordance with some specific criterion or measure.	In data communications, a sequence of signals counted as one unit in accordance with some specific criterion or measure.	ITA/Jan 2001	
bursty traffic	An uneven pattern of data transmission.	A data communications term referring to an uneven pattern of data transmission.	ITA/Jan 2001	
CAR	committed access rate. A traffic policing and marking mechanism.	committed access rate. The CAR and DCAR (distributed CAR) services limit the input or output transmission rate on an interface or subinterface based on a flexible set of criteria.	ITA/Jan 2001	
CBWFQ	class-based weighted fair queueing	class-based weighted fair queueing. CBWFQ extends the standard WFQ functionality to provide support for user-defined traffic classes. Allows the user to define traffic classes based on customer-defined match criteria such as access control lists (ACLs), input interfaces, protocol, and quality-of-service (QoS) label. When traffic classes have been defined, they can be assigned a bandwidth, queue limit, or drop policy such as weighted random early detection (WRED).	ITA/Jan 2001 The ABCs of Cisco IOS Software Acronyms	
CDT	Congestive Discard Threshold	Congestive Discard Threshold. In WFQ, the CDT is defined as the number of messages allowed in each queue. When a conversation reaches this threshold, new message packets are discarded.	Cisco IOS 12.2 documentation: http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122cgcr/fqos_r/qrfcmd1.htm#1098249	

Acronym or Term	Short Definition for Mouseover	Expansion and Definition of Acronym or Term	Source for Definition	Additional Explanation/ Description for this Course Context
CEF	Cisco Express Forwarding	Cisco Express Forwarding. Increases performance by adopting a new caching mechanism that optimizes Internet traffic and enhances network scalability.	The ABCs of Cisco IOS Software Acronyms	
CE router	customer edge router	customer edge router. A router that is part of a customer network and that interfaces to a provider edge (PE) router.	ITA/Jan 2001	
CIR	committed information rate	committed information rate. The rate at which a Frame Relay network agrees to transfer information under normal conditions, averaged over a minimum increment of time. CIR, measured in bits per second, is one of the key negotiated tariff metrics.	ITA/Jan 2001	
Cisco IOS Software	Cisco operating system software that runs on routers and switches	Cisco IOS Software, the leading and most widely deployed network system software, delivers intelligent network services on a flexible networking infrastructure that enables the rapid deployment of Internet applications.	The ABCs of Cisco IOS Software Acronyms	
class map	In the Modular QoS CLI, the class-map command is used to define a traffic class.	In the Modular QoS CLI, the class-map command is used to define a traffic class. The purpose of a traffic class is to classify traffic.	Cisco IOS 12.2 documentation: http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122cgcr/fqos_c/fqcprt8/qcfmdcli.htm	
classification	Identifying traffic as belonging to a specific traffic class.	Classification entails using a traffic descriptor to categorize a packet within a specific group to define that packet and make it accessible for QoS handling on the network.	Cisco IOS 12.2 documentation: http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122cgcr/fqos_c/fqcprt1/qcfclass.htm	
CLI	command-line interface	command-line interface. An interface that allows the user to interact with the operating system by entering commands and optional arguments. The UNIX operating system and DOS provide CLIs.	ITA/Jan 2001	

Acronym or Term	Short Definition for Mouseover	Expansion and Definition of Acronym or Term	Source for Definition	Additional Explanation/ Description for this Course Context
CLP	cell loss priority. Field in the ATM cell header that determines the probability of a cell being dropped if the network becomes congested.	cell loss priority. Field in the ATM cell header that determines the probability of a cell being dropped if the network becomes congested. Cells with CLP = 0 are insured traffic, which is unlikely to be dropped. Cells with CLP = 1 are best-effort traffic, which might be dropped in congested conditions to free up resources to handle insured traffic.	ITA/Jan 2001	
codec	coder-decoder. Used to compress / decompress speech or audio signals.	coder-decoder. Integrated circuit device that typically uses pulse code modulation to transform analog signals into a digital bit stream and digital signals back into analog signals. In Voice over IP, Voice over Frame Relay, and Voice over ATM, a DSP software algorithm used to compress/decompress speech or audio signals.	ITA/Jan 2001	
compression	Technique used to reduce the size of transmitted data.	The running of a data set through an algorithm that reduces the space required to store or the bandwidth required to transmit the data set.	ITA/Jan 2001	
confluence	The flowing together of two or more streams (of traffic)	The flowing together of two or more streams (of traffic)	Merriam-Webster online Dictionary	
conforming [traffic]	Traffic that is within predefined bandwidth limits.	[traffic] obedient or compliant (to predefined bandwidth limits)	Merriam-Webster online Dictionary	
congestion	Traffic in excess of network capacity.	Traffic in excess of network capacity.	ITA/Jan 2001	
congestion avoidance	Congestion avoidance techniques monitor network traffic loads in an effort to anticipate and avoid congestion at common network bottlenecks.	Mechanism by which an ATM network controls the traffic entering the network to minimize delays. To use resources most efficiently, lower-priority traffic is discarded at the edge of the network if conditions indicate that it cannot be delivered.	ITA/Jan 2001 and http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122cgcr/fqos_c/fqcprt3/qcfconav.htm	
converged network	The integration of the telephone system with IP-based data networks.	The integration of the telephone system with IP-based data networks.	Merriam-Webster online Dictionary	

Acronym or Term	Short Definition for Mouseover	Expansion and Definition of Acronym or Term	Source for Definition	Additional Explanation/ Description for this Course Context
CoS	class of service	class of service. An indication of how an upper-layer protocol requires a lower-layer protocol to treat its messages.	ITA/Jan 2001	
CQ	custom queuing	custom queuing. A Cisco IOS congestion management feature where bandwidth is allocated proportionally for each different class of traffic. CQ allows you to specify the number of bytes or packets to be drawn from the queue, which is especially useful on slow interfaces.	http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122cgcr/fqos_c/fqcprt2/qcfconmg.htm#48685	
CS	class selector	To preserve backward-compatibility with any IP precedence scheme currently in use on the network, DiffServ has defined a DSCP value in the form xxx000, where x is either 0 or 1. These DSCP values are called Class-Selector Code Points.	Cisco IOS 12.2 documentation: http://www.cisco.com/en/US/products/sw/iosswrel/ps1839/products_feature_guide09186a0080087afb.html#xtocid2	
DE	discard eligible. Frame Relay congestion control mechanism.	discard eligible. If the network is congested, DE traffic can be dropped to ensure the delivery of higher priority traffic.	ITA/Jan 2001	
delay	The time required to move a packet from source to destination over a given path.	The time between the initiation of a transaction by a sender and the first response received by the sender. Also, the time required to move a packet from source to destination over a given path.	ITA/Jan 2001	
DiffServ	differentiated service	differentiated service. A paradigm for providing QoS on the Internet by employing a small, well-defined set of building blocks from which a variety of services can be built.	ITA/Jan 2001	
DLCI	data-link connection identifier. Value that specifies a PVC or an SVC in a Frame Relay network.	data-link connection identifier. Value that specifies a PVC or an SVC in a Frame Relay network. In the basic Frame Relay specification, DLCIs are locally significant (connected devices might use different values to specify the same connection). In the LMI extended specification, DLCIs are globally significant (DLCIs specify individual end devices).	ITA/Jan 2001	

Acronym or Term	Short Definition for Mouseover	Expansion and Definition of Acronym or Term	Source for Definition	Additional Explanation/ Description for this Course Context
DSCP	DiffServ Code Point.	Differentiated service code point. Six bits in the type-of-service (ToS) field.	The ABCs of Cisco IOS Software Acronyms	
ECN	Explicit Congestion Notification	A congestion indication for incipient congestion where the notification can sometimes be through marking packets rather than dropping them	RFC 3168 http://www.rfc-editor.org/rfc/rfc3168.txt	
EF	Expedited Forwarding	Expedited Forwarding. A DSCP PHB intended to provide a building block for low delay, low jitter and low loss services by ensuring that the EF aggregate is served at a certain configured rate.	IETF RFC 3246	
ESP	Encapsulating Security Payload. Security protocol providing data encapsulation protection.	Encapsulating Security Payload. Security protocol that provides data privacy services, optional data authentication, and anti-replay services. ESP encapsulates the data to be protected.	ITA/Jan 2001	
exceeding traffic	Traffic which goes beyond a limit	Traffic which goes beyond a limit set	Merriam-Webster online Dictionary	
FECN	forward explicit congestion notification. Frame Relay congestion signaling mechanism.	forward explicit congestion notification. Bit set by a Frame Relay network to inform DTE receiving the frame that congestion was experienced in the path from source to destination. DTE receiving frames with the FECN bit set can request that higher-level protocols take flow-control action as appropriate.	ITA/Jan 2001	
FIFO	first-in, first-out	first-in, first-out. Refers to a buffering scheme where the first byte of data entering the buffer is the first byte retrieved by the CPU. In telephony, FIFO refers to a queueing scheme where the first calls received are the first calls processed.	ITA/Jan 2001	
FIFO queueing	first-in, first-out queueing	first-in, first-out queueing. Involves buffering and forwarding of packets in the order of arrival. FIFO embodies no concept of priority or classes of traffic. There is only one queue, and all packets are treated equally. Packets are sent out an interface in the order in which they arrive.	ITA/Jan 2001	

Acronym or Term	Short Definition for Mouseover	Expansion and Definition of Acronym or Term	Source for Definition	Additional Explanation/ Description for this Course Context
flow	Stream of data traveling between two endpoints across a network.	Stream of data traveling between two endpoints across a network (for example, from one LAN station to another). Multiple flows can be transmitted on a single circuit.	ITA/Jan 2001	
Frame Relay	Industry-standard, switched data link layer protocol.	Industry-standard, switched data link layer protocol that handles multiple virtual circuits using HDLC encapsulation between connected devices. Frame Relay is more efficient than X.25, the protocol for which it generally is considered a replacement.	ITA/Jan 2001	
FRF.11	Frame Relay Forum implementation agreement for Voice over Frame Relay.	Frame Relay Forum implementation agreement for Voice over Frame Relay (v1.0 May 1997). This specification defines multiplexed data, voice, fax, DTMF digit-relay, and CAS/Robbed-bit signaling frame formats but does not include call setup, routing, or administration facilities. *See also* www.frforum.com.	ITA/Jan 2001	
FRF.12	Frame Relay Forum implementation agreement for Frame Relay Fragmentation.	The FRF.12 Implementation Agreement (also known as FRF.11 Annex C) was developed to allow long data frames to be fragmented into smaller pieces and interleaved with real-time frames. In this way, real-time voice and non real-time data frames can be carried together on lower speed links without causing excessive delay to the real-time traffic. *See also* www.frforum.com.	ITA/Jan 2001	
FRTS	Frame Relay traffic shaping	Frame Relay traffic shaping. Queueing method that uses queues on a Frame Relay network to limit surges that can cause congestion. Data is buffered and sent into the network in regulated amounts to ensure that the traffic can fit within the promised traffic envelope for the particular connection.	ITA/Jan 2001	
GRE	generic routing encapsulation. A tunneling protocol developed by Cisco for carrying traffic inside IP tunnels.	generic routing encapsulation. Tunneling protocol developed by Cisco that can encapsulate a wide variety of protocol packet types inside IP tunnels, creating a virtual point-to-point link to Cisco routers at remote points over an IP internetwork. By connecting multiprotocol subnetworks in a single-protocol backbone environment, IP tunneling using GRE allows network expansion across a single-protocol backbone environment.	ITA/Jan 2001	

Acronym or Term	Short Definition for Mouseover	Expansion and Definition of Acronym or Term	Source for Definition	Additional Explanation/ Description for this Course Context
GTS	generic traffic shaping	Generic traffic shaping. Provides a mechanism to control the traffic flow on a particular interface. It reduces outbound traffic flow to avoid congestion by constraining specified traffic to a particular bit rate (also known as the token bucket approach), while queuing bursts of the specified traffic. Thus, traffic adhering to a particular profile can be shaped to meet downstream requirements, eliminating bottlenecks in topologies with data-rate mismatches.	The ABCs of Cisco IOS Software Acronyms	
header	Control information placed before data when encapsulating that data for network transmission.	Control information placed before data when encapsulating that data for network transmission.	ITA/Jan 2001	
IEEE	Institute of Electrical and Electronics Engineers	Institute of Electrical and Electronics Engineers. Professional organization whose activities include the development of communications and network standards. IEEE LAN standards are the predominant LAN standards today.	ITA/Jan 2001	
IETF	Internet Engineering Task Force	Internet Engineering Task Force. Task force consisting of over 80 working groups responsible for developing Internet standards. The IETF operates under the auspices of ISOC.	ITA/Jan 2001	
IntServ	integrated services	The Internet integrated services framework provides the ability for applications to choose among multiple, controlled levels of delivery service for their data packets.	RFC 2210 http://www.ietf.org/rfc/rfc2210.txt	
IP	Internet Protocol. Network layer protocol in the TCP/IP stack offering a connectionless internetwork service.	Internet Protocol. Network layer protocol in the TCP/IP stack offering a connectionless internetwork service. IP provides features for addressing, type-of-service specification, fragmentation and reassembly, and security. Defined in RFC 791.	ITA/Jan 2001	
IP precedence	A 3-bit value in the type of service (TOS) byte used for assigning precedence to IP packets.	A 3-bit value in the type of service (TOS) byte used for assigning precedence to IP packets.	ITA/Jan 2001	

Acronym or Term	Short Definition for Mouseover	Expansion and Definition of Acronym or Term	Source for Definition	Additional Explanation/ Description for this Course Context
IPSec	IP Security. A framework of open standards that provides data confidentiality, data integrity, and data authentication between participating peers.	IP Security. A framework of open standards that provides data confidentiality, data integrity, and data authentication between participating peers. IPSec provides these security services at the IP layer. IPSec uses IKE to handle the negotiation of protocols and algorithms based on local policy and to generate the encryption and authentication keys to be used by IPSec. IPSec can protect one or more data flows between a pair of hosts, between a pair of security gateways, or between a security gateway and a host.	ITA/Jan 2001	
ISL	Inter-Switch Link	Inter-Switch Link. Cisco-proprietary protocol that maintains VLAN information as traffic flows between switches and routers.	Cisco.com	
jitter	The interpacket delay variance; that is, the difference between interpacket arrival and departure.	The interpacket delay variance; that is, the difference between interpacket arrival and departure. Jitter is an important QoS metric for voice and video applications.	ITA/Jan 2001	
LAN	local-area network	local-area network. High-speed, low-error data network covering a relatively small geographic area (up to a few thousand meters). LANs connect workstations, peripherals, terminals, and other devices in a single building or other geographically limited area. LAN standards specify cabling and signaling at the physical and data link layers of the OSI model. Ethernet, FDDI, and Token Ring are widely used LAN technologies.	ITA/Jan 2001	
latency	Delay used to measure packet transmission time.	1. Delay between the time a device requests access to a network and the time it is granted permission to transmit. 2. Delay between the time a device receives a frame and the time that frame is forwarded out the destination port.	ITA/Jan 2001	
LFI	link fragmentation and interleaving	Link fragmentation and interleaving reduces delay on slower-speed links by breaking up large datagrams and interleaving low-delay traffic packets with the smaller packets resulting from the fragmented datagram.	Cisco.com http://www.cisco.com/en/US/products/sw/iosswrel/ps1831/products_configuration_guide_chapter09186a00800ca6d4.html	

Acronym or Term	Short Definition for Mouseover	Expansion and Definition of Acronym or Term	Source for Definition	Additional Explanation/ Description for this Course Context
LLQ	low-latency queuing	low-latency queuing. Brings strict priority queuing to class-based weighted fair queuing (CBWFQ). Strict priority queuing allows delay-sensitive data such as voice to be de-queued and sent first (before packets in other queues are de-queued), giving delay-sensitive data preferential treatment over other traffic.	The ABCs of Cisco IOS Software Acronyms	
marking	To label so as to indicate membership in a specific traffic class.	To designate as if by a mark; to make or leave a mark on; or to label so as to indicate quality (in this case, the quality of service level associated with a packet).	Merriam-Webster online Dictionary	
MDRR	Modified Deficit Round Robin	Modified Deficit Round Robin (MDRR). A variant of Deficit Round Robin (DRR). Regular DRR selects packets from each virtual output queue in a regular round-robin mechanism, thus providing every class-of-service (CoS) queue equal scheduling into the fabric. In MDRR, all queues are also serviced in a round-robin fashion, with the exception of one of the queues.	The ABCs of Cisco IOS Software Acronyms	
metering	To measure by means of a meter.	To measure by means of a meter.	Merriam-Webster online Dictionary	
MIB	Management Information Base. Database of network management information that is used and maintained by a network management protocol.	Management Information Base. Database of network management information that is used and maintained by a network management protocol, such as SNMP or CMIP. The value of a MIB object can be changed or retrieved using SNMP or CMIP commands, usually through a GUI network management system. MIB objects are organized in a tree structure that includes public (standard) and private (proprietary) branches.	ITA/Jan 2001	
MPLS	Multiprotocol Label Switching	Multiprotocol Label Switching. Switching method that forwards IP traffic using a label. This label instructs the routers and the switches in the network where to forward the packets based on preestablished IP routing information.	ITA/Jan 2001	

Acronym or Term	Short Definition for Mouseover	Expansion and Definition of Acronym or Term	Source for Definition	Additional Explanation/ Description for this Course Context
MQC	Modular QoS CLI	Modular QoS CLI (MQC), the simple, scalable, and powerful QoS framework in IOS. The MQC allows for the clear separation of classification, from the policy applied on the classes, to the application of a QoS policy on an interface or sub-interface.	http://www.cisco.com/cmc/cc/pd/iosw/prodlit/iosq_ds.htm#xtocid4	
NBAR	network-based application recognition	network-based application recognition (NBAR). A new classification engine that can recognize a wide variety of application-level protocols, including HTTP via Universal Resource Locator/Multipurpose Internet Mail Extensions (URL/MIME) type and protocols that utilize dynamic port assignments. When the traffic is classified by NBAR, appropriate quality-of-service (QoS) policies can be applied to the traffic classes using existing Cisco IOS QoS features.	The ABCs of Cisco IOS Software Acronyms	
P	provider router	provider router. A router in the core of a service provider network that interfaces to other P routers or to a PE router. P routers only do MPLS switching across the Internet service provider (ISP) backbone based on label swapping of the first-level label to the BGP next-hop address.		
packet	Logical grouping of information that includes a header containing control information and (usually) user data.	Logical grouping of information that includes a header containing control information and (usually) user data. Packets most often are used to refer to network layer units of data. The terms datagram, frame, message, and segment also are used to describe logical information groupings at various layers of the OSI reference model and in various technology circles.	ITA/Jan 2001	
payload	Portion of a cell, frame, or packet that contains upper-layer information (data).	Portion of a cell, frame, or packet that contains upper-layer information (data).	ITA/Jan 2001	
PBR	Policy Based Routing	Policy-based routing. Routing scheme that forwards packets to specific interfaces based on user-configured policies. Such policies might specify that traffic sent from a particular network should be forwarded out one interface, while all other traffic should be forwarded out another interface.	The ABCs of Cisco IOS Software Acronyms	

Acronym or Term	Short Definition for Mouseover	Expansion and Definition of Acronym or Term	Source for Definition	Additional Explanation/ Description for this Course Context
PDLM	Packet Description Language Module	An external Packet Description Language Module (PDLM) can be loaded at run time to extend the NBAR list of recognized protocols. PDLMs can also be used to enhance an existing protocol recognition capability. PDLMs allow NBAR to recognize new protocols without requiring a new Cisco IOS image or a router reload.	Cisco.com http://www.cisco.com/en/US/products/sw/iosswrel/ps1835/products_configuration_guide_chapter09186a00800c75d1.html	
PE router	provider edge router	Provider edge router. A router that is part of a service provider's network and that is connected to a customer edge (CE) router. The PE router function is a combination of an MLS edge label switch router (LSR) function with some additional functions to support VPNs.	http://www.cisco.com/univercd/cc/td/doc/product/rtrmgmt/vpnsc/mpls/2_1/prov_gd/mplsgl.pdf	
PHB	per-hop behavior	per-hop behavior. A description of the externally observable forwarding treatment applied at a differentiated services-compliant node to a behavior aggregate.	IETF RFC2475	
ping	packet internet groper. ICMP echo message and its reply.	packet internet groper. ICMP echo message and its reply. Often used in IP networks to test the reachability of a network device.	ITA/Jan 2001	
PIR	peak information rate. Maximum rate, in kilobits per second, at which a virtual circuit can transmit.	peak information rate. Maximum rate, in kilobits per second, at which a virtual circuit can transmit.	ITA/Jan 2001	
policing	Traffic policing allows control over the maximum rate of traffic sent or received on an interface.	To supervise the operation, execution, or administration of to prevent or detect and prosecute violations of rules and regulations (in this case, regarding pre-defined traffic bandwidth). Traffic policing allows you to control the maximum rate of traffic sent or received on an interface, and to partition a network into multiple priority levels or class of service (CoS).	Merriam-Webster online Dictionary and http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122cgcr/fqos_c/fqcprt4/qcfpolsh.htm#22120	

Acronym or Term	Short Definition for Mouseover	Expansion and Definition of Acronym or Term	Source for Definition	Additional Explanation/ Description for this Course Context
policy	Any defined rule that determines the use of resources within the network.	Any defined rule that determines the use of resources within the network. A policy can be based on a user, a device, a subnetwork, a network, or an application.	ITA/Jan 2001	
policy map	A policy map is a traffic policy.	A policy map is a traffic policy. The purpose of a traffic policy is to configure 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.	Cisco.com http://www.cisco.com/en/US/products/sw/iosswrel/ps1835/products_configuration_guide_chapter09186a00800bd908.html	
PPP	Point-to-Point Protocol	Point-to-Point Protocol. Successor to SLIP that provides router-to-router and host-to-network connections over synchronous and asynchronous circuits. Whereas SLIP was designed to work with IP, PPP was designed to work with several network layer protocols, such as IP, IPX, and ARA. PPP also has built-in security mechanisms, such as CHAP and PAP. PPP relies on two protocols: LCP and NCP.	ITA/Jan 2001	
PQ	priority queuing	priority queuing. Routing feature in which frames in an output queue are prioritized based on various characteristics, such as packet size and interface type.	ITA/Jan 2001	
PQ/CBWFQ	priority queueing/class-based weighted fair queueing	priority queueing/class-based weighted fair queueing (PQ/CBWFQ). Feature that brings strict priority queueing to CBWFQ. Strict priority queueing 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.	ITA/Jan 2001	

Acronym or Term	Short Definition for Mouseover	Expansion and Definition of Acronym or Term	Source for Definition	Additional Explanation/ Description for this Course Context
predictor	One of the two most commonly used compression algorithms on internetworking devices.	The two most commonly used compression algorithms on internetworking devices are the Stacker compression and the predictor data compression algorithms. The predictor compression algorithm tries to predict the next sequence of characters in a data stream by using an index to look up a sequence in the compression dictionary.	Cisco.com http://www.cisco.com/en/US/products/hw/modules/ps2957/products_module_installation_guide_chapter09186a008007e5da.html#xtocid2	
processing delay	The time it takes to manipulate data in the computer	The time it takes to manipulate data in the computer. The computer is said to be processing no matter what action it is taking upon the data; whether the data is actually being updated in a database or just being displayed on screen.	TechEncyclopedia http://www.techweb.com/encyclopedia/defineterm?term=propagation&x=12&y=9	
propagation delay	The time it takes to transmit a signal from one place to another.	The time it takes to transmit a signal from one place to another. Propagation delay is dependent solely on distance and two thirds the speed of light. Signals going through a wire or fiber generally travel at two thirds the speed of light.	TechEncyclopedia http://www.techweb.com/encyclopedia/defineterm?term=propagation&x=12&y=9	
provisioning	Creation of an active subscriber account, or modification of parameters for an existing subscriber account.	Creation of an active subscriber account, or modification of parameters for an existing subscriber account. Provisioning of a subscriber account includes subscriber account registration and device activation.	ITA/Jan 2001	
QoS	quality of service	Quality of service (QoS). The goal of QoS is to provide better and more predictable network service by providing dedicated bandwidth, controlled jitter and latency, and improved loss characteristics. QoS achieves these goals by providing tools for managing network congestion, shaping network traffic, using expensive wide-area links more efficiently, and setting traffic policies across the network.	The ABCs of Cisco IOS Software Glossary	

Acronym or Term	Short Definition for Mouseover	Expansion and Definition of Acronym or Term	Source for Definition	Additional Explanation/ Description for this Course Context
QPPB	QoS Policy Propagation on BGP	The quality of service (QoS) policy propagation via Border Gateway Protocol (BGP) feature allows you to classify packets based on access lists, BGP community lists, and BGP autonomous system (AS) paths.	Cisco.com http://www.cisco.com/en/US/products/sw/iosswrel/ps1820/products_feature_guide09186a00800f4898.html	
queue	A sequence of messages or jobs held in auxiliary storage awaiting transmission.	A sequence of messages or jobs held in auxiliary storage awaiting transmission or processing or a data structure that consists of a list of records such that records are added at one end and removed from the other.	Merriam-Webster online Dictionary	
queuing delay	Queuing delay is time spent being held in a queue.	Queuing is the process of lining up events in the order you want them processed. Whether it refers to packets in an IP network that search for the most optimal path to their destination, or telephone callers sitting in a "hold queue" waiting to be answered, queuing means the same thing: deciding on priorities through bottle-necked passageways. Queuing delay is time spent being held in a queue.	TechEncyclopedia http://www.techweb.com/encyclopedia/defineterm?term=propagation&x=12&y=9 and http://www.cisco.com/en/US/tech/tk652/tk698/technologies_white_paper09186a00800a8993.shtml#codeprocessdelay	
RED	random early detection	Random early detection. This class of algorithms is designed to avoid congestion in internetworks before it becomes a problem. RED works by monitoring traffic load at points in the network and stochastically discarding packets if the congestion begins to increase. The result of the drop is that the source detects the dropped traffic and slows its transmission. RED is designed to work primarily with TCP in IP internetwork environments.	The ABCs of Cisco IOS Software Glossary	

Acronym or Term	Short Definition for Mouseover	Expansion and Definition of Acronym or Term	Source for Definition	Additional Explanation/ Description for this Course Context
RSVP	Resource Reservation Protocol. Protocol that supports the reservation of resources across an IP network.	Resource Reservation Protocol. Protocol that supports the reservation of resources across an IP network. Applications running on IP end systems can use RSVP to indicate to other nodes the nature (bandwidth, jitter, maximum burst, and so on) of the packet streams they want to receive. RSVP depends on IPv6. Also known as Resource Reservation Setup Protocol.	ITA/Jan 2001	
RTP	Real-Time Transport Protocol. RTP is designed to provide end-to-end network transport functions for applications transmitting real-time data.	Real-Time Transport Protocol. Commonly used with IP networks. RTP is designed to provide end-to-end network transport functions for applications transmitting real-time data, such as audio, video, or simulation data, over multicast or unicast network services. RTP provides such services as payload type identification, sequence numbering, timestamping, and delivery monitoring to real-time applications.	ITA/Jan 2001	
scheduling	Method used to schedule data for transmission.	A (scheduling algorithm is a) method used to schedule jobs for execution (or, in this case, data for transmission). Priority, length of time in the job queue and available resources are examples of criteria used.	TechEncyclopedia http://www.techweb.com/encyclopedia/defineterm?term=scheduling	
serialization delay	The time it takes in a computer system to transmit the bits of a byte sequentially over a single wire.	The time it takes in a computer system to transmit the bits of a byte sequentially over a single wire.	Merriam-Webster online Dictionary	
service class	Collection of service types required for a specific service offered.	Collection of service types required for a specific service offered. Each service class includes the attributes and values that define the type or quality of service associated with a given class. For example, data connectivity is a service class you might define that includes the service type data-bandwidth.	ITA/Jan 2001	
shaping	Adapt the traffic so that it fits a previously defined bandwidth.	To adapt in shape so as to fit neatly and closely (in this case to adapt the traffic so that it fits a previously defined bandwidth)	Merriam-Webster online Dictionary	
SLA	service-level agreement	service-level agreement. A service contract between a customer and a service provider that specifies the forwarding service a customer should receive.	IETF RFC 2475	

Acronym or Term	Short Definition for Mouseover	Expansion and Definition of Acronym or Term	Source for Definition	Additional Explanation/ Description for this Course Context
stacker	One of the two most commonly used compression algorithms on internetworking devices.	The two most commonly used compression algorithms on internetworking devices are the stacker compression and the Predictor data compression algorithms. The stacker algorithm uses an encoded dictionary that replaces a continuous stream of characters with codes. The symbols represented by the codes are stored in memory in a dictionary-style list.	http://www.cisco.com/en/US/products/hw/modules/ps2957/products_module_installation_guide_chapter09186a008007e5da.html#xtocid2	
starvation	To suffer from a lack of service.	To perish from lack of food or to suffer extreme hunger (in this case from a lack of service)	Merriam-Webster online Dictionary http://www.m-w.com/cgi-bin/dictionary	
tail-drop	Method for avoiding congestion in which packets are dropped until the congestion is eliminated and the queue is no longer full.	Method for avoiding congestion. Tail drop treats all traffic equally and does not differentiate between classes of service. Queues fill during periods of congestion. When the output queue is full and tail drop is in effect, packets are dropped until the congestion is eliminated and the queue is no longer full.	Cisco.com http://www.cisco.com/en/US/products/sw/iosswrel/ps1828/products_configuration_guide_chapter09186a00800ca59c.html	
Tc	Time interval	Time interval. Also called the measurement interval, it specifies the time quantum in seconds per burst.	http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122cgcr/fqos_c/fqcprt4/qcfpolsh.htm	
TCP	Transmission Control Protocol. Transport layer protocol that provides reliable full-duplex data transmission.	Transmission Control Protocol. Connection-oriented transport layer protocol that provides reliable full-duplex data transmission. TCP is part of the TCP/IP protocol stack.	ITA/Jan 2001	

Acronym or Term	Short Definition for Mouseover	Expansion and Definition of Acronym or Term	Source for Definition	Additional Explanation/ Description for this Course Context
token bucket	A token bucket is used to manage a device that regulates the rate of a flow.	A formal definition of a rate of transfer. A token bucket has three components: a burst size, a mean rate, and a time interval (Tc). A token bucket is used to manage a device that regulates the rate of a flow.	ITA/Jan 2001	
ToS	type of service	A field in an IP packet (IP datagram) that is used for quality of service (QoS).	TechEncyclopedia http://www.techweb.com/encyclopedia/defineterm?term=tos&x=30&y=12	
trust boundary	The trust boundary is established by the access device that either discards or trusts a marker in a frame or packet.	When classifying traffic types in an enterprise network, a trust boundary must be established. The boundary is established by the access device, which either classifies traffic that it allows into the network itself or trusts classification that has already been applied by an end station, such as an IP phone.	Cisco.com http://www.cisco.com/en/US/about/ac123/ac114/ac173/ac205/about_cisco_packet_feature09186a00801513.html	
tunnel	Secure communication path between two peers, such as two routers.	Secure communication path between two peers, such as two routers.	ITA/Jan 2001	
UDP	User Datagram Protocol. Connectionless transport layer protocol in the TCP/IP protocol stack.	User Datagram Protocol. Connectionless transport layer protocol in the TCP/IP protocol stack. UDP is a simple protocol that exchanges datagrams without acknowledgments or guaranteed delivery, requiring that error processing and retransmission be handled by other protocols. UDP is defined in RFC 768.	ITA/Jan 2001	
VC	virtual circuit. Logical circuit created to ensure reliable communication between two network devices.	virtual circuit. Logical circuit created to ensure reliable communication between two network devices. A virtual circuit is defined by a VPI/VCI pair, and can be either permanent (PVC) or switched (SVC). Virtual circuits are used in Frame Relay and X.25. In ATM, a virtual circuit is called a virtual channel. Sometimes abbreviated VC.	ITA/Jan 2001	
violating traffic	Traffic not following limits set for traffic bandwidth.	To fail to show proper respect for (in this case, not following limits set for traffic bandwidth)	Merriam-Webster online Dictionary	

Acronym or Term	Short Definition for Mouseover	Expansion and Definition of Acronym or Term	Source for Definition	Additional Explanation/ Description for this Course Context
VIP	Versatile Interface Processor. Interface card used in Cisco 7000 and Cisco 7500 series routers.	Versatile Interface Processor. Interface card used in Cisco 7000 and Cisco 7500 series routers. The VIP provides multilayer switching and runs Cisco IOS. The most recent version of the VIP is VIP2.	ITA/Jan 2001	
VLAN	virtual LAN. Group of devices on one or more LANs that are configured so that they can communicate as if they were attached to the same wire.	virtual LAN. Group of devices on one or more LANs that are configured (using management software) so that they can communicate as if they were attached to the same wire, when in fact they are located on a number of different LAN segments. Because VLANs are based on logical instead of physical connections, they are extremely flexible.	ITA/Jan 2001	
VoIP	Voice over IP. The capability to carry normal telephony-style voice over an IP-based internet with POTS-like functionality, reliability, and voice quality.	Voice over IP. The capability to carry normal telephony-style voice over an IP-based internet with POTS-like functionality, reliability, and voice quality. VoIP enables a router to carry voice traffic (for example, telephone calls and faxes) over an IP network. In VoIP, the DSP segments the voice signal into frames, which then are coupled in groups of two and stored in voice packets. These voice packets are transported using IP in compliance with ITU-T specification H.323.	ITA/Jan 2001	
VPN	Virtual Private Network	Virtual Private Network. Enables IP traffic to travel securely over a public TCP/IP network by encrypting all traffic from one network to another. A VPN uses "tunneling" to encrypt all information at the IP level.	ITA/Jan 2001	
WAN	wide-area network	wide-area network. Data communications network that serves users across a broad geographic area and often uses transmission devices provided by common carriers. Frame Relay, SMDS, and X.25 are examples of WANs.	ITA/Jan 2001	
WFQ	weighted fair queuing	weighted fair queuing. Congestion management algorithm that identifies conversations (in the form of traffic streams), separates packets that belong to each conversation, and ensures that capacity is shared fairly between these individual conversations. WFQ is an automatic way of stabilizing network behavior during congestion and results in increased performance and reduced retransmission.	ITA/Jan 2001	

Acronym or Term	Short Definition for Mouseover	Expansion and Definition of Acronym or Term	Source for Definition	Additional Explanation/ Description for this Course Context
WRED	weighted random early detection	weighted random early detection. Queueing method that ensures that high-precedence traffic has lower loss rates than other traffic during times of congestion.	ITA/Jan 2001	
WRR	weighted round robin	weighted round robin. Packets are assigned a class (real-time, file transfer, etc.) and placed into the queue for that class of service. Packets are accessed round-robin style, but classes can be given priorities based on “weights” assigned to the queues. For example, four packets from a high-priority class might be serviced, followed by two from a middle-priority class and then one from a low-priority class.	TechEncyclopedia http://www.techweb.com/encyclopedia/defineterm?term=trafficengineering methods	

Lab Guide

Overview

Use the exercises here to complete the lab exercises for this course. The solutions information is found in the Lab Exercise Answer Key.

Outline

This Lab Guide includes these exercises:

- Lab Exercise 2-1: QoS Lab Setup and Initialization
- Lab Exercise 2-2: Baseline QoS Measurement
- Lab Exercise 3-1: Configuring QoS with AutoQoS
- Lab Exercise 4-1: Classification and Marking Using MQC
- Lab Exercise 4-2: Classification Using NBAR
- Lab Exercise 4-3: Configuring QoS Pre-Classify
- Lab Exercise 4-4: LAN-Based Packet Classification and Marking
- Lab Exercise 5-1: Configuring Basic Queuing
- Lab Exercise 5-2: Configuring LLQ
- Lab Exercise 5-3: Queuing on Catalyst Switches
- Lab Exercise 6-1: Configuring DSCP-Based WRED
- Lab Exercise 7-1: Configuring Class-Based Policing
- Lab Exercise 7-2: Configuring Class-Based Shaping
- Lab Exercise 8-1: Configuring Class-Based Header Compression
- Lab Exercise 8-2: Configuring LFI

Lab Exercise 2-1: QoS Lab Setup and Initialization

Complete this lab initialization procedure to prepare your student workgroup pod prior to completing any laboratory practice exercises.

Exercise Objective

In this exercise, you will prepare your student workgroup for use with the laboratory practice exercises that accompany the *Implementing Cisco Quality of Service* (QOS) v2.0 course. After completing this exercise, you will be able to meet these objectives:

- Configure your workgroup routers for basic network connectivity
- Configure your workgroup switch for basic network connectivity
- Verify network connectivity using the Cisco IOS tools: ping and trace

Visual Objective

The following figures illustrate what you will accomplish in this exercise.

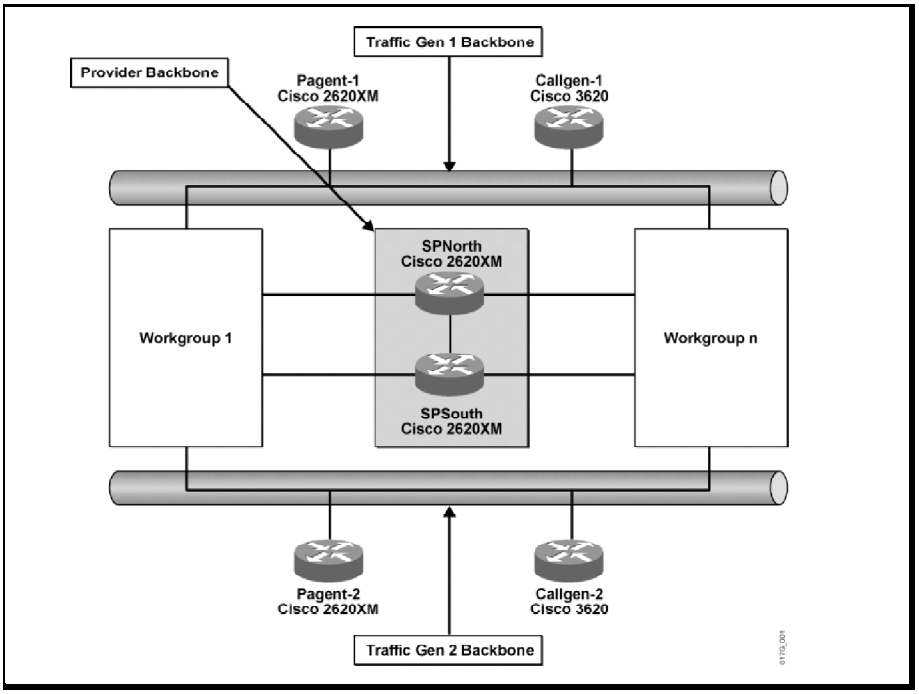

Figure 1: Logical Lab Backbone Topology

As shown in Figure 1, the lab topology for the course is split into a number of workgroups and three separate backbones.

Each workgroup is designated to service two students and has been designed to interface with two traffic generation backbones named, “Traffic Gen 1” and “Traffic Gen 2” and a shared provider backbone named “Provider.”

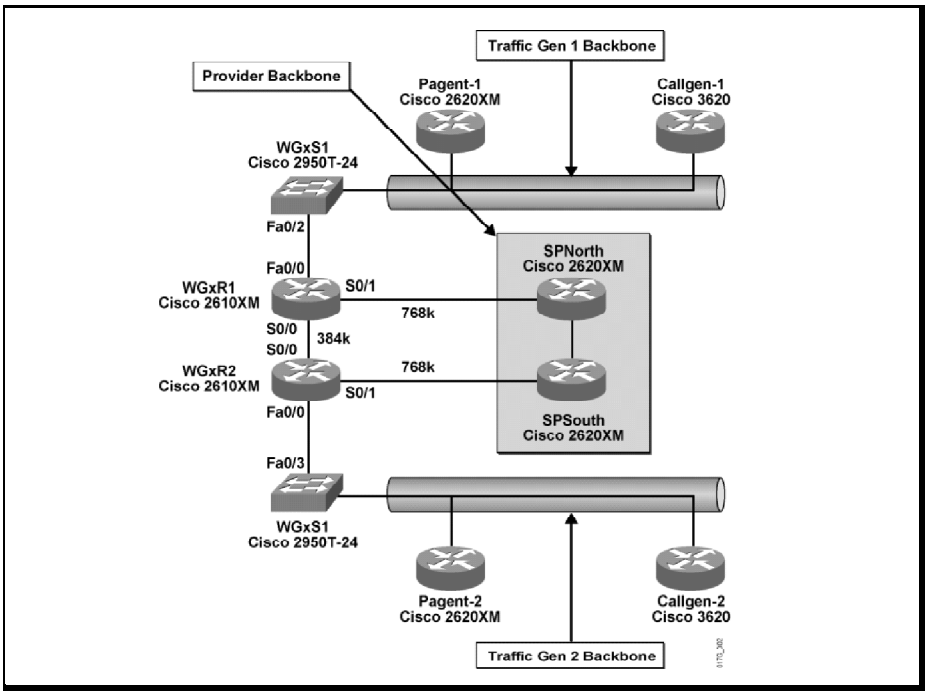

Figure 2: Physical Lab Workgroup Topology

Figure 2 shows the physical topology of a single workgroup and its connectivity into the three lab backbones. Each workgroup consists of two user-controlled Cisco 2610XM routers and one user-controlled Cisco 2950T-24 workgroup switch. Each student workgroup connects to each backbone as shown in Figure 2.

Although depicted as two different workgroup switches in Figure 2, each student workgroup consists of a single workgroup switch configured to support two different virtual LANs. In figure two, the single workgroup switch has been depicted as two different switches to simplify the diagram only. Notice that the same name has been used to identify the Cisco 2950T indicating it is in fact the same device.

Traffic for each pod can bypass the high-speed service provider backbone (using the slow 384-kbps link) or travel via the high-speed provider backbone itself (using the fast 768-kbps serial link). Traffic flow through both the slow and fast serial links will be tested in the QOS v2.0 labs.

In Figure 3, the logical topology configuration of each workgroup and the devices contained within each of the three laboratory backbones is shown. In the Provider Backbone in Figure 3, each of the backbone routers (SPNorth and SPSouth) contains a serial connection to each workgroup.

The QoS lab uses the two routers called Pagent-1 and Pagent-2 to generate traffic from different applications including SQL, Napster, FTP, Citrix, HTTP, Outlook, and Kazaa. The two routers called Callgen-1 and Callgen-2 are used in the lab to generate (G.711) VoIP traffic.

Note:	The SPNorth router, the SPSouth router, the core switch, and the traffic generation routers (Pagent-1, Pagent-2, Callgen-1, and Callgen-2) are preconfigured and managed by the instructor.

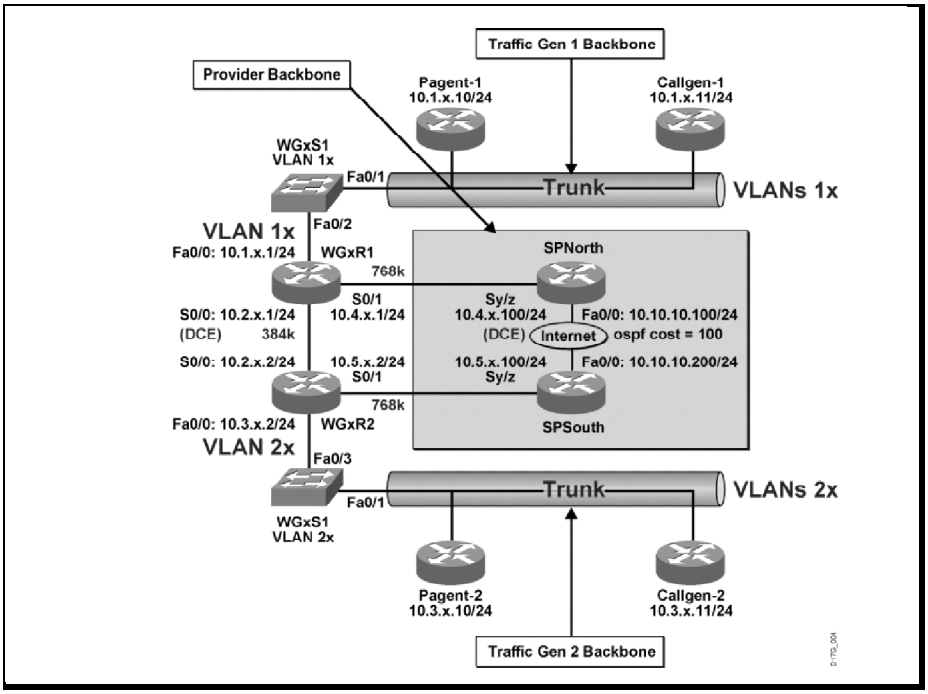

Figure 3: Logical Lab Topology

Each Pagent and Callgen router is set up with connections to 8 VLANs (one for each lab workgroup) as follows:

- Pagent-1 and Callgen-1 are configured with VLANs 11 to 18.
- Pagent-2 and Callgen-2 are configured with VLANs 21 to 28.

Traffic flow to and from the Pagent and Callgen lab routers is designed to traverse the network through each workgroup as follows:

- Pagent-1/Callgen-1 (VLAN 11) sends traffic to Pagent-2/Callgen-2 (VLAN 21) via pod 1.
- Pagent-1/Callgen-1 (VLAN 12) sends traffic to Pagent-2/Callgen-2 (VLAN 22) via pod 2.
- Pagent-1/Callgen-1 (VLAN 13) sends traffic to Pagent-2/Callgen-2 (VLAN 23) via pod 3.
- Pagent-1/Callgen-1 (VLAN 14) sends traffic to Pagent-2/Callgen-2 (VLAN 24) via pod 4.
- Pagent-1/Callgen-1 (VLAN 15) sends traffic to Pagent-2/Callgen-2 (VLAN 25) via pod 5.
- Pagent-1/Callgen-1 (VLAN 16) sends traffic to Pagent-2/Callgen-2 (VLAN 26) via pod 6.
- Pagent-1/Callgen-1 (VLAN 17) sends traffic to Pagent-2/Callgen-2 (VLAN 27) via pod 7.
- Pagent-1/Callgen-1 (VLAN 18) sends traffic to Pagent-2/Callgen-2 (VLAN 28) via pod 8.

The logical configuration of each of these VLANs is as follows:

Lab VLAN Logical Address Assignments

Workgroup Pod	VLANs	Assigned IP Subnets
1	11 and 21	10.1.1.0/24 (VLAN 11) and 10.3.1.0/24 (VLAN 21)
2	12 and 22	10.1.2.0/24 (VLAN 12) and 10.3.2.0/24 (VLAN 22)
3	13 and 23	10.1.3.0/24 (VLAN 13) and 10.3.3.0/24 (VLAN 23)
4	14 and 24	10.1.4.0/24 (VLAN 14) and 10.3.4.0/24 (VLAN 24)
5	15 and 25	10.1.5.0/24 (VLAN 15) and 10.3.5.0/24 (VLAN 25)
6	16 and 26	10.1.6.0/24 (VLAN 16) and 10.3.6.0/24 (VLAN 26)
7	17 and 27	10.1.7.0/24 (VLAN 17) and 10.3.7.0/24 (VLAN 27)
8	18 and 28	10.1.8.0/24 (VLAN 18) and 10.3.8.0/24 (VLAN 28)

Required Resources

These are the resources and equipment required to complete this exercise:

- Lab topology configured for QOS
- Student workgroup consisting of two user-controlled Cisco 2610XM routers and one user-controlled Cisco 2950T-24 workgroup switch
- Classroom reference materials as follows:
  - QOS Student Guide
  - QOS Lab Guide
- Student pod workstation with Telnet or console access to workstation pod devices

Command List

The commands used in this exercise are described in the table here.

QoS Lab Setup and Initialization Lab Router Commands

Command	Description
hostname *name*	To specify or modify the host name
enable secret *password*	Password for users to enter enable mode
interface *interface-id*	Enter interface configuration mode and the physical interface identification
[no] ip address ip-address mask	To set a primary or secondary IP address for an interface
clock rate *bps*	To configure the clock rate for the hardware connections on serial interfaces
bandwidth *kbps*	To set and communicate to higher-level protocols the current bandwidth value for an interface
encapsulation *encapsulation-type*	To set the encapsulation method used by the interface
show ip interface [brief] [*type***] [***number***]**	To list a summary of an interface IP information and status
router ospf *process-id*	To configure an OSPF routing process
network *ip-address wildcard-mask* **area** *area-id*	To define the interfaces on which OSPF runs and to define the area ID for those interfaces
show ip ospf neighbor	To display OSPF-neighbor information on a per-interface basis
shutdown	To disable an interface
copy running-config startup-config	Save your entries in the configuration file

QoS Lab Setup and Initialization Lab Switch Commands

Command	Description
hostname *name*	To specify or modify the host name
enable secret *password*	Password for users to enter enable mode
interface ***interface-id***	Enter interface configuration mode and the physical interface identification
show ip interface [brief] [*type***] [***number***]**	To list a summary of an interface IP information and status
show interfaces [*interface-id***]**	Displays the administrative and operational status of all interfaces or a specified interface
copy running-config startup-config	Save your entries in the configuration file

Job Aids

These job aids are available to help you complete the laboratory exercise:

- Your assigned workgroup pod number provided by the instructor

Exercise Procedure

Complete these steps:

Step 1 Configure the hostname and passwords on both of the workgroup routers in your assigned workgroup pod as shown in the following table (where *x* is your assigned workgroup pod number):

Hostname	Enable Secret Password	VTY Login Password
WGxR1	cisco	cisco
WGxR2	cisco	cisco

Step 2 Configure the IP address on the S0/0 and S0/1 and Fa0/0 interfaces of the workgroup routers in your assigned workgroup pod as shown in the following table:

Interface	WGxR1	WGxR2
Fa0/0	10.1.x.1/24	10.3.x.2/24
S0/0	10.2.x.1/24	10.2.x.2/24
S0/1	10.4.x.1/24	10.5.x.2/24

Step 3 Configure the clock rate on the S0/0 serial interface of your WGxR1 router to 384 kbps.

Interface	WGxR1 S0/0 clock rate
S0/0	384000 bps

Note: In the service provider backbone, each of the backbone routers (SPNorth, SPSouth) contains a serial connection to each workgroup router. The following table lists the IP addressing requirements of these connections. Both service provider routers are the DCE with the clock rate configured by the instructor as 768 kbps.

Sy/x	SPNorth IP Address	SPSouth IP Address
S0/0 – to pod 1	10.4.1.100	10.5.1.100
S0/1 – to pod 2	10.4.2.100	10.5.2.100
S0/2 – to pod 3	10.4.3.100	10.5.3.100
S0/3 – to pod 4	10.4.4.100	10.5.4.100
S1/0 – to pod 5	10.4.5.100	10.5.5.100
S1/1 – to pod 6	10.4.6.100	10.5.6.100
S1/2 – to pod 7	10.4.7.100	10.5.7.100
S1/3 – to pod 8	10.4.8.100	10.5.8.100

Step 4 Configure the S0/0 and S0/1 serial interfaces of your workgroup routers for PPP encapsulation and set the bandwidth to match the clock rate configured in Step 3 of this lab exercise.

Step 5 Administratively enable the S0/0, S0/1, and Fa0/0 interfaces on both of your workgroup routers and verify that they are all in the up state (administratively up, line protocol up).

If the Fa0/0 interface is down, login to your workgroup switch to ensure the switch port is also administratively enabled.

```
WGxR1#show ip interface brief

Interface        IP-Address OK?  Method Status Protocol
FastEthernet0/0  10.1.x.1   YES  NVRAM  up     up
Serial0/0        10.2.x.1   YES  NVRAM  up     up
Serial0/1        10.4.x.1   YES  NVRAM  up     up
Virtual-Access1  unassigned YES  unset  up     up
```

```
WGxR2#show ip interface brief

Interface        IP-Address OK?  Method Status Protocol
FastEthernet0/0  10.3.x.2   YES  NVRAM  up     up
Serial0/0        10.2.x.2   YES  NVRAM  up     up
Serial0/1        10.5.x.2   YES  NVRAM  up     up
Virtual-Access1  unassigned YES  unset  up     up
```


Step 6 Configure an OSPF routing process on your workgroup routers and place the S0/0, S0/1, and Fa0/0 interfaces into OSPF area 0.

```
router ospf 1
 network 10.0.0.0 0.255.255.255 area 0
```

Step 7 Verify that both OSPF neighbors of your workgroup routers are in the FULL state. Each of your workgroup routers should have a FULL neighbor relationship to the service provider router and to the other workgroup router in your pod.

```
WGxR1#show ip ospf neighbor

Neighbor ID   Pri       State   Dead Time Address
Interface
10.10.10.100  0         FULL/-  00:00:37  10.4.x.100
Serialy/z
10.5.x.2      0         FULL/-  00:00:36  10.2.x.2
Serial0/0
```

```
WGxR2#show ip ospf neighbor

Neighbor ID   Pri   State   Dead Time Address     Interface
10.10.10.200  0         FULL/-  00:00:31  10.5.x.100
Serialy/z
10.4.x.1      0         FULL/-  00:00:33  10.2.x.1
Serial0/0
```

Step 8 Verify that the serial ports on both of your workgroup routers (WGxR1 and WGxR2) have their queuing strategy set to weighted fair queuing (WFQ).

Step 9 Configure the hostname and password on the workgroup switch in your assigned workgroup pod as shown in the following table:

Hostname	Enable Secret Password
WGxS1	cisco

Step 10 Configure the VLANs (1 x and 2 x) on your workgroup switch by adding VLANs 1x and 2x from within vlan database mode.

```
vlan database
vtp domain qos
vtp transparent
!
vlan 1x name vlan1x
vlan 2x name vlan2x
!
exit
```

Step 11 Configure 802.1Q trunking and access ports on the workgroup switch by configuring the Fa0/1, Fa0/2, and Fa0/3 interfaces of your workgroup switch as follows:

- Fa0/1 is an 802.1Q trunk connected to the core switch. Only VLAN 1x and 2x should be allowed on the trunk.
- Fa0/2 should be an access port in VLAN 1x connected to the WGxR1 router.
- Fa0/3 should be an access port in VLAN 2x connected to the WGxR2 router.

```
interface FastEthernet0/1
 description - to core sw
 switchport trunk allowed vlan 1x,2x
 switchport mode trunk
 no ip address
!
interface FastEthernet0/2
 description - to WGxR1
 switchport access vlan 1x
 switchport mode access
 no ip address
!
interface FastEthernet0/3
 description - to WGxR2
 switchport access vlan 2x
 switchport mode access
 no ip address
```

Step 12 Verify that the Fa0/1, Fa0/2, and Fa0/3 interfaces on the workgroup switch are all up. Administratively enable any interfaces in the shutdown state.

```
WGxS1#sh ip int brief

Interface       IP-Address      OK?   Method    Status
          Protocol
Vlan1           unassigned  YES manual    administratively
down down
FastEthernet0/1 unassigned  YES unset     up
up
FastEthernet0/2 unassigned  YES unset     up
up
FastEthernet0/3 unassigned  YES unset     up
up
FastEthernet0/4 unassigned  YES unset     down
down

[rest omitted]
```

Note: In the lab, there is no requirement to ping to or from the workgroup switch. The workgroup switch will not need an IP address configured on Interface VLAN1 and will not need an IP default-gateway configuration.

Step 13 From your WGxS1 switch, use the show interface fa0/x switchport command to verify that the Fa0/1 interface 802.1Q trunking is on and only allow VLANs 1x and 2x on the trunk.

```
WGxS1#sh int fa 0/1 switchport
Name: Fa0/1
Switchport: Enabled
Administrative Mode: trunk
Operational Mode: trunk
Administrative Trunking Encapsulation: dot1q
Operational Trunking Encapsulation: dot1q
Negotiation of Trunking: On
Access Mode VLAN: 1 (default)
Trunking Native Mode VLAN: 1 (default)
Administrative private-vlan host-association: none
Administrative private-vlan mapping: none
Operational private-vlan: none
Trunking VLANs Enabled: 1x,2x
Pruning VLANs Enabled: 2-1001

Protected: false

Voice VLAN: none (Inactive)
Appliance trust: none
```

Step 14 Verify that the Fa0/2 interface is in VLAN1x.

```
WGxS1#sh int fa 0/2 switchport
Name: Fa0/2
Switchport: Enabled
Administrative Mode: static access
Operational Mode: static access
Administrative Trunking Encapsulation: dot1q
Operational Trunking Encapsulation: native
Negotiation of Trunking: Off
Access Mode VLAN: 1x (VLAN001x)
Trunking Native Mode VLAN: 1 (default)
Administrative private-vlan host-association: none
Administrative private-vlan mapping: none
Operational private-vlan: none
Trunking VLANs Enabled: ALL
Pruning VLANs Enabled: 2-1001

Protected: false

Voice VLAN: none (Inactive)
Appliance trust: none
```

Step 15 Verify that the Fa0/3 interface is in VLAN2x.

```
WGxS1#sh int fa 0/3 switchport
Name: Fa0/3
Switchport: Enabled
Administrative Mode: static access
Operational Mode: static access
Administrative Trunking Encapsulation: dot1q
Operational Trunking Encapsulation: native
Negotiation of Trunking: Off
Access Mode VLAN: 2x (VLAN002x)
```

```
Trunking Native Mode VLAN: 1 (default)
Administrative private-vlan host-association: none
Administrative private-vlan mapping: none
Operational private-vlan: none
Trunking VLANs Enabled: ALL
Pruning VLANs Enabled: 2-1001

Protected: false

Voice VLAN: none (Inactive)
Appliance trust: none
```

Step 16 From the WGxR1 router, perform the following pings to confirm connectivity and routing protocol operation:

- Ping the SPNorth router (10.4.x.100)

```
WGxR1#ping 10.4.x.100

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.4.x.100, timeout is 2
seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max =
4/4/4 ms
```

- Ping the SPSouth router internet connection (10.10.10.200)

```
WGxR1#ping 10.10.10.200

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.10.10.200, timeout is 2
seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max =
4/4/8 ms
```

- Ping the WGxR2 router (10.2.x.2)

```
WGxR1#ping 10.2.x.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.2.x.2, timeout is 2
seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max =
16/16/16 ms
```

- Ping the Pagent-1 (10.1.x.10) and Callgen-1 (10.1.x.11) routers.

```
WGxR1#ping 10.1.x.10

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.x.10, timeout is 2
seconds:
!!!!!
```



```
Success rate is 100 percent (5/5), round-trip min/avg/max =
40/47/56 ms

WGxR1#ping 10.1.x.11

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.x.11, timeout is 2
seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max =
1/2/4 ms
```

Step 17 From the WGxR2 router, perform the following pings to confirm connectivity and routing protocol operation:

- Ping the SPSouth router (10.5.x.100)

```
WGxR2#ping 10.5.x.100

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.5.x.100, timeout is 2
seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max =
4/4/4 ms
```

- Ping the SPNorth router internet connection (10.10.10.100)

```
WGxR2#ping 10.10.10.100

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.10.10.100, timeout is 2
seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max =
4/5/8 ms
Ping the Pagent-2 (10.3.x.10) and Callgen-2 (10.3.x.11)
routers.
```

- Ping the Pagent-2 (10.3.x.10) and Callgen-2 (10.3.x.11) routers.

```
WGxR2#ping 10.3.x.10

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.3.x.10, timeout is 2
seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max =
40/47/56 ms

WGxR2#ping 10.3.x.11

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.3.x.11, timeout is 2
seconds:
!!!!!
```

```
Success rate is 100 percent (5/5), round-trip min/avg/max =
1/2/4 ms
```

Step 18 On the WGxR1 and WGxR2 routers, administratively disable (shut) the serial 0/0 interface.

Step 19 From the WGxR1 router, Telnet to the Pagent-1 router (10.1.x.10) and perform a trace to the Pagent-2 router (10.3.x.10) to confirm the path from Pagent-1 to Pagent-2 flows through your pod (via the SPNorth and SPSouth router). Use the diagram in Figure 3 to verify your trace.

```
WGxR1#telnet 10.1.x.10
Trying 10.1.x.10 ... Open

User Access Verification
Username: super
Password: bowl

pagent-1>trace 10.3.x.10

Type escape sequence to abort.
Tracing the route to 10.3.x.10

  1 10.1.x.1 68 msec 64 msec 60 msec
  2 10.4.x.100 44 msec 56 msec 56 msec
  3 10.10.10.200 48 msec 52 msec 48 msec
  4 10.5.x.2 44 msec 52 msec 52 msec
  5 10.3.x.10 32 msec *  44 msec

pagent-1>exit
[Connection to 10.1.1.10 closed by foreign host]
WGxR1#
```

Step 20 From the WGxR1 router, Telnet to the Callgen-1 (10.1.x.11) router and perform a trace to the Callgen-2 router (10.3.x.11) to confirm the path from Callgen-1 to Callgen-2 flows through your pod (via the SPNorth and SPSouth router). Use the diagram in Figure 3 to verify your trace.

```
WGxR1>telnet 10.1.x.11
Trying 10.1.x.11 ... Open

User Access Verification
Username: super
Password: bowl

callgen-1>trace 10.3.x.11

Type escape sequence to abort.
Tracing the route to 10.3.x.11

  1 10.1.x.1 68 msec 64 msec 60 msec
  2 10.4.x.100 44 msec 56 msec 56 msec
  3 10.10.10.200 48 msec 52 msec 48 msec
  4 10.5.x.2 44 msec 52 msec 52 msec
  5 10.3.x.11 32 msec *  44 msec
```



```
callgen-1>exit
[Connection to 10.1.x.10 closed by foreign host]
WGxR1#
```

Step 21 On the WGxR1 and WGxR2 routers, administratively enable (no shut) the serial 0/0 interface.

Step 22 Telnet to the Pagent-1 router (10.1.x.10) and perform a trace to the Pagent-2 router (10.3.x.10) to confirm the path from Pagent-1 to Pagent-2 now flows through your pod and via the slow 384-kbps serial connection between your WGxR1 and WGxR2 router. Use the diagram in Figure 3 to verify your trace.

```
WGxR1#telnet 10.1.x.10

Trying 10.1.x.10 ... Open

User Access Verification
Username: super
Password: bowl

pagent-1>trace 10.3.x.10

Type escape sequence to abort.
Tracing the route to 10.3.x.10

  1 10.1.x.1 56 msec 52 msec 60 msec
  2 10.2.x.2 116 msec 252 msec 48 msec
  3 10.3.x.10 128 msec *  104 msec

pagent-1>exit
[Connection to 10.1.x.10 closed by foreign host]
WGxR1#
```

Step 23 Save your running configurations of the workgroup routers and the workgroup switch to the startup-config in NVRAM.

Step 24 Notify your instructor when you have completed this initial setup lab.

Exercise Verification

You have completed this exercise when you attain these results:

- Pings from the WGxR1 router to the SPNorth, WGxR2, Pagent-1, and Callgen-1 routers are successful
- Pings from the WGxR2 router to the SPSouth, Pagent-2, and Callgen-2 routers are successful
- A trace from Pagent-1 router to Pagent-2 router flows through your pod (via the SPNorth and SPSouth routers) with the S0/0 interface in the shutdown state
- A trace from Callgen-1 router to Callgen-2 router flows through your pod (via the SPNorth and SPSouth routers) with the S0/0 interface in the shutdown state
- A trace from the Pagent-1 router to the Pagent-2 router flows through your pod 384-kbps serial link with all WGxR1 and WGxR2 serial interfaces administratively enabled

Lab Exercise 2-2: Baseline QoS Measurement

Complete this lab exercise to practice what you learned in the related lesson.

Exercise Objective

In this exercise, you will create a baseline measurement of network traffic for use in evaluating the effectiveness of applied QoS mechanisms. After completing this exercise, you will be able to meet these objectives:

- Clear interface counters on Cisco routers and switches
- Identify interface statistics, which are meaningful in traffic baselines
- Use Cisco IOS monitoring commands and network connectivity tools (ping) to gather network response time data

Visual Objective

The figure illustrates what you will accomplish in this exercise.

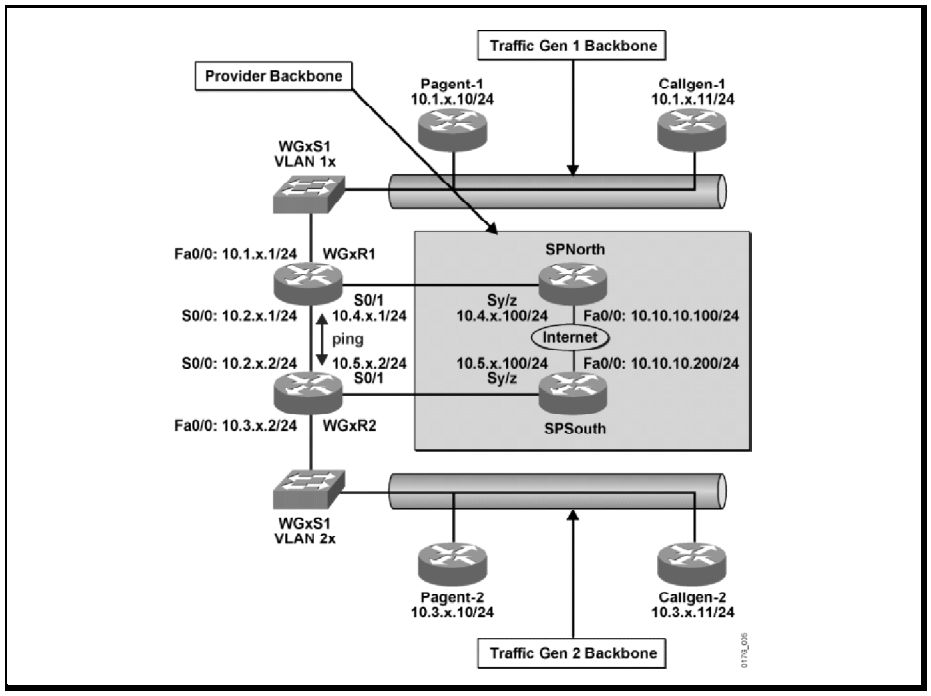

In this lab, when the connectivity is properly established for your pod, you will record traffic statistics without any QoS configuration on your workgroup routers or workgroup switch. This record of the workgroup traffic statistics will form a rough baseline QoS measurement for your pod.

For this lab, no special tools such as QPM, will be used to monitor QoS statistics. Instead, Cisco IOS show commands and extended pings are used to form a rough baseline measurement.

Company Background

E-Commerce University is one of the most respected private universities in California based on its progressive educational offerings, which include a Master of Science degree in E-Commerce Administrations and Implementations. E-Commerce University has a northern and southern campus. The northern campus is located in Seattle, Washington, and the southern campus is located in Santa Monica, California. Both the northern and southern campuses have populations of 500 students and a faculty staff of 50 professors and administrators.

Customer Situation

The E-Commerce University network currently has limited bandwidth capacity on their 384-kbps lease line PPP WAN link that connects the northern and southern campus and they do not envision being able to increase bandwidth in the near future. Both campuses also have a 786-kbps Internet connection. The preferred traffic path between the E-Commerce University campuses is the 384-kbps link as the 768-kbps link connects to the Internet and will have a high path cost, even though it is a directly connected link with a higher bandwidth of 768 kbps.

The University has recently implemented the following three new applications:

- An IP telephony system between the northern and southern campus
- An Oracle (SQL) student administrations database system
- Wireless Internet access for the students and faculties

Some of the other key applications currently running on the E-Commerce University network that the University IT staff is aware of include:

- Faculty remote access (Citrix)
- Outlook e-mail (Exchange)
- Web server for accessing University Information (HTTP)
- Online courseware transfer between the northern and southern campus (FTP)

Because the deployment of these applications, the E-Commerce University has been encountering increasingly serious problems with their network.

- Users of the Oracle (SQL) student administrations database system have been complaining of unacceptable response times. Their subsecond response time has now stretched to multiple seconds in many cases and up to a minute in some cases.
- User of the new IP telephony devices are the most upset. The quality of their calls is very poor and their calls often just drop.

Customer Requirements

At this point, the E-Commerce University is most concerned about the low VoIP voice quality and has called upon you (the new network engineer they hired in the last month) to perform a baseline measurement of the VoIP traffic via the low speed 384-kbps lease line connection between the northern and southern campus.

Required Resources

These are the resources and equipment required to complete this exercise:

- Lab topology configured for QOS
- Student workgroup consisting of two user-controlled Cisco 2610XM routers and one user-controlled Cisco 2950T-24 workgroup switch
- Classroom reference materials as follows:
  - QOS Student Guide
  - QOS Lab Guide
- Student pod workstation with Telnet or console access to workstation pod devices

Command List

The commands used in this exercise are described in the table here.

Baseline QoS Measurement Lab Commands

Command	Description
show ip interface [brief] [*type***] [***number***]**	To list a summary of an interface IP information and status
clear counters	To clear the interface counters
show interfaces [*interface-id***]**	Displays the administrative and operational status of all interfaces or a specified interface
shutdown	To disable an interface
copy running-config startup-config	Save your entries in the configuration file

Job Aids

These job aids are available to help you complete the laboratory exercise:

- Your assigned workgroup pod number provided by the instructor

Exercise Procedure

Complete these steps:

Step 1 Verify that the S0/0 and S0/1 interfaces on both of your workgroup routers (WGxR1 and WGxR2) routers are administratively enabled.

Step 2 Clear the interface counters on both your workgroup routers using the clear counters command.

Note: At this time, your instructor has not yet started the Pagent and Callgen router traffic generations.

Step 3 From the WGxR1 workgroup router, perform an extended ping to the WGxR2 router serial 0/0 interface then record the ping response time in the table at the end of the lab. For the extended ping, use a repeat count of 100 and a datagram size of 160.

```
WGxR1#ping
Protocol [ip]:
Target IP address: 10.2.x.2
Repeat count [5]: 100
Datagram size [100]: 160
Timeout in seconds [2]:
Extended commands [n]:
Sweep range of sizes [n]:
Type escape sequence to abort.
Sending 100, 160-byte ICMP Echos to 10.2.x.2, timeout is 2
seconds:
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
Success rate is 100 percent (100/100), round-trip min/avg/max
= 8/8/13 ms
```

Step 4 Repeat the extended ping two more times and record your results in the table at the end of this lab.

Step 5 Repeat Steps 3 and 4, but ping from the WGxR2 to WGxR1 serial 0/0 interface and record the response time results in the table at the end of this lab.

Step 6 From both of your workgroup routers, issue the show interfaces serial 0/0 command and record the highlighted stats below in the table at the end of the lab.

```
WGxR1#show interfaces serial0/0
Serial0/0 is up, line protocol is up
  Hardware is PowerQUICC Serial
  Description: to WGxR2
  Internet address is 10.2.3.1/24
  MTU 1500 bytes, BW 384 Kbit, DLY 20000 usec,
     reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation PPP, LCP Open
  Open: CDPCP, IPCP, loopback not set
  Last input 00:00:03, output 00:00:01, output hang never
  Last clearing of "show interface" counters 00:00:12
  Input queue: 0/75/0/0 (size/max/drops/flushes); Total output
drops: 0
  Queueing strategy: weighted fair
  Output queue: 0/1000/64/0 (size/max total/threshold/drops)
     Conversations  0/21/32 (active/max active/max total)
     Reserved Conversations 0/0 (allocated/max allocated)
     Available Bandwidth 288 kilobits/sec
  5 minute input rate 0 bits/sec, 0 packets/sec
  5 minute output rate 0 bits/sec, 0 packets/sec
     3 packets input, 429 bytes, 0 no buffer
     Received 0 broadcasts, 0 runts, 0 giants, 0 throttles
     0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored, 0
abort
     3 packets output, 184 bytes, 0 underruns
     0 output errors, 0 collisions, 0 interface resets
     0 output buffer failures, 0 output buffers swapped out
     0 carrier transitions
     DCD=up  DSR=up  DTR=up  RTS=up  CTS=up
```

Step 7 Clear the interface counters on both your workgroup routers using the clear counters command.

Step 8 Notify your instructor when you are done with the above steps.

Step 9 From the WGxR1 workgroup router, perform an extended ping to the WGxR2 router serial 0/0 interface then record the ping response time in the table at the end of the lab. For the extended ping, use a repeat count of 100 and a datagram size of 160.

Caution: Before initiating the extended ping, wait for the Pagent/Callgen traffic to run for at least 1 minute so that the traffic generation can stabilize.

```
WGxR1#ping
Protocol [ip]:
Target IP address: 10.2.x.2
Repeat count [5]: 100
Datagram size [100]: 160
Timeout in seconds [2]:
Extended commands [n]:
Sweep range of sizes [n]:
Type escape sequence to abort.
Sending 100, 160-byte ICMP Echos to 10.2.x.2, timeout is 2
seconds:
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
Success rate is 100 percent (100/100), round-trip min/avg/max
= 12/62/220 ms
```

Caution: After the Pagent traffic generation has been started, it is important that the Pagent routers not be used for network measurements using ping and trace, because the Pagent routers maintain a very high CPU load in generating the traffic demands for the QoS labs.

Step 10 Repeat the extended ping two more times and record your results in the table at the end of this lab.

Step 11 Repeat steps 9 and 10, but ping from the WGxR2 to WGxR1 serial 0/0 interface and record the response time results in the table at the end of this lab.

Step 12 From both of your workgroup routers, issue the show interfaces serial 0/0 command and record the highlighted stats below in the table at the end of the lab.

```
WGxR1#show interfaces serial 0/0
  Hardware is PowerQUICC Serial
  Description: to wgxr1
  Internet address is 10.2.x.1/24
  MTU 1500 bytes, BW 384 Kbit, DLY 20000 usec,
     reliability 255/255, txload 130/255, rxload 37/255
  Encapsulation PPP, LCP Open
  Open: CDPCP, IPCP, loopback not set
  Last input 00:00:02, output 00:00:00, output hang never
  Last clearing of "show interface" counters 00:01:20
  Input queue: 0/75/0/0 (size/max/drops/flushes); Total output
drops: 1023
  Queueing strategy: weighted fair
  Output queue: 0/1000/64/1023 (size/max
total/threshold/drops)
```



```
     Conversations  0/32/256 (active/max active/max total)
     Reserved Conversations 0/0 (allocated/max allocated)
     Available Bandwidth 288 kilobits/sec
  5 minute input rate 57000 bits/sec, 116 packets/sec
  5 minute output rate 196000 bits/sec, 211 packets/sec
     6595 packets input, 800344 bytes, 0 no buffer
     Received 0 broadcasts, 0 runts, 0 giants, 0 throttles
     0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored, 0
abort
     9418 packets output, 826272 bytes, 0 underruns
     0 output errors, 0 collisions, 0 interface resets
     0 output buffer failures, 0 output buffers swapped out
     0 carrier transitions
     DCD=up  DSR=up  DTR=up  RTS=up  CTS=up
```

Note: The traffic sent between Pagent-1 and Pagent-2 is set up such that the traffic rate varies constantly and will be different between the Pagent-1 and Pagent-2. When the traffic rate from Pagent-1 to Pagent-2 is high and increasing then the traffic rate from Pagent-2 to Pagent-1 will be low and decreasing and vice versa. As a result, the drop rate on your workgroup router serial 0/0 interface may be different between your workgroup R1 and R2 routers.

Note: For Callgen, both Callgen routers will generate VoIP calls at a constant rate.

Step 13 Compare the resulting statistics with and without the Pagent and Callgen traffic generations enabled.

You should notice that many of the pings would have a longer response time due to congestion on the low bandwidth 384-kbps PPP serial link.

Baseline QoS Measurement Lab Results

The following four tables are used to record the results of your lab testing in this exercise. Record your extended ping results in the two tables below:

WGxR1 to WGxR2 QoS Baseline ping Results

Packet Size	Without Pagent/Callgen	With Pagent/Callgen
160 bytes	min/avg/max	min/avg/max
Extended ping 1		
Extended ping 2		
Extended ping 3		
	success rate %	success rate %
Extended ping 1		
Extended ping 2		
Extended ping 3		

WGxR2 to WGxR1 QoS Baseline ping Results

Packet Size	Without Pagent/Callgen	With Pagent/Callgen
160 bytes	min/avg/max	min/avg/max
Extended ping 1		
Extended ping 2		
Extended ping 3		
	success rate %	success rate %
Extended ping 1		
Extended ping 2		
Extended ping 3		

Record your WGxR1 show interfaces serial 0/0 results in the following table:

WGxR1 QoS Baseline show interfaces Results

	Without Pagent/Callgen	With Pagent/Callgen
Queuing Strategy		
Reliability, Txload, Rxload		
Total Output Drops		
Output Queue: size/max total		
Output Queue: threshold/drops		
Packets Output		
Drop % (Calculated by you as: Total Output Drop / Packets Output)		

Record your WGxR2 show interfaces serial 0/0 results in the following table:

WGxR2 QoS Baseline show interfaces Results

	Without Pagent/Callgen	With Pagent/Callgen
Queuing Strategy		
Reliability, Txload, Rxload		
Total Output Drops		
Output Queue: size/max total		
Output Queue: threshold/drops		
Packets Output		
Drop % (Calculated by you as: Total Output Drop / Packets Output)		

Exercise Verification

You have completed this exercise when you attain these results:

- You have successfully completed the QoS baseline measurement by recording ping and interface statistics both before and after network traffic generation.

Lab Exercise 3-1: Configuring QoS with AutoQoS

Complete this lab exercise to practice what you learned in the related lesson.

Exercise Objective

In this exercise, you will configure QoS for VoIP on Cisco IOS routers and Catalyst switches using AutoQoS. After completing this exercise, you will be able to meet these objectives:

- Configure AutoQoS on Cisco IOS routers
- Configure AutoQoS on the Catalyst 2950 workgroup switch
- Use Cisco IOS monitoring commands and network connectivity tools (ping) to gather network response time data

Visual Objective

The figure illustrates what you will accomplish in this exercise.

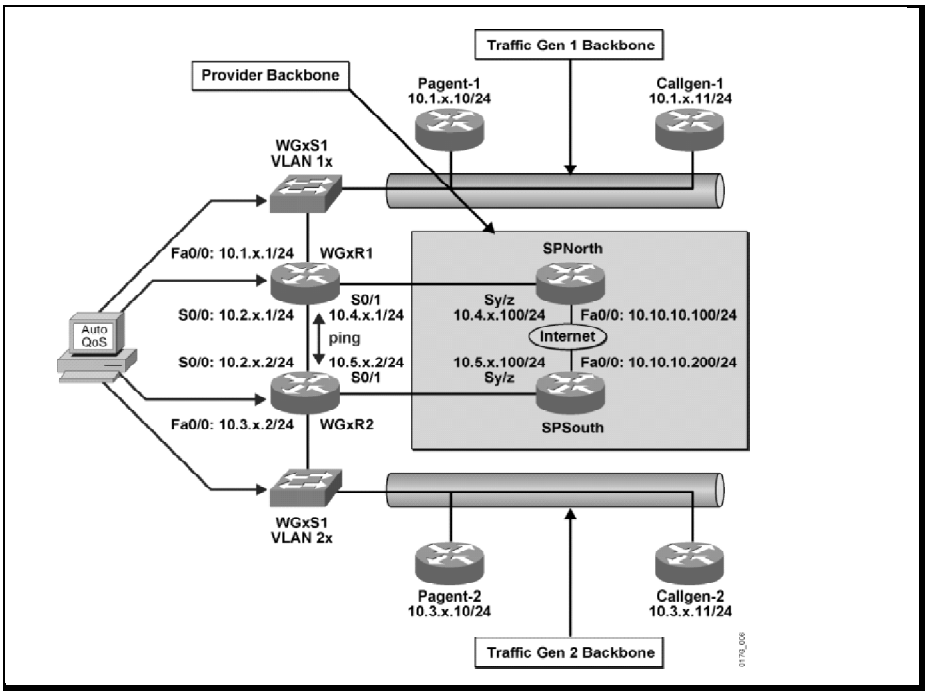

From the baseline measurement results, the E-Commerce University IT staff has determined that the drop rate and the latency of the VoIP traffic must be improved. At this point, the E-Commerce University has called upon you (the new CCNA network engineer they hired in the last month) to improve the voice quality as quickly as possible over the weekend.

By the way, the E-Commerce University network is built using Cisco Catalyst 2950 switches, Cisco 2610XM routers, and Cisco Aironet Wireless Access Points.

Through Cisco online e-learning, you discovered the new AutoQoS for VoIP feature that allows automated configuration of quality of service (QoS) on the network and provides a means for simplifying the implementation and provisioning of QoS for VoIP traffic.

Because you have only a limited amount of time to implement a solution, you have decided to go ahead and use AutoQoS and then test and compare the VoIP QoS results to see if AutoQoS can be used to solve problem.

Required Resources

These are the resources and equipment required to complete this exercise:

- Lab topology configured for QOS
- Student workgroup consisting of two user-controlled Cisco 2610XM routers and one user-controlled Cisco 2950T-24 workgroup switch
- Classroom reference materials as follows:
  - QOS Student Guide
  - QOS Lab Guide
- Student pod workstation with Telnet or console access to workstation pod devices

Command List

The commands used in this exercise are described in the table here.

Configuring QoS with AutoQoS Lab Router Commands

Command	Description
show running-config	To display the contents of the currently running configuration file
ip cef	To enable Cisco Express Forwarding (CEF) on the router
interface ***interface-id***	Enter interface configuration mode and the physical interface identification
auto qos voip	To configure the AutoQoS-VoIP feature on an interface
show auto qos [interface [*interface-type***]]**	To display the configurations created by the AutoQoS-VoIP feature on a specific interface or all interfaces
show ip interface [brief] [*type***] [***number***]**	To list a summary of an interface IP information and status
show interfaces multilink [*interface-id***]**	Displays the administrative and operational status of all interfaces or a specified interface
clear counters	To clear the interface counters
encapsulation *encapsulation-type*	To set the encapsulation method used by the interface
copy running-config startup-config	Save your entries in the configuration file

Configuring QoS with AutoQoS Lab Switch Commands

Command	Description
interface ***interface-id***	Enter interface configuration mode and the physical interface identification
auto qos voip	To configure automatic quality of service (AutoQoS) for voice over IP (VoIP) within a QoS domain
show auto qos [interface [*interface-id***]]**	Displays the automatic quality of service (AutoQoS) configuration that is applied
copy running-config startup-config	Save your entries in the configuration file

Job Aids

These job aids are available to help you complete the lab exercise:

- Your assigned workgroup pod number provided by the instructor

Task 1: Configuring AutoQoS on Cisco IOS Routers

In this task, you will enable the AutoQoS for VoIP feature on your workgroup router low bandwidth PPP serial interface.

Exercise Procedure

Complete these steps:

Step 1 Display and examine the running configuration of your WGxR1 router.

Step 2 Enable CEF on your WGxR1 router.

Step 3 Enable the AutoQoS for VoIP feature for traffic on the S0/0 interface of WGxR1 only. Do not configure AutoQoS to trust differentiated services code point (DSCP) markings.

Step 4 Display and examine the resulting AutoQoS configuration after enabling AutoQoS. The following example outputs are from WG1.

```
 ip access-list extended AutoQoS-VoIP-RTCP
  permit udp any any range 16384 32767
!
 ip access-list extended AutoQoS-VoIP-Control
  permit tcp any any eq 1720 (3 matches)
  permit tcp any any range 11000 11999
  permit udp any any eq 2427
  permit tcp any any eq 2428
  permit tcp any any range 2000 2002
  permit udp any any eq 1719
  permit udp any any eq 5060
!
 class-map match-any AutoQoS-VoIP-RTP-UnTrust
  match protocol rtp audio
  match access-group name AutoQoS-VoIP-RTCP
!
 class-map match-any AutoQoS-VoIP-Control-UnTrust
  match access-group name AutoQoS-VoIP-Control
!
 class-map match-any AutoQoS-VoIP-Remark
  match ip dscp ef
  match ip dscp cs3
  match ip dscp af31
!
  policy-map AutoQoS-Policy-UnTrust
   class AutoQoS-VoIP-RTP-UnTrust
    priority percent 70
    set dscp ef
   class AutoQoS-VoIP-Control-UnTrust
    bandwidth percent 5
    set dscp af31
   class AutoQoS-VoIP-Remark
    set dscp default
   class class-default
    fair-queue
Serial0/0 -
!
 interface Serial0/0
  no ip address
```

```
 encapsulation ppp
 no fair-queue
 ppp multilink
 multilink-group 2001100114
!
interface Multilink2001100114
 bandwidth 384
 ip address 10.2.1.1 255.255.255.0
 service-policy output AutoQoS-Policy-UnTrust
 ppp multilink
 ppp multilink fragment-delay 10
 ppp multilink interleave
 ip rtp header-compression iphc-format
!
 rmon event 33333 log trap AutoQoS description "AutoQoS SNMP
traps for Voice Drops" owner AutoQoS
 rmon alarm 33334 cbQosCMDropBitRate.1145.1147 30 absolute
rising-threshold 1 33333 falling-threshold 0 owner AutoQoS
```

Step 5 Repeat Steps 1 through 4 for WGxR2.

Step 6 Issue the **show ip interface brief** command on WGxR1 and ensure the Multilink interface is up. The Multilink interface is required for PPP multilink and interleaving operation. Notice that the S0/0 IP address assignment is automatically moved to the multilink interface.

```
WGxR1#show ip interface brief
Interface              IP-Address      OK? Method Status
Protocol
FastEthernet0/0        10.1.1.1        YES NVRAM  up
up
Serial0/0              unassigned      YES unset  up
up
Serial0/1              10.4.1.1        YES NVRAM
administratively down down
Virtual-Access1        unassigned      YES unset  up
up
Multilink2001100114    10.2.1.1        YES unset  up
up
```

Note: Because Callgen is used to generate the VoIP traffic, the voice quality of the VoIP Phone calls cannot be tested directly. Therefore, after AutoQoS has been enabled, you will modify the resulting QoS configurations to make the ping traffic (icmp echo and reply) to have the same EF PHB (per-hop behavior) as the VoIP traffic. This way, you can compare the extended ping responses with AutoQoS enabled to before AutoQoS was enabled.

Step 7 On the WGxR1 router, modify the ip access-list extended AutoQoS-VoIP-RTCP to include the ping traffic (icmp echo and echo reply).

```
ip access-list extended AutoQoS-VoIP-RTCP
 permit udp any any range 16384 32767
 permit icmp any any echo
 permit icmp any any echo-reply
```


Step 8 Repeat Step 7 for the WGxR2 router.

Exercise Verification

You have completed this exercise when you attain these results:

- You have successfully enabled the AutoQoS for VoIP feature on both WGxR1 and WGxR2
- You have configured ping (ICMP echo and reply) to belong to the same traffic class as VoIP traffic

Task 2: Configuring AutoQoS on the Catalyst 2950 Switch

In this task, you will enable the AutoQoS for VoIP feature on your workgroup Catalyst 2950 switch.

Exercise Procedure

Complete these steps:

Step 1 Display and examine the running configuration of your WGxS1 switch.

Step 2 Enable the AutoQoS for VoIP feature for traffic on the Fa0/1 interface of WGxS1 and trust the CoS markings from the core switch.

Step 3 Display and examine the resulting AutoQoS configuration after enabling AutoQoS. Notice that the 2950 is now configured for WRR queuing with queue 4 setup as the expedite queue (weight = 0). WRR queuing will be covered in the Congestion Management module.

```
Initial configuration applied by AutoQoS:
wrr-queue bandwidth 20 1 80 0
no wrr-queue cos-map
wrr-queue cos-map 1 0 1 2 4
wrr-queue cos-map 3 3 6 7
wrr-queue cos-map 4 5
mls qos map cos-dscp 0 8 16 26 32 46 48 56
!
interface FastEthernet0/1
 mls qos trust cos
 auto qos voip trust
```

Exercise Verification

You have completed this exercise when you attain these results:

- You have successfully enabled the AutoQoS for VoIP feature on WGxS1

Task 3: QoS Baseline with AutoQoS

In this task, you will use Cisco IOS monitoring commands and network connectivity tools (ping) to gather network response time data. You will compare the results of the traffic statistics of the network baseline statistics captured in Lab Exercise 2-2 and the network statistics after the application of AutoQoS.

Exercise Procedure

Complete these steps:

Step 1 Refer to Lab 2-2 and copy the baseline traffic information with the Pagent and Callgen traffic generation running into the tables at the end of this lab.

Step 2 From the WGxR1 workgroup router, perform an extended ping to the WGxR2 router serial 0/0 interface then record the ping response time in the table at the end of the lab. For the extended ping, use a repeat count of 100 and a datagram size of 160.

```
WGxR1#ping
Protocol [ip]:
Target IP address: 10.2.x.2
Repeat count [5]: 100
Datagram size [100]: 160
Timeout in seconds [2]:
Extended commands [n]:
Sweep range of sizes [n]:
Type escape sequence to abort.
Sending 100, 160-byte ICMP Echos to 10.2.x.2, timeout is 2
seconds:
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
Success rate is 100 percent (100/100), round-trip min/avg/max
= 28/49/89 ms
```

Step 3 Repeat the extended ping two more times and record your results in the table at the end of this lab.

Step 4 Repeat Steps 2 and 3, but ping from the WGxR2 to WGxR1 serial 0/0 interface and record the response time results in the table at the end of this lab.

Step 5 Clear the interface counters on both your workgroup routers using the clear counters command.

Step 6 Wait for the interface counters to accumulate traffic statistics for at least 1 minute.

Step 7 From both of your workgroup routers, issue the show interfaces multilink command and record the highlighted stats below in the table at the end of the lab.

```
WGxR1#show interfaces Multilink2001100114
MMultilink2001100114 is up, line protocol is up
  Hardware is multilink group interface
  Internet address is 10.2.x.2/24
  MTU 1500 bytes, BW 384 Kbit, DLY 100000 usec,
     reliability 255/255, txload 23/255, rxload 5/255
  Encapsulation PPP, LCP Open, multilink Open
  Open: IPCP, loopback not set
  DTR is pulsed for 2 seconds on reset
  Last input 00:00:09, output never, output hang never
  Last clearing of "show interface" counters 00:01:15
  Input queue: 0/75/0/0 (size/max/drops/flushes); Total output
drops: 1415
  Queueing strategy: weighted fair
  Output queue: 151/1000/64/1415/3541 (size/max
total/threshold/drops/interleaves)
     Conversations  21/30/128 (active/max active/max total)
     Reserved Conversations 1/1 (allocated/max allocated)
     Available Bandwidth 1 kilobits/sec
  5 minute input rate 9000 bits/sec, 19 packets/sec
  5 minute output rate 35000 bits/sec, 37 packets/sec
     4926 packets input, 217664 bytes, 0 no buffer
     Received 0 broadcasts, 0 runts, 0 giants, 0 throttles
     0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored, 0
abort
     11486 packets output, 1358207 bytes, 0 underruns
     0 output errors, 0 collisions, 4 interface resets
     0 output buffer failures, 0 output buffers swapped out
```

Step 8 Compare the results of the traffic statistics from Lab 2-2 QoS Baseline Measurement to the results from this lab.

Is the ping maximum response time shorter than before AutoQoS was enabled? Explain. _____

Is the drop rate higher or lower or about the same than before AutoQoS was enabled? Explain. _____

Step 9 Remove the AutoQoS configuration on your workgroup switch and workgroup routers.

```
WGxS1(config)#int fa0/1
WGxS1(config-if)#no auto qos voip
WGxS1(config-if)#end

WGxS1#show auto qos
AutoQoS is disabled

WGxR1(config)#int s0/0
WGxR1(config-if)#no auto qos voip
```



```
WGxR2(config)#int s0/0
WGxR2(config-if)#no auto qos voip
```

Note: When removing AutoQoS from WGxR1 and WGxR2, the encapsulation on your serial interfaces may be returned to HDLC, the Cisco default serial interface encapsulation. Be sure to reconfigure your serial interfaces for PPP encapsulation on both WGxR1 and WGxR2.

Step 10 On both of your workgroup routers, issue the **show ip interface brief** command and ensure the serial 0/0 interface is up. The multilink interface should be removed when AutoQoS has been disabled.

```
WGxR1#sh ip int brief
Interface       IP-Address OK? Method     Status
Protocol
FastEthernet0/0 10.1.1.1   YES NVRAM      up
up
Serial0/0       10.2.x.1   YES unset      up
up
Serial0/1       10.4.1.1   YES NVRAM      administratively
down down
Virtual-Access1 unassigned YES unset      up
up
```

Step 11 Save your running configurations of the workgroup routers and the workgroup switch to the startup-config in NVRAM.

Configuring QoS with AutoQoS Lab Results

Record your extended ping results in the two tables below:

WGxR1 to WGxR2 AutoQoS ping Results

Packet Size	Without AutoQoS (From Lab 2-2)	With AutoQoS (This Lab)
160 bytes	min/avg/max	min/avg/max
Extended ping 1		
Extended ping 2		
Extended ping 3		
	success rate %	success rate %
Extended ping 1		
Extended ping 2		
Extended ping 3		

WGxR2 to WGxR1 AutoQoS ping Results

Packet Size	Without AutoQoS (From Lab 2-2)	With AutoQoS (This Lab)
160 bytes	min/avg/max	min/avg/max
Extended ping 1		
Extended ping 2		
Extended ping 3		
	success rate %	success rate %
Extended ping 1		
Extended ping 2		
Extended ping 3		

Record your WGxR1 show interfaces multilink results in the following table:

WGxR1 AutoQoS show interfaces Results

	Without AutoQoS (From Lab 2-2) show interface s0/0	With AutoQoS (This Lab) show interface multilink
Queuing Strategy		
Reliability, Txload, Rxload		
Total Output Drops		
Output Queue: size/max total		
Output Queue: threshold/drop/interleaves		
Packets Output		
Drop % (Calculated by you as: Total Output Drop / Packets Output)		

Record your WGxR2 show interfaces serial 0/0 results in the following table:

WGxR2 AutoQoS show interfaces Results

	Without AutoQoS (From Lab 2-2) show interface s0/0	With AutoQoS (This Lab) show interface multilink
Queuing Strategy		
Reliability, Txload, Rxload		
Total Output Drops		
Output Queue: size/max total		
Output Queue: threshold/drop/interleaves		
Packets Output		
Drop % (Calculated by you as: Total Output Drop / Packets Output)		

Exercise Verification

You have completed this exercise when you attain these results:

- You have successfully completed the QoS baseline measurement after enabling AutoQoS by recording ping and interface statistics.
- You have compared the results of the traffic measurement to those from the QoS Baseline Lab and correctly answered the questions contained within the lab.
- You have successfully remove AutoQoS configuration from your workgroup routers and switch.

Lab Exercise 4-1: Classification and Marking Using MQC

Complete this lab exercise to practice what you learned in the related lesson.

Exercise Objective

In this exercise, you will configure classification using Modular QoS command-line interface (CLI) (MQC) and marking using class-based marking. After completing this exercise, you will be able to meet these objectives:

- Configure an IP extended access-list matching specific traffic for use in MQC classification
- Configure MQC classification
- Configure class-based marking

Visual Objective

The figure illustrates what you will accomplish in this exercise.

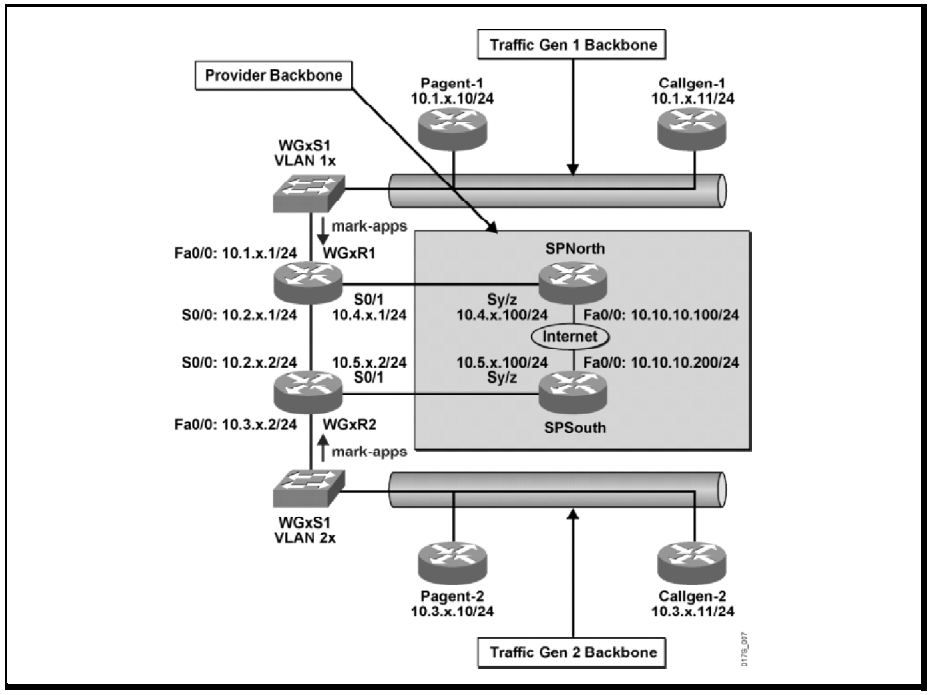

From the AutoQoS results, the E-Commerce University IT staff has determined that the VoIP voice quality is now satisfactory. Because you did such a great job so far, the E-Commerce University is now calling upon you to also improve the response time of the Oracle (SQL) student administration database application over their spring break. Because AutoQoS only works for voice traffic currently, you decided to remove the AutoQoS configurations and to implement the proper QoS mechanisms manually using MQC.

As you have learned in the Cisco QoS course, one of the first steps to implement QoS is to properly classify and mark the traffic. Therefore, you decide the first step now for you is to implement the proper classification and marking.

To ease into the QoS implementation, you decide to first approach the classification and marking of the applications you feel are the bulk of the University traffic (FTP and HTTP).

Required Resources

These are the resources and equipment required to complete this exercise:

- Lab topology configured for QOS
- Student workgroup consisting of two user-controlled Cisco 2610XM routers and one user-controlled Cisco 2950T-24 workgroup switch
- Classroom reference materials as follows:
  - QOS Student Guide
  - QOS Lab Guide
- Student pod workstation with Telnet or console access to workstation pod devices

Command List

The commands used in this exercise are described in the table here.

Classification and Marking Using MQC Lab Commands

Command	Description
access-list *access-list-number* {**deny** | **permit**} **tcp src** *src-wildcard* [**operator** [*port*]] **dest** *dest-wildcard* [**operator** [*port*]]	To define an extended IP access list for Transmission Control Protocol (TCP) based traffic
class-map *class-map-name*	To create a class map to be used for matching packets to a specified class
match access-group {*access-group* | **name** *access-group-name*}	To configure the match criteria for a class map on the basis of the specified access control list (ACL)
policy-map *policy-map-name*	To create or modify a policy map that can be attached to one or more interfaces
class {*class-name* | **class-default**}	To specify the name of the class whose policy you want to create or change or to specify the default class
set dscp *dscp-value*	To mark a packet by setting the differentiated services code point (DSCP)
service-policy {**input** | **output**} *policy-map-name*	To attach a policy map to an input interface or virtual circuit (VC), or an output interface or VC
show access-lists [*access-list-number* | *access-list-name*]	To display the contents of current access lists
show class-map [*class-map-name*]	To display all class maps and their matching criteria
show policy-map [*policy-map*]	To display the configuration of all classes for a specified service policy map or all classes for all existing policy maps
show policy-map interface *interface-name* [**input** | **output**] [**class** *class-map-name*]	To display the packet statistics of all classes that are configured for all service policies either on the specified interface or subinterface
copy running-config startup-config	Save your entries in the configuration file.

Job Aids

These job aids are available to help you complete the laboratory exercise:

- Your assigned workgroup pod number provided by the instructor

Exercise Procedure

Complete these steps:

Step 1 Connect to the WGxR1 router. Configure an IP extended access-list to match all FTP traffic.

Step 2 On the WGxR1 router, configure a second IP extended access-list to match all HTTP (WWW) traffic.

Step 3 Display and verify your IP extended access-list configuration.

```
Extended IP access list 101
    10 permit tcp any any eq ftp
    20 permit tcp any any eq ftp-data
Extended IP access list 102
    10 permit tcp any any eq www
```

Step 4 Create two new class maps called, **match-ftp** and **match-www** to match the FTP and WWW traffic, respectively.

Step 5 Display and verify your class-map configuration.

```
Class Map match-any class-default (id 0)
 Match any

Class Map match-all match-ftp (id 1)
  Match access-group 101

Class Map match-all match-www (id 2)
  Match access-group 102
```

Step 6 Create a policy map on your workgroup WGxR1 router, named **mark-apps**, that includes the two newly configured traffic classes (match-ftp and match-www). Use class-based marking to mark the FTP traffic to AF 11 and the WWW traffic to DSCP 0.

Step 7 Display and verify your policy-map configuration.

```
Policy Map mark-apps
  Class match-ftp
    set dscp af11
  Class match-www
    set dscp default
```

Step 8 Apply the policy map to the FastEthernet 0/0 interface of your WGxR1 router in the inbound direction.

Step 9 Display and verify your service policy.

```
FastEthernet0/0

  Service-policy input: mark-apps

    Class-map: match-ftp (match-all)
      320 packets, 19200 bytes
      5 minute offered rate 1000 bps, drop rate 0 bps
      Match: access-group 101
      QoS Set
        dscp af11
          Packets marked 320

    Class-map: match-www (match-all)
      172 packets, 10320 bytes
      5 minute offered rate 1000 bps, drop rate 0 bps
      Match: access-group 102
      QoS Set
        dscp default
          Packets marked 172

    Class-map: class-default (match-any)
```

```
            12591 packets, 827819 bytes
            5 minute offered rate 26000 bps, drop rate 0 bps
            Match: any
```

Step 10 How many packets have been matched and marked on WGxR1 for each of the traffic class?

Class ftp _____

Class www _____

Class class-default _____

Step 11 Repeat Steps 1 through 9 for workgroup router WGxR2.

Step 12 How many packets have been matched and marked on WGxR2 for each of the traffic class?

Class ftp _____

Class www _____

Class class-default _____

Step 13 Save your running configurations of the workgroup routers and the workgroup switch to the startup-config in NVRAM.

Exercise Verification

You have completed this exercise when you attain these results:

- You have correctly created MQC classification for FTP and HTTP traffic.
- You have correctly configured MQC marking to mark FTP traffic as AF11 and HTTP traffic as DSCP 0.
- You have correctly enabled MQC classification and marking by applying the service policy to both workgroup routers.

Lab Exercise 4-2: Classification Using NBAR

Complete this lab exercise to practice what you learned in the related lesson.

Exercise Objective

In this exercise, you will configure classification using network-based application recognition. After completing this exercise, you will be able to meet these objectives:

- Discover network applications and traffic using NBAR protocol discovery
- Configure classification of discovered applications using NBAR classification
- Configure class-based marking

Visual Objective

The figure illustrates what you will accomplish in this exercise.

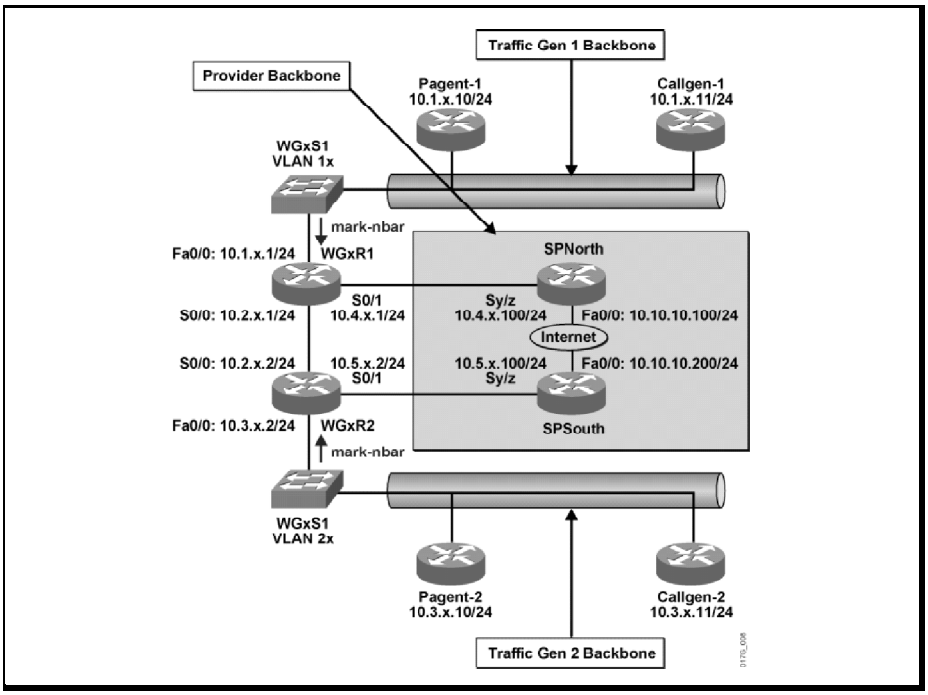

After studying your current classification and marking strategy, you realize that using extended IP access lists cannot properly classify all the traffic now running on the network. One issue is that there are too many applications being classified into the class-default. Therefore, you decide to configure network-based application recognition (NBAR) for your classification requirements. Before you configure protocol matching with NBAR, you decide first to analyze the network using NBAR protocol discovery to validate your assumptions about the traffic currently traversing the network. After all traffic has been properly identified, you plan to configure a new class-based marking policy to classify and mark the applications running on the network.

Required Resources

These are the resources and equipment required to complete this exercise:

- Lab topology configured for QOS
- Student workgroup consisting of two user-controlled Cisco 2610XM routers and one user-controlled Cisco 2950T-24 workgroup switch
- Classroom reference materials as follows:
  - QOS Student Guide
  - QOS Lab Guide
- Student pod workstation with Telnet or console access to workstation pod devices

Command List

The commands used in this exercise are described in the table here.

Classification and Marking Using NBAR Lab Commands

Command	Description
no service-policy {input | output} *policy-map-name*	To remove a service policy from an input or output interface
[no] ip nbar protocol-discovery	To configure NBAR to discover traffic for all protocols known to NBAR on a particular interface
show ip nbar protocol-discovery [interface *interface-spec*]	To display the statistics gathered by the NBAR protocol discovery feature
clear ip nbar protocol-discovery	Clear NBAR protocol discovery statistics
ip access-list {standard | extended} *access-list-name*	To define an IP access list by name
permit tcp *source source-wildcard destination destination-wildcard* **[operator [port]]**	To set conditions to allow a packet to pass a named IP access list
permit udp *source source-wildcard destination destination-wildcard* **[operator [port]]**	To set conditions to allow a packet to pass a named IP access list
class-map *class-map-name*	To create a class map to be used for matching packets to a specified class
match protocol *protocol-name*	To configure the match criteria for a class map on the basis of the specified protocol
match access-group {access-group | name *access-group-name*}	To configure the match criteria for a class map on the basis of the specified ACL
policy-map *policy-map-name*	To create or modify a policy map that can be attached to one or more interfaces
class {*class-name* | **class-default**}	To specify the name of the class whose policy you want to create or change or to specify the default class
set dscp *dscp-value*	To mark a packet by setting the differentiated services code point (DSCP)
service-policy {input | output} *policy-map-name*	To attach a policy map to an input interface or virtual circuit (VC), or an output interface or VC
show class-map [*class-map-name***]**	To display all class maps and their matching criteria
show policy-map [*policy-map***]**	To display the configuration of all classes for a specified service policy map or all classes for all existing policy maps
show policy-map interface *interface-name* **[input | output] [class** *class-map-name*]	To display the packet statistics of all classes that are configured for all service policies either on the specified interface or subinterface
copy running-config startup-config	Save your entries in the configuration file.

Job Aids

These job aids are available to help you complete the laboratory exercise:

- Your assigned workgroup pod number provided by the instructor

Exercise Procedure

Complete these steps:

Step 1 Disable the existing inbound service policy on the FastEthernet 0/0 interface of your WGxR1 router.

Step 2 Verify that Cisco Express Forwarding is enabled on your WGxR1 router.

Step 3 Enable NBAR protocol discovery on the FastEthernet 0/0 interface of your WGxR1 router.

Step 4 Clear the NBAR protocol discovery counters on WGxR1.

Step 5 Wait for the interface counters to accumulate traffic statistics for at least 1 minute.

Step 6 Display the NBAR protocol discovery traffic statistics of all NBAR discovered protocols.

```
FastEthernet0/0
                    Input                     Output
   Protocol         Packet Count              Packet Count
                    Byte Count                Byte Count
                    5 minute bit rate (bps)   5 minute bit rate
(bps)
   ---------------- ------------------------ ------------------
------
   sqlnet             157                       404
                      9420                      60678
                      1000                      3000
   citrix             201                       547
                      13642                     51549
                      1000                      3000
   http               86                        255
                      13353                     59838
                      0                         3000
   napster            98                        272
                      5880                      54727
                      0                         3000

[rest omitted]
```


Step 7 In the space provided below, list the protocols discovered by NBAR protocol discovery:

1.	2.	3.
4.	5.	6.
7.	8.	9.
10.	11.	

Step 8 Disable NBAR protocol discovery from the FastEthernet 0/0 interface on WGxR1.

Step 9 Repeat Steps 1 through 8 for workgroup router WGxR2.

Step 10 On the WGxR1 router, configure NBAR classification and MQC marking to classify inbound traffic on the FastEthernet 0/0 interface and mark it as outlined in the table below. Completion of this lab step will require the configuration of 5 new class maps (one for each service class) and the configuration of a policy map (called **mark-nbar**) that marks traffic in each class appropriately.

Note: Remember by default, a class map is set to match all. If you are matching multiple protocols into the same class, remember to use match-any instead of match-all.

Class Name (class-map name)	Protocol	PHB
real-time	rtp/rtcp	EF
real-time	icmp	EF
mission-critical	sqlnet voice-control (port specific tcp and udp)	AF 31
interactive	citrix	AF 21
bulk	ftp	AF 11
scavenger	kazaa2	CS 1
scavenger	napster	CS 1
class-default	all others	BE

When using NBAR to match RTP packets, there is one limitation in that protocol matching for RTP does not match control packets. This is somewhat of an advantage in that it is preferred that voice bearer traffic be separated from voice control traffic as each traffic type receives different QoS treatment (EF for voice bearer and AF31 for voice control).

Create the following named access list for matching RTCP traffic:

```
ip access-list extended VoIP-RTCP
    permit udp any any range 16384 32767
```

Match the named access-list **VoIP-RTCP** into the real-time traffic class along with the RTP and ICMP traffic.

The mission-critical class should be configured to contain both the sqlnet traffic and voice-control traffic. Use the following ACL to match voice control traffic when creating the mission-critical class on your router:

```
ip access-list extended Voice-Control
 permit tcp any any eq 1720
 permit tcp any any range 11000 11999
 permit udp any any eq 2427
 permit tcp any any eq 2428
 permit tcp any any range 2000 2002
 permit udp any any eq 1719
 permit udp any any eq 5060
```

Note: Recall that Internet Control Message Protocol (ICMP) traffic has been placed in the real-time class as a means of measuring QoS performance.

Step 11 Display and verify your class-map configuration.

```
Class Map match-all bulk (id 8)
  Match protocol ftp

Class Map match-any real-time (id 5)
  Match protocol rtp
  Match protocol icmp
  Match access-group name VoIP-RTCP

Class Map match-any mission-critical (id 6)
  Match protocol sqlnet
  Match access-group name Voice-Control

Class Map match-all interactive (id 7)
  Match protocol citrix

Class Map match-any scavenger (id 9)
  Match protocol kazaa2
  Match protocol napster

[rest omitted]
```

Step 12 Display and verify your policy-map configuration.

```
Policy Map mark-nbar
  Class real-time
    set dscp ef
  Class mission-critical
    set dscp af31
  Class interactive
    set dscp af21
```



```
  Class bulk
    set dscp af11
  Class scavenger
    set dscp cs1
  Class class-default
    set dscp default
```

Step 13 Apply the policy map to the FastEthernet 0/0 interface of your WGxR1 router in the inbound direction.

Step 14 Wait for the interface counters to accumulate traffic statistics for at least 1 minute.

Step 15 Display and verify your service policy on interface FastEthernet 0/0.

```
FastEthernet0/0
  Service-policy input: mark-nbar

  Class-map: real-time (match-any)
      542 packets, 115388 bytes
      5 minute offered rate 4000 bps, drop rate 0 bps
      Match: protocol rtp
        4 packets, 856 bytes
        5 minute rate 0 bps
      Match: protocol icmp
        1 packets, 70 bytes
        5 minute rate 0 bps
      Match: access-group name AutoQoS-VoIP-RTCP
        537 packets, 114462 bytes
        5 minute rate 4000 bps
      QoS Set
        dscp ef
          Packets marked 542


    Class-map: mission-critical (match-any)
      366 packets, 54424 bytes
      5 minute offered rate 4000 bps, drop rate 0 bps
      Match: protocol sqlnet
        331 packets, 50194 bytes
        5 minute rate 3000 bps
      Match: access-group name Voice-Control
        35 packets, 4230 bytes
        5 minute rate 2000 bps
      QoS Set
        dscp af31
          Packets marked 366

    Class-map: interactive (match-all)
      262 packets, 15720 bytes
      5 minute offered rate 2000 bps, drop rate 0 bps
      Match: protocol citrix
      QoS Set
        dscp af21
          Packets marked 295

[rest omitted]
```

Step 16 How many packets have been matched and marked for each of the traffic class?

Class real-time _____

Class mission-critical _____

Class interactive _____

Class bulk _____

Class scavenger _____

Class class-default _____

Note: If the real-time, mission-critical, and scavenger classes have no matches, verify your class-map configuration to ensure they are set to match-any and not match-all.

Step 17 Repeat Steps 10 through 15 for workgroup router WGxR2.

Step 18 How many packets have been matched and marked for each of the traffic class?

Class real-time _____

Class mission-critical _____

Class interactive _____

Class bulk _____

Class scavenger _____

Class class-default _____

Step 19 Save your running configurations of the workgroup routers and the workgroup switch to the startup-config in NVRAM.

Exercise Verification

You have completed this exercise when you attain these results:

- You have properly configured NBAR protocol discovery to identify network applications.
- You have correctly configured NBAR classification.
- You have correctly configured class-based marking of NBAR classified traffic.

Lab Exercise 4-3: Configuring QoS Pre-Classify

Complete this lab exercise to practice what you learned in the related lesson.

Exercise Objective

In this exercise, you will configure and examine QoS pre-classify. After completing this exercise, you will be able to meet these objectives:

- Configure QoS pre-classify on a Cisco IOS router using a GRE tunnel
- Monitor QoS pre-classify configurations

Visual Objective

The figure illustrates what you will accomplish in this exercise.

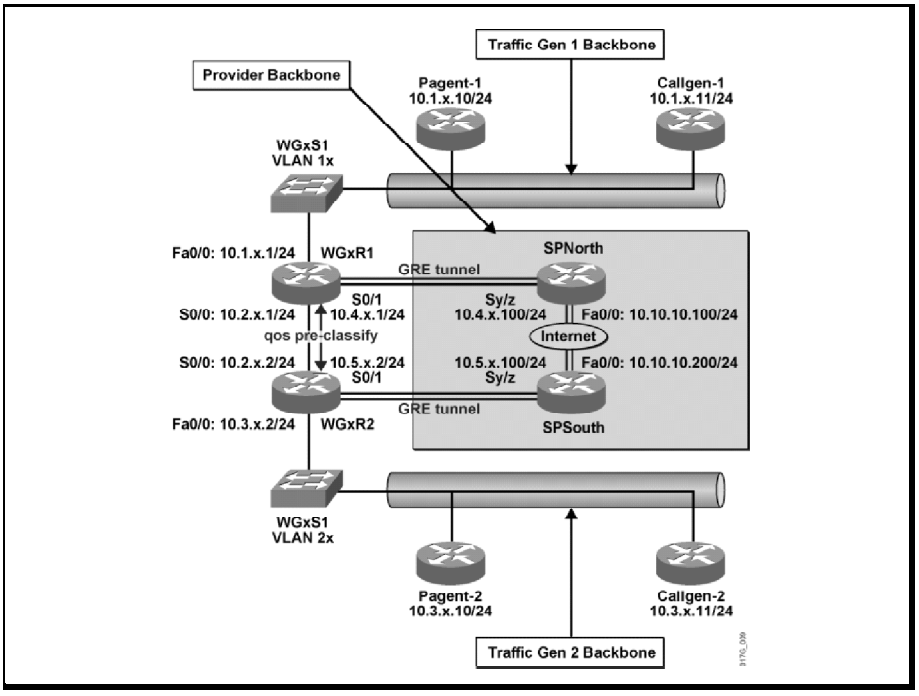

The E-Commerce University IT staff has decided to implement a generic routing encapsulation (GRE) tunnel between the north and south campus via the existing 768-kbps Internet connection. After the GRE tunnel is set up and working properly, IPSec will also be enabled over the GRE tunnel. At this point, the E-Commerce University IT staff needs you to first configure and test the GRE tunnel (without IPSec). They university plans to send different types of traffic over the tunnel and would like to be able to differentiate between different traffic flows so that QoS can be applied. In this lab, you will configure and verify the qos pre-classify feature for traffic classification over a GRE tunnel.

Required Resources

These are the resources and equipment required to complete this exercise:

- Lab topology configured for QOS
- Student workgroup consisting of two user-controlled Cisco 2610XM routers and one user-controlled Cisco 2950T-24 workgroup switch
- Classroom reference materials as follows:
  - QOS Student Guide
  - QOS Lab Guide
- Student pod workstation with Telnet or console access to workstation pod devices

Command List

The commands used in this exercise are described in the table here.

Configuring VPN QoS Lab Commands

Command	Description
interface *interface-id*	Enter interface configuration mode and the physical interface identification
qos pre-classify	To enable QoS preclassification
ip cef	To enable Cisco Express Forwarding (CEF) on the router
[no] shutdown	To disable an interface
show ip interface [brief] [*type***] [***number***]**	To list a summary of an interface IP information and status
copy running-config startup-config	Save your entries in the configuration file.

Job Aids

These job aids are available to help you complete the laboratory exercise:

- Your assigned workgroup pod number provided by the instructor

Exercise Procedure

Complete these steps:

Step 1 Verify the S0/0 and S0/1 interfaces of your workgroup WGxR1 and WGxR2 routers are UP.

```
WGxR1#show ip interface brief

Interface        IP-Address      OK? Method Status
Protocol
FastEthernet0/0  10.3.x.2        YES NVRAM  up
up
Serial0/0        10.2.x.2        YES unset  up
up
Serial0/1        10.5.x.2        YES NVRAM  up
up
```

Step 2 Verify CEF switching is still enabled on both workgroup routers in your pod.

Step 3 Configure a GRE tunnel between your WGxR1 and WGxR2 router via the service provider core as follows:

```
! WGxR1
!
 interface Tunnel0
 ip unnumbered fastethernet0/0
 tunnel source Serial0/1
 tunnel destination 10.5.x.2


! WGxR2
!
 interface Tunnel0
 ip unnumbered fastethernet0/0
 tunnel source Serial0/1
 tunnel destination 10.4.x.1
```

Step 4 Configure a static route via the tunnel 0 interface to the 10.1.x.0 or 10.3.x.0 subnet on the respective WGxR router as follows:

```
! WGxR1
!
 ip route 10.3.x.0 255.255.255.0 Tunnel0


! WGxR2
!
 ip route 10.1.x.0 255.255.255.0 Tunnel0
```

Step 5 Display the tunnel interface and verify it is UP and operational.

```
WGxR2#show interface tunnel 0
Tunnel0 is up, line protocol is up
  Hardware is Tunnel
  Interface is unnumbered. Using address of FastEthernet0/0
(10.3.x.2)
  MTU 1514 bytes, BW 9 Kbit, DLY 500000 usec,
     reliability 255/255, txload 81/255, rxload 196/255
  Encapsulation TUNNEL, loopback not set
  Keepalive not set
  Tunnel source 10.5.x.2 (Serial0/1), destination 10.4.x.1
  Tunnel protocol/transport GRE/IP, key disabled, sequencing
disabled
  Tunnel TTL 255
  Checksumming of packets disabled,  fast tunneling enabled
  Last input 00:00:03, output 00:00:04, output hang never
  Last clearing of "show interface" counters never
  Input queue: 0/75/0/0 (size/max/drops/flushes); Total output
drops: 0
  Queueing strategy: fifo
  Output queue: 0/0 (size/max)
  5 minute input rate 25000 bits/sec, 9 packets/sec
  5 minute output rate 39000 bits/sec, 23 packets/sec
     87865 packets input, 12254589 bytes, 0 no buffer
           Received 0 broadcasts, 0 runts, 0 giants, 0
throttles
           0 input errors, 0 CRC, 0 frame, 0 overrun, 0
ignored, 0 abort
           282049 packets output, 54114890 bytes, 0 underruns
           0 output errors, 0 collisions, 0 interface resets
```

Step 6 Telnet to the Callgen-1 (10.1.x.11) router and perform a trace to the Callgen-2 (10.3.x.11) router. Verify the path goes thru your GRE tunnel.

```
Callgen-1#trace 10.3.x.11

      Type escape sequence to abort.
      Tracing the route to 10.3.x.11

      1 10.1.x.1 4 msec 0 msec 0 msec
      2 10.3.x.11 140 msec *  8 msec
```

Step 7 Telnet to the Callgen-2 router and perform a trace to the Callgen-1 router. Verify the path goes through your GRE tunnel.

```
Callgen-2#trace 10.1.x.11

      Type escape sequence to abort.
      Tracing the route to 10.1.x.11

      1 10.3.x.2 60 msec 56 msec 64 msec
      2 10.1.x.11 60 msec *  56 msec
```


Step 8 Issue the **show queue serial0/1** command to display the WFQ information for the serial 0/1 interface.

```
WGxR1#show queue serial0/1


  Input queue: 0/75/0/0 (size/max/drops/flushes); Total output
drops: 9269
  Queueing strategy: weighted fair
  Output queue: 42/1000/64/9269 (size/max
total/threshold/drops)
     Conversations  1/31/256 (active/max active/max total)
     Reserved Conversations 0/0 (allocated/max allocated)
     Available Bandwidth 576 kilobits/sec


  (depth/weight/total drops/no-buffer drops/interleaves)
41/32384/9269/0/0
  Conversation 95, linktype: ip, length: 206
  source: 10.5.8.2, destination: 10.4.8.1, id: 0xC2FA, ttl:
255, prot: 47
```

Note: You may have to run the **show queue** *interface* command a few times until you catch an active flow. If you are not able to see packets in the queue after several attempts, log onto your other workgroup router and repeat Step 8.

How many active flows does the router see? _____

Notice that when QoS pre-classify is not configured, the output interface sees only one flow, which is Protocol 47 (GRE).

Step 9 Configure QoS pre-classify feature on the tunnel interface.

Step 10 Issue the **show queue serial0/1** command again to display the WFQ information for the serial 0/1 interface.

```
WGxR1#show queue serial0/1


  Input queue: 0/75/0/0 (size/max/drops/flushes); Total output
drops: 117791
  Queueing strategy: weighted fair
  Output queue: 157/1000/64/117791 (size/max
total/threshold/drops)
     Conversations  21/31/256 (active/max active/max total)
     Reserved Conversations 0/0 (allocated/max allocated)
     Available Bandwidth 576 kilobits/sec


  (depth/weight/total drops/no-buffer drops/interleaves)
7/32384/0/0/0
```

```
  Conversation 119, linktype: ip, length: 72
  source: 10.1.1.10, destination: 10.3.1.10, id: 0xFF4C, ttl:
254,
  TOS: 104 prot: 6, source port 1063, destination port 1521


  (depth/weight/total drops/no-buffer drops/interleaves)
7/32384/0/0/0
  Conversation 212, linktype: ip, length: 196
  source: 10.1.1.10, destination: 10.3.1.10, id: 0xE642, ttl:
127,
  TOS: 72 prot: 6, source port 1102, destination port 1494


  (depth/weight/total drops/no-buffer drops/interleaves)
7/32384/0/0/0
  Conversation 64, linktype: ip, length: 79
  source: 10.1.1.10, destination: 10.3.1.10, id: 0x908A, ttl:
126,
  TOS: 184 prot: 17, source port 49590, destination port 49602


 <output omitted>
```

Note: You may have to run the **show queue** *interface* command a few times until you catch some active flows.

What happens when QoS pre-classify is configured on the tunnel interface?

_____

Can WFQ now distinguish between different application flows? Or does WFQ still only see one flow (protocol 47=GRE)?

_____

Step 11 On the WGxR1 and WGxR2 routers, remove the static route via the tunnel interface.

Step 12 Save your running configurations of the workgroup routers and the workgroup switch to the startup-config in NVRAM.

Exercise Verification

You have completed this exercise when you attain these results:

- You have successfully configured a tunnel interface between WGxR1 and WGxR2 through the service provider backbone
- You have successfully enabled the QoS pre-classify feature on the tunnel interface
- You have successfully examined the difference in traffic flows with and without the QoS pre-classify feature enabled

Lab Exercise 4-4: LAN-Based Packet Classification and Marking

Complete this lab exercise to practice what you learned in the related lesson.

Exercise Objective

In this exercise, you will configure LAN based classification and marking. After completing this exercise, you will be able to meet these objectives:

- Configure a trust boundary on a Catalyst 2950 switch to only trust Cisco IP Phones
- Configure CoS-to-DSCP mapping on a Catalyst 2950 switch
- Configure IP access lists and class-based marking on a Catalyst 2950 switch to mark traffic
- Verify the QoS markings from the workgroup switch using the workgroup routers

Visual Objective

The figure illustrates what you will accomplish in this exercise.

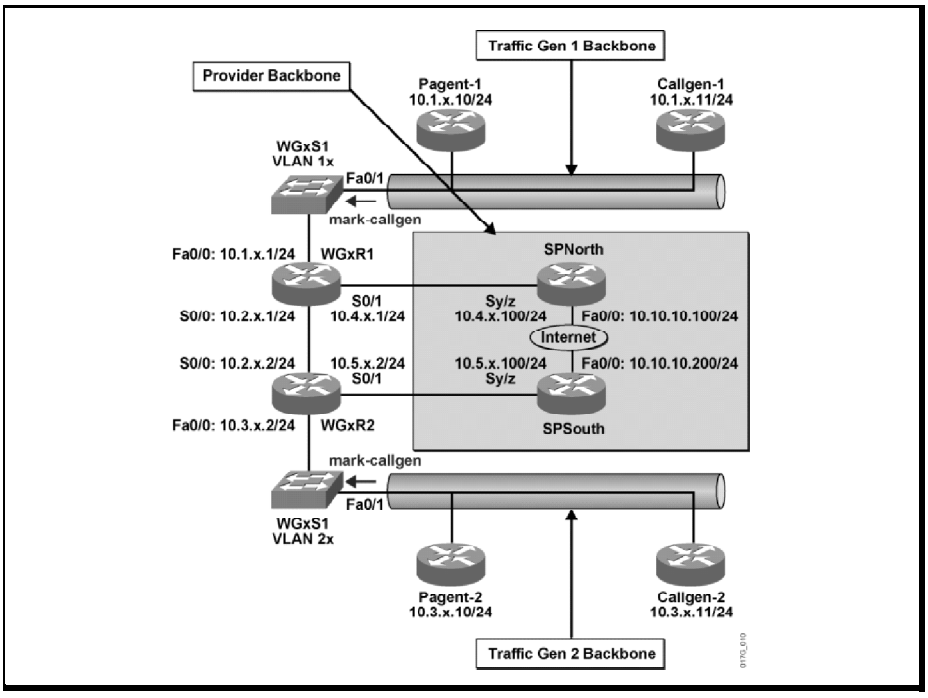

Your research into classification and marking policy has you concerned that students might inject their traffic into the network with CoS and DSCP markings that are not in accordance with the new QoS policy. Not forgetting the QoS requirements of the Cisco IP Phones connected in the wiring closet, you decide to establish a trust boundary that extends to these phones.

In addition, you also wish to classify and mark traffic as close to the source as possible. Studying the default CoS to DSCP marking maps on your Catalyst 2950 switches, you discover that CoS 5 (which is set in your IP Phone endpoints automatically) is currently mapped to

DSCP 40. You would like this to map to EF (DSCP 46) and decide to change the default CoS to DSCP mapping in your Catalyst 2950 switches.

One more issue arises in that the network also contains IP voice gateway devices that generate G.711 voice traffic, but cannot mark CoS or DSCP. You will have to implement classification and marking at your trust boundary using access lists to accommodate these voice gateways.

Required Resources

These are the resources and equipment required to complete this exercise:

- Lab topology configured for QOS
- Student workgroup consisting of two user-controlled Cisco 2610XM routers and one user-controlled Cisco 2950T-24 workgroup switch
- Classroom reference materials as follows:
  - QOS Student Guide
  - QOS Lab Guide
- Student pod workstation with Telnet or console access to workstation pod devices

Command List

The commands used in this exercise are described in the table here.

LAN Based Packet Classification and Marking Lab Commands

Command	Description
mls qos trust [cos | device cisco-phone | dscp]	Configure the port trust state
switchport priority extend {cos *value* **| trust}**	Sets a port priority for the incoming untagged frames or the priority of frames received by the IP Phone connected to the specified port
mls qos map {cos-dscp *dscp1...dscp8* **| dscp-cos** *dscp-list* **to** *cos***}**	Define the class of service (CoS)-to-DSCP map or DSCP-to-CoS map
access-list *access-list-number* **{deny | permit | remark} {source** *source-wildcard* **| host** *source* **| any}**	Configures a standard IP ACL
class-map *class-map-name*	Create a class map to be used for matching packets to the class whose name you specify
match {access-group *acl-index* **| access-group name** *acl-name* **| ip dscp** *dscp-list***}**	Define the match criteria to classify traffic
policy-map *policy-map-name*	Create or modify a policy map that can be attached to multiple interfaces
class {*class-name* **| class-default}**	To specify the name of the class whose policy you want to create or change or to specify the default class
set ip dscp *new-dscp*	Classify IP traffic by setting a DSCP value
service-policy input *policy-map-name*	Apply a policy map defined by the **policy-map** command to the input of a particular interface
show access-lists [*name* **|** *number***]**	Display ACLs configured on the switch
show class-map [*class-map-name***]**	Display QoS class maps
show policy-map [*policy-map-name***]**	Display QoS policy maps
show mls qos maps [cos-dscp | dscp-cos]	Display QoS mapping information
show mls qos interface [*interface-id***]**	Display QoS information at the interface level
copy running-config startup-config	Save your entries in the configuration file.

Job Aids

These job aids are available to help you complete the laboratory exercise:

- Your assigned workgroup pod number provided by the instructor

Exercise Procedure

Complete these steps:

Step 1 Assuming there is a Cisco IP Phone connected to your workgroup WGxS1 switch FastEthernet 0/1 port, establish a trust boundary on your workgroup WGxS1 switch FastEthernet 0/1 interface by setting your workgroup WGxS1 switch FastEthernet 0/1 interface to only trust cos from the Cisco IP Phone.

Note: Because there is not a Cisco IP Phone connected to your workgroup switch in the lab, all the incoming frames will have the CoS value set to the default port CoS of 0 by your workgroup switch.

Step 2 Display the default CoS-to-DSCP mapping on your workgroup WGxS1 switch.

```
Dscp-cos map:
dscp:   0  8 10 16 18 24 26 32 34 40 46 48 56
          ---------------------------------------
 cos:   0  1  1  2  2  3  3  4  4  5  5  6  7

Cos-dscp map:
 cos:   0  1   2   3   4  5   6   7
          ------------------------------
dscp:   0  8  16  24  32 40  48  56
```

What is the default mapping of CoS 5?

CoS 5 = DSCP _____

What is the DSCP per-hop behavior (PHB) represented by the default mapping of CoS 5 on the Catalyst 2950? _____

Step 3 Notice from Step 2 that the Catalyst 2950 default CoS-to-DSCP mapping does not map CoS 5 to DSCP 46.

On your WGxS1 switch, change the default CoS-to-DSCP mapping to map CoS 5 to DSCP 46 as the downstream devices are expecting voice (CoS 5) traffic to be marked with DSCP 46 (EF). All other CoS-to-DSCP mappings can remain at their default values.

Step 4 Display and verify the new CoS-to-DSCP mapping on your workgroup WGxS1 switch. Verify CoS 5 is now being mapped to DSCP 46.

```
Dscp-cos map:
dscp:   0  8 10 16 18 24 26 32 34 40 46 48 56
          ---------------------------------------
 cos:   0  1  1  2  2  3  3  4  4  5  5  6  7

Cos-dscp map:
 cos:   0  1   2   3   4  5   6   7
          ------------------------------
dscp:   0  8  16  24  32 46  48  56
```


Step 5 Assuming the Callgen routers (10.1.x.11 and 10.3.x.11) are the voice gateways described earlier and the Callgen routers are not marking voice traffic with DSCP 46 (Expedited Forwarding [EF]).

On your WGxS1 switch, create a policy map called **mark-callgen**, which uses two access lists and class-based marking to mark all the traffic from the Callgen routers (10.1.x.11 and 10.3.x.11) to DSCP 46 (EF). Name the two new service classes **callgen1** and **callgen2**.

Apply your class-based marking policy to the FastEtherent 0/1 interface on your WGxS1 switch in the inbound direction.

Step 6 Display and verify your IP access-list configuration.

```
Standard IP access list 1
 permit 10.1.x.11
Standard IP access list 2
 permit 10.3.x.11
```

Step 7 Display and verify your class-map configuration.

```
Class Map match-any class-default (id 0)
 Match any

Class Map match-all callgen1 (id 2)
 Match access-group 1

Class Map match-all callgen2 (id 3)
 Match access-group 2
```

Step 8 Display and verify the configuration of your class-based marking policy.

```
Policy Map mark-callgen

  class  callgen1
   set ip dscp 46

  class  callgen2
   set ip dscp 46
```

Step 9 Display and verify that your service policy is properly applied to the Fa0/1 interface in the inbound direction and also verify the trust setting and the default CoS value.

```
FastEthernet0/1
Attached policy-map for Ingress: mark-callgen
trust state: not trusted
trust mode: trust cos
COS override: dis
default COS: 0
pass-through: none
trust device: cisco-phone
```

Step 10 To verify that the Catalyst 2950 classification and marking configuration is working correctly, you will configure a class map on your workgroup router to count marked packets. These verification steps should be performed on both workgroup routers WGxR1 and WGxR2.

Step 11 On WGxR1, configure a new class map called **match-sw-ef** to match DSCP 46 traffic from the WGxS1 switch.

Step 12 On WGxR1, configure a second new class map called **match-sw-be** to match DSCP 0 traffic from the WGxS1 switch.

Step 13 Display and verify the two new class maps configured on WGxR1.

```
Class Map match-all match-sw-ef (id 4)
  Match ip dscp ef

Class Map match-all match-sw-be (id 5)
  Match ip dscp default

Class Map match-any class-default (id 0)
  Match any
```

Step 14 Configure a new policy map called **verify-mark** on WGxR1, which contains the two new traffic classes, match-sw-ef and match-sw-be.

Step 15 Display and verify the policy-map configuration on WGxR1.

```
Policy Map verify-mark
  Class match-sw-ef
  Class match-sw-be
```

Step 16 Disable the existing inbound service policy on the FastEthernet 0/0 interface of WGxR1.

Caution: Do not remove the actual policy map, which contains the real-time, mission-critical, interactive, bulk, and scavenger traffic classes, as you will use this policy map in later labs.

Step 17 Apply the new **verify-mark** policy to the FastEthernet 0/0 interface of WGxR1 in the inbound direction.

Step 18 Display and verify the policy map is being applied correctly on the FastEthernet 0/0 interface.

```
FastEthernet0/0

 Service-policy input: verify-mark

   Class-map: match-sw-ef (match-all)
     332 packets, 69798 bytes
     5 minute offered rate 0 bps
     Match: ip dscp ef

   Class-map: match-sw-be (match-all)
     10854 packets, 652290 bytes
     5 minute offered rate 26000 bps
     Match: ip dscp default

    Class-map: class-default (match-any)
     0 packets, 0 bytes
     5 minute offered rate 0 bps, drop rate 0 bps
     Match: any
```


How many packets have been matched for each of the traffic class?

Traffic Class	WGxR1	WGxR2
match-sw-ef		
match-sw-be		
default		

Step 19 Repeat Steps 11 through 18 for router WGxR2.

Step 20 Remove the inbound service policy from the FastEthernet 0/0 interface on both the WGxR1 and WGxR2 routers.

Step 21 Save your running configurations of the workgroup routers and the workgroup switch to the startup-config in NVRAM.

Exercise Verification

You have completed this exercise when you attain these results:

- You have properly configured a trust boundary on a the WGxS1 switch to only trust Cisco IP Phones
- You have properly configured the CoS-to-DSCP mapping on WGxS1 to map CoS 5 to DSCP 46
- You have correctly configured IP access lists and class-based marking on WGxS1 to mark traffic from the Callgen routers DSCP 46
- You have correctly verified the QoS markings from WGxS1 by configuring a QoS policy on WGxR1 and WGxR2 to count marked packets

Lab Exercise 5-1: Configuring Basic Queuing

Complete this lab exercise to practice what you learned in the related lesson.

Exercise Objective

In this exercise, you will configure WFQ on a router to improve QoS. After completing this exercise, you will be able to meet these objectives:

- Configure FIFO queuing on Cisco routers
- Configure WFQ on Cisco routers
- Use the proper show commands to monitor and verify the WFQ operation
- Use Cisco IOS monitoring commands and network connectivity tools (ping) to gather network response time data

Visual Objective

The figure illustrates what you will accomplish in this exercise.

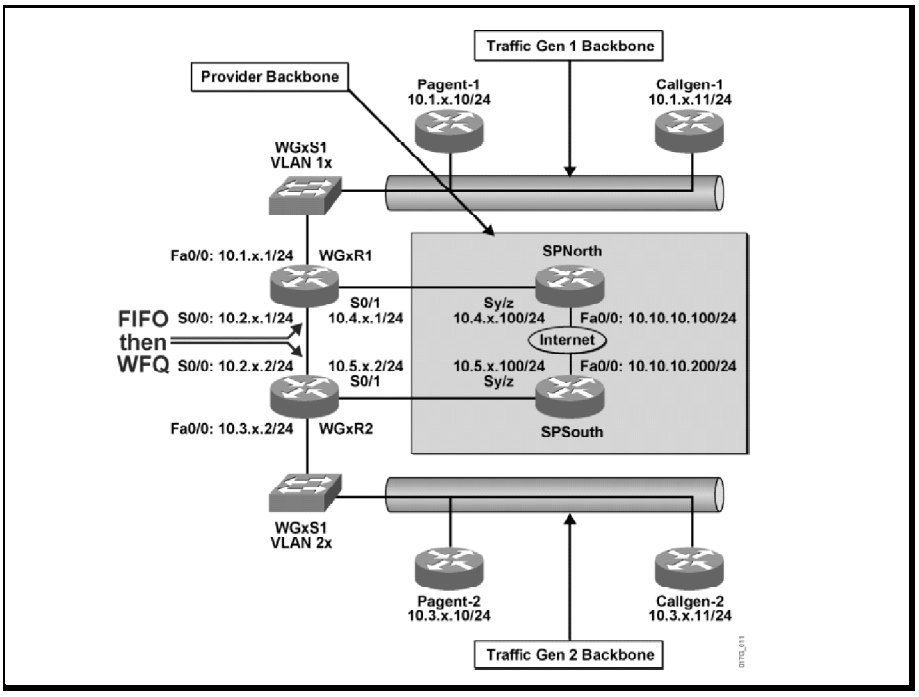

Once the proper classification and marking are implemented, the next step in completing the E-Commerce University QoS policy is to implement queuing mechanisms. Being an adventurous network engineer, and to gain a better understanding of the various queuing mechanisms, you decided to first explore two of the more basic queuing methods: FIFO and WFQ.

Required Resources

These are the resources and equipment required to complete this exercise:

- Lab topology configured for QOS
- Student workgroup consisting of two user-controlled Cisco 2610XM routers and one user-controlled Cisco 2950T-24 workgroup switch
- Classroom reference materials as follows:
  - QOS Student Guide
  - QOS Lab Guide
- Student pod workstation with Telnet or console access to workstation pod devices

Command List

The commands used in this exercise are described in the table here.

Configuring Basic Queuing Lab Commands

Command	Description
shutdown	To disable an interface
show interfaces [*type number***]**	To display statistics for all interfaces configured on the router or access server
interface *interface-id*	Enter interface configuration mode and the physical interface identification
[no] fair-queue [*congestive-discard-threshold* **[***dynamic-queues* **[***reservable-queues***]]]**	To enable WFQ for an interface
clear counters	To clear the interface counters
show queue *interface-name interface-number*	To display the contents of packets inside a queue for a particular interface
copy running-config startup-config	Save your entries in the configuration file

Job Aids

These job aids are available to help you complete the laboratory exercise:

- Your assigned workgroup pod number provided by the instructor

Exercise Procedure

Complete these steps:

Step 1 Knowing that the S0/0 interface on WGxR1 is 384 kbps, what will be the default queuing mechanism? _____

Step 2 Verify your answer by displaying the current queuing mechanism on the S0/0 interface of WGxR1.

```
Serial0/0 is up, line protocol is up
  Hardware is PowerQUICC Serial
  Description: to WGxR2
  Internet address is 10.2.1.1/24
  MTU 1500 bytes, BW 384 Kbit, DLY 20000 usec,
     reliability 255/255, txload 73/255, rxload 170/255
  Encapsulation HDLC, loopback not set
  Keepalive set (10 sec)
  Last input 00:00:00, output 00:00:00, output hang never
  Last clearing of "show interface" counters 00:03:57
  Input queue: 0/75/0/0 (size/max/drops/flushes); Total output
drops: 1358
  Queueing strategy: weighted fair
  Output queue: 0/1000/64/1358 (size/max
total/threshold/drops)
     Conversations  0/31/128 (active/max active/max total)
     Reserved Conversations 0/0 (allocated/max allocated)
     Available Bandwidth 288 kilobits/sec
  5 minute input rate 256000 bits/sec, 281 packets/sec
  5 minute output rate 111000 bits/sec, 177 packets/sec
     76914 packets input, 9858156 bytes, 0 no buffer
     Received 28 broadcasts, 0 runts, 0 giants, 0 throttles
     0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored, 0
abort
     44293 packets output, 2764269 bytes, 0 underruns
     0 output errors, 0 collisions, 0 interface resets
     0 output buffer failures, 0 output buffers swapped out
     0 carrier transitions
     DCD=up  DSR=up  DTR=up  RTS=up  CTS=up
```

Step 3 Disable WFQ on the serial 0/0 interface of the WGxR1 and WGxR2 routers.

Which queuing method is S0/0 using now after WFQ was disabled?

_____

Step 4 Display the current queuing mechanism on the workgroup WGxR1 router S0/0 interface to verify your answer.

```
Serial0/0 is up, line protocol is up
  Hardware is PowerQUICC Serial
  Description: to WGxR2
  Internet address is 10.2.8.1/24
  MTU 1500 bytes, BW 384 Kbit, DLY 20000 usec,
     reliability 255/255, txload 239/255, rxload 48/255
  Encapsulation PPP, LCP Open
  Open: CDPCP, IPCP, loopback not set
  Last input 00:00:04, output 00:00:00, output hang never
  Last clearing of "show interface" counters 21:01:49
  Input queue: 0/75/0/0 (size/max/drops/flushes); Total output
drops: 3041005
```



```
  Queueing strategy: fifo
  Output queue: 18/40 (size/max)
  5 minute input rate 73000 bits/sec, 207 packets/sec
  5 minute output rate 360000 bits/sec, 353 packets/sec
     19158035 packets input, 1358704729 bytes, 0 no buffer
     Received 0 broadcasts, 0 runts, 0 giants, 0 throttles
     1 input errors, 0 CRC, 1 frame, 0 overrun, 0 ignored, 0
abort
     25555103 packets output, 2871522779 bytes, 0 underruns
     0 output errors, 0 collisions, 5 interface resets
     0 output buffer failures, 0 output buffers swapped out
     0 carrier transitions
     DCD=up  DSR=up  DTR=up  RTS=up  CTS=up
```

Step 5 Perform an extended ping from WGxR1 to WGxR2 using a repeat count of 100 and a size of 160 bytes and record the results in the table at the end of this lab.

```
WGxR1#ping
Protocol [ip]:
Target IP address: 10.2.x.2
Repeat count [5]: 100
Datagram size [100]: 160
Timeout in seconds [2]:
Extended commands [n]:
Sweep range of sizes [n]:
Type escape sequence to abort.
Sending 100, 160-byte ICMP Echos to 10.2.1.2, timeout is 2
seconds:
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
Success rate is 100 percent (100/100), round-trip min/avg/max
= 8/15/168 ms
```

Step 6 Repeat the extended ping two more times and record your results in the table at the end of this lab.

Step 7 Repeat Steps 5 and 6, but ping from WGxR2 to WGxR1 serial 0/0 interface and record the response time results in the table at the end of this lab.

Step 8 What are some disadvantages of with using FIFO queuing for the voice and mission-critical traffic? _____

Step 9 Re-enable WFQ on the serial 0/0 interface of both WGxR1 and WGxR2. Use the default congestion discard threshold and the default maximum conversations.

Step 10 Display the interface statistics of the serial 0/0 interface.

```
Serial0/0 is up, line protocol is up
  Hardware is PowerQUICC Serial
  Description: to WGxR2
  Internet address is 10.2.8.1/24
  MTU 1500 bytes, BW 384 Kbit, DLY 20000 usec,
     reliability 255/255, txload 243/255, rxload 48/255
  Encapsulation PPP, LCP Open
  Open: CDPCP, IPCP, loopback not set
  Last input 00:00:01, output 00:00:00, output hang never
  Last clearing of "show interface" counters 19:29:36
  Input queue: 0/75/0/0 (size/max/drops/flushes); Total output
drops: 2678190
  Queueing strategy: weighted fair
```

```
  Output queue: 184/1000/64/2677828 (size/max
total/threshold/drops)
     Conversations  21/26/128 (active/max active/max total)
     Reserved Conversations 0/0 (allocated/max allocated)
     Available Bandwidth 288 kilobits/sec
  5 minute input rate 73000 bits/sec, 206 packets/sec
  5 minute output rate 366000 bits/sec, 347 packets/sec
     17972092 packets input, 1304053724 bytes, 0 no buffer
     Received 0 broadcasts, 0 runts, 0 giants, 0 throttl
22:46:55: %SYS-5-CONFIG_I: Configured from console by
consolees
     1 input errors, 0 CRC, 1 frame, 0 overrun, 0 ignored, 0
abort
     23553094 packets output, 2617787163 bytes, 0 underruns
     0 output errors, 0 collisions, 3 interface resets
     0 output buffer failures, 0 output buffers swapped out
     0 carrier transitions
     DCD=up  DSR=up  DTR=up  RTS=up  CTS=up
```

Step 11 What is the default Congestion Discard Threshold (CDT)?

_____

Step 12 What is the default number of maximum conversations on the S0/0 interface?

_____

Step 13 Change the maximum number of conversations on the S0/0 interface of WGxR1 and WGxR2 to 1024.

Step 14 Display the interface statistics of the serial 0/0 interface and verify that the maximum conversations have been corrected changed.

```
Serial0/0 is up, line protocol is up
  Hardware is PowerQUICC Serial
  Description: to WGxR2
  Internet address is 10.2.8.1/24
  MTU 1500 bytes, BW 384 Kbit, DLY 20000 usec,
     reliability 255/255, txload 243/255, rxload 48/255
  Encapsulation PPP, LCP Open
  Open: CDPCP, IPCP, loopback not set
  Last input 00:00:06, output 00:00:00, output hang never
  Last clearing of "show interface" counters 19:32:41
  Input queue: 0/75/0/0 (size/max/drops/flushes); Total output
drops: 2689808
  Queueing strategy: weighted fair
  Output queue: 101/1000/64/2689446 (size/max
total/threshold/drops)
     Conversations  15/29/1024 (active/max active/max total)
     Reserved Conversations 0/0 (allocated/max allocated)
     Available Bandwidth 288 kilobits/sec
  5 minute input rate 73000 bits/sec, 211 packets/sec
  5 minute output rate 367000 bits/sec, 361 packets/sec
     18011990 packets input, 1305889643 bytes, 0 no buffer
     Received 0 broadcasts, 0 runts, 0 giants, 0 thrott
22:50:00: %SYS-5-CONFIG_I: Configured from console by
consoleles
     1 input errors, 0 CRC, 1 frame, 0 overrun, 0 ignored, 0
abort
     23620090 packets output, 2626292704 bytes, 0 underruns
     0 output errors, 0 collisions, 4 interface resets
     0 output buffer failures, 0 output buffers swapped out
     0 carrier transitions
```



```
     DCD=up  DSR=up  DTR=up  RTS=up  CTS=up
```

Step 15 What is a benefit of increasing the maximum conversations?

_____

Step 16 Clear the interface counters on your WGxR1 and WGxR2 routers.

Step 17 Examine the **active** dynamic queues set up by WFQ. You may only see flows on only one of your workgroup routers because of the traffic flow patterns between the Pagent-1 and Pagent-2 router.

```
 Input queue: 0/75/0/0 (size/max/drops/flushes); Total output
drops: 3227345
  Queueing strategy: weighted fair
  Output queue: 34/1000/64/3226983 (size/max
total/threshold/drops)
     Conversations  5/30/1024 (active/max active/max total)
     Reserved Conversations 0/0 (allocated/max allocated)
     Available Bandwidth 288 kilobits/sec

  (depth/weight/total drops/no-buffer drops/interleaves)
1/32384/26/0/0
  Conversation 85, linktype: ip, length: 44
  source: 10.1.1.10, destination: 10.3.1.10, id: 0xFF4C, ttl:
254,
  TOS: 0 prot: 6, source port 1063, destination port 1521

  (depth/weight/total drops/no-buffer drops/interleaves)
27/32384/416/0/0
  Conversation 97, linktype: ip, length: 97
  source: 10.1.1.10, destination: 10.3.1.10, id: 0xE642, ttl:
127,
  TOS: 0 prot: 6, source port 1102, destination port 1494

  (depth/weight/total drops/no-buffer drops/interleaves)
8/32384/350/0/0
  Conversation 57, linktype: ip, length: 279
  source: 10.1.1.10, destination: 10.3.1.10, id: 0x908A, ttl:
126,
  TOS: 0 prot: 17, source port 49590, destination port 49602
```

How many packets have been dropped? _____

What is the number of active conservations? _____

What is the number of maximum active conservations? _____

What is the maximum number of total conservations? _____

What is the available bandwidth and how is it calculated?

_____

What is the significance of the weight in these conversations?

_____

What factor or factors can influence the weight? _____

Step 18 Perform an extended ping from WGxR1 to WGxR2 using a repeat count of 100 and a size of 160 bytes and record the results in the table at the end of this lab.

```
WGxR1#ping
Protocol [ip]:
Target IP address: 10.2.x.2
Repeat count [5]: 100
Datagram size [100]: 160
Timeout in seconds [2]:
Extended commands [n]:
Sweep range of sizes [n]:
Type escape sequence to abort.
Sending 100, 160-byte ICMP Echos to 10.2.1.2, timeout is 2
seconds:
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
Success rate is 100 percent (100/100), round-trip min/avg/max
= 12/57/192 ms
```

Step 19 Repeat the extended ping two more times and record your results in the table at the end of this lab.

Step 20 Repeat Steps 18 and 19, but ping from the WGxR2 to WGxR1 serial 0/0 interface and record the response time results in the table at the end of this lab.

Did the ping response time improve compared to when FIFO was the queuing method? _____

Step 21 Save your running configurations of the workgroup routers and the workgroup switch to the startup-config in NVRAM.

WGxR1 RTT Measurement Worksheet

Network Situation	Ping 1 RTT (min/ave/max)	Ping 2 RTT (min/ave/max)	Ping 3 RTT (min/ave/max)
Before Pagnet / Callgen traffic (from 2-2: Baseline QoS Measurement)			
After Pagnet/Callgen traffic (from 2-2: Baseline QoS Measurement)			
After enabling Auto-QoS (from 3-1: Configuring QoS with AutoQoS)			
FIFO (from 5-1: Configuring Basic Queuing)			
WFQ (from 5-1: Configuring Basic Queuing)			

WGxR2 RTT Measurement Worksheet

Network Situation	Ping 1 RTT (min/ave/max)	Ping 2 RTT (min/ave/max)	Ping 3 RTT (min/ave/max)
Before Pagnet / Callgen traffic (from 2-2: Baseline QoS Measurement)			
After Pagnet/Callgen traffic (from 2-2: Baseline QoS Measurement)			
After enabling Auto-QoS (from 3-1: Configuring QoS with AutoQoS)			
FIFO (from 5-1: Configuring Basic Queuing)			
WFQ (from 5-1: Configuring Basic Queuing)			

Exercise Verification

You have completed this exercise when you attain these results:

- You have successfully configured FIFO queuing on WGxR1 and WGxR2
- You have successfully configured WFQ on WGxR1 and WGxR2
- You have examined WFQ router parameters and monitoring command output and successfully answered the questions contained in this lab exercise
- You have compared the round-trip response times for packets traversing a serial interface using FIFO queuing to an interface using WFQ

Lab Exercise 5-2: Configuring LLQ

Complete this lab exercise to practice what you learned in the related lesson.

Exercise Objective

In this exercise, you will configure low-latency queuing (LLQ) on a router to improve QoS. After completing this exercise, you will be able to meet these objectives:

- Configure LLQ on a Cisco router to provide bandwidth guarantees
- Use the proper show commands to monitor and verify the LLQ operation
- Use Cisco IOS monitoring commands and network connectivity tools (ping) to gather network response time data

Visual Objective

The figure illustrates what you will accomplish in this exercise.

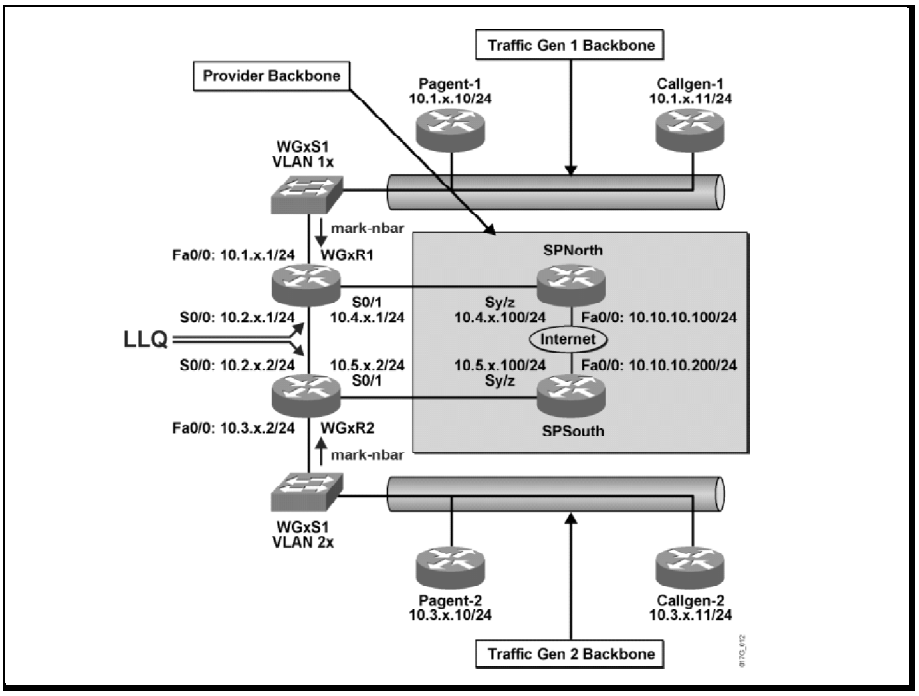

In Lab Exercise 5-1, you examined both FIFO queuing and WFQ.

FIFO is a first-come-first-serve queuing strategy. As such, FIFO does not classify traffic into different flows, or provide differentiated treatment to packets. FIFO does not fairly allocate bandwidth among multiple flows. Some flows may receive more bandwidth because they contain larger packets or they send more packets (aggressive flows). FIFO does not give priority to voice traffic or mission-critical traffic, as it cannot differentiate between packets from different flows. Therefore, if FIFO is implemented on the serial WAN links, the configuration will not meet the E-Commerce University requirements for the voice and Oracle (SQL) traffic.

After FIFO queuing, WFQ was enabled on the workgroup routers. With WFQ, packets are automatically classified into a particular dynamic queue based on information contained within the protocol headers. WFQ automatically assigns a weight to the packets. The weight is based on the IP precedence value and the size of the packets. Higher IP precedence and smaller size packets will receive better serviced than lower IP precedence and larger size packets. However, WFQ does not provide any hard bandwidth guarantees to voice traffic or to mission-critical traffic. From the traffic measurement results (round-trip time [RTT] observations), it was determined that WFQ provided too much latency for EF traffic types. Therefore, in order to improve voice quality and the response time of mission-critical traffic, LLQ must be implemented on the workgroup routers.

Required Resources

These are the resources and equipment required to complete this exercise:

- Lab topology configured for QOS
- Student workgroup consisting of two user-controlled Cisco 2610XM routers and one user-controlled Cisco 2950T-24 workgroup switch
- Classroom reference materials as follows:
  - QOS Student Guide
  - QOS Lab Guide
- Student pod workstation with Telnet or console access to workstation pod devices

Command List

The commands used in this exercise are described in the table here.

Configuring Queuing Hybrids Lab Commands

Command	Description
class-map *class-map-name*	Create a class map to be used for matching packets to the class whose name you specify
match {access-group *acl-index* | **access-group name** *acl-name* | **ip dscp** *dscp-list*}	Define the match criteria to classify traffic
match [ip] dscp *dscp-value*	To identify a specific IP DSCP value as a match criterion
match protocol *protocol-name*	To configure the match criteria for a class map on the basis of the specified protocol
policy-map *policy-map-name*	Create or modify a policy map that can be attached to multiple interfaces
class {*class-name* | **class-default**}	To specify the name of the class whose policy you want to create or change or to specify the default class
set ip dscp *new-dscp*	Classify IP traffic by setting a DSCP value
priority{*bandwidth-kbps* | percent percentage} [burst]	To give priority to a class of traffic belonging to a policy map
bandwidth {*bandwidth-kbps* | remaining percent percentage | **percent** percentage}	To specify or modify the bandwidth allocated for a class belonging to a policy map
service-policy input *policy-map-name*	Apply a policy map defined by the policy-map command to the input of a particular interface
access-list *access-list-number* {**deny** | **permit**} **tcp src** *src-wildcard* **[operator** [*port*]**] dest** *dest-wildcard* **[operator** [*port*]**]**	To define an extended IP access list for Transmission Control Protocol (TCP) based traffic
show access-lists [*access-list-number* | *access-list-name*]	To display the contents of current access lists
show class-map [*class-map-name*]	To display all class maps and their matching criteria
show policy-map [*policy-map*]	To display the configuration of all classes for a specified service policy map or all classes for all existing policy maps
show policy-map interface *interface-name* **[input** | **output] [class** *class-map-name*]	To display the packet statistics of all classes that are configured for all service policies either on the specified interface or subinterface
copy running-config startup-config	Save your entries in the configuration file

Job Aids

These job aids are available to help you complete the laboratory exercise:

- Your assigned workgroup pod number provided by the instructor

Exercise Procedure

Complete these steps:

Step 1 Connect to your WGxR1 router and verify that the QoS policy map named **mark-nbar** is still configured. Apply the existing **mark-nbar** policy map to the FastEthernet 0/0 interface in the inbound direction. If you do not have the mark-nbar policy, it is shown below:

```
class-map match-any real-time
 match protocol rtp
 match protocol icmp
 match access-group name VoIP-RTCP
class-map match-any mission-critical
 match protocol sqlnet
 match access-group name Voice-Control
class-map interactive
 match protocol citrix
class-map bulk
 match protocol ftp
class-map match-any scavenger
 match protocol kazaa2
 match protocol napster
!
policy-map mark-nbar
 class real-time
  set ip dscp ef
 class mission-critical
  set ip dscp af31
 class interactive
  set ip dscp af21
 class bulk
  set ip dscp af11
 class scavenger
  set ip dscp cs1
```

Step 2 At this point, what is the expected effect on traffic flow out the low speed 384-kbps serial link (S0/0) when the mark-nbar service-policy is applied on the FastEthernet 0/0 interface of WGxR1 router? Explain.

_____

Step 3 Configure five new class maps on your workgroup WGxR1 router as described on the following table. When matching ICMP traffic, use an extended ACL instead of NBAR as NBAR only matched transit traffic.

Class Name (class map name)	Match Criteria
ef-traffic	EF, ICMP echo and reply
af31-traffic	AF 31
af21-traffic	AF 21
af11-traffic	AF 11
cs1-traffic	CS 1

Step 4 Display and verify your extended IP access list for matching the ping traffic.

```
Extended IP access list 100
  10 permit icmp any any echo-reply (100 matches)
  20 permit icmp any any echo
```

Step 5 Display and verify the five new class maps configured on your workgroup WGxR1 router.

```
 Class Map match-any ef-traffic (id 10)
   Match   dscp ef
   Match access-group 100

 Class Map match-all af21-traffic (id 12)
   Match   dscp af21

 Class Map match-all af31-traffic (id 11)
   Match   dscp af31

 Class Map match-all af11-traffic (id 13)
   Match   dscp af11

 Class Map match-all cs1-traffic (id 14)
   Match   dscp cs1
```


Step 6 Configure a new policy map called **llq-policy** on your workgroup WGxR1 router where each traffic class gets the following bandwidth guarantee.

Traffic Class	Bandwidth Guarantee
ef-traffic	168 kbps maximum
af31-traffic	40% of remaining bandwidth mininum
af21-traffic	20% of remaining bandwidth mininum
af11-traffic	13% of remaining bandwidth mininum
cs1-traffic	2% of remaining bandwidth mininum
class-default	25% of remaining bandwidth mininum

Step 7 Display and verify the llq-policy policy map on your workgroup WGxR1 router.

```
Policy Map llq-policy
  Class ef-traffic
    Strict Priority
    Bandwidth 168 (kbps) Burst 4200 (Bytes)
  Class af31-traffic
    Bandwidth remaining 40 (%) Max Threshold 64 (packets)
  Class af21-traffic
    Bandwidth remaining 20 (%) Max Threshold 64 (packets)
  Class af11-traffic
    Bandwidth remaining 13 (%) Max Threshold 64 (packets)
  Class cs1-traffic
    Bandwidth remaining 2 (%) Max Threshold 64 (packets)
  Class class-default
    Bandwidth remaining 25 (%) Max Threshold 64 (packets)
```

Step 8 Apply the new llq-policy policy map on your workgroup WGxR1 router S0/0 interface in the outbound direction.

Why must the policy be applied in the outbound direction and not the inbound direction? Explain. _____

Step 9 Repeat the above LLQ configuration, Steps 1 through 8, for the WGxR2 router.

Step 10 Wait for the interface counters to accumulate traffic statistics for at least 1 minute.

Step 11 Display and verify the outbound service policy on your workgroup WGxR1 router S0/0 interface.

```
Serial0/0

  Service-policy output: llq-policy

    Class-map: ef-traffic (match-any)
      332 packets, 66452 bytes
      5 minute offered rate 2000 bps, drop rate 0 bps
      Match:  dscp ef
        332 packets, 66452 bytes
        5 minute rate 2000 bps
      Match: access-group 100
        0 packets, 0 bytes
        5 minute rate 0 bps
      Queueing
        Strict Priority
        Output Queue: Conversation 136
        Bandwidth 168 (kbps) Burst 4200 (Bytes)
        (pkts matched/bytes matched) 332/66452
        (total drops/bytes drops) 0/0

    Class-map: af31-traffic (match-all)
      1067 packets, 149314 bytes
      5 minute offered rate 4000 bps, drop rate 0 bps
      Match:  dscp af31
      Queueing
        Output Queue: Conversation 137
        Bandwidth remaining 40 (%) Max Threshold 64 (packets)
        (pkts matched/bytes matched) 1248/172788
        (depth/total drops/no-buffer drops) 19/0/0

    Class-map: af21-traffic (match-all)
```



```
    1452 packets, 122100 bytes
    5 minute offered rate 4000 bps, drop rate 0 bps
    Match:  dscp af21
    Queueing
      Output Queue: Conversation 138
      Bandwidth remaining 20 (%) Max Threshold 64 (packets)
      (pkts matched/bytes matched) 1672/140600
      (depth/total drops/no-buffer drops) 33/0/0

  Class-map: af11-traffic (match-all)
    510 packets, 32499 bytes
    5 minute offered rate 2000 bps, drop rate 0 bps
    Match:  dscp af11
    Queueing
      Output Queue: Conversation 139
      Bandwidth remaining 13 (%) Max Threshold 64 (packets)
      (pkts matched/bytes matched) 618/39384
      (depth/total drops/no-buffer drops) 0/0/0

  Class-map: cs1-traffic (match-all)
    891 packets, 197225 bytes
    5 minute offered rate 6000 bps, drop rate 4000 bps
    Match:  dscp cs1
    Queueing
      Output Queue: Conversation 140
      Bandwidth remaining 2 (%) Max Threshold 64 (packets)
      (pkts matched/bytes matched) 1089/241224
      (depth/total drops/no-buffer drops) 64/725/0

  Class-map: class-default (match-any)
    2638 packets, 442008 bytes
    5 minute offered rate 10000 bps, drop rate 4000 bps
    Match: any
    Queueing
      Output Queue: Conversation 141
      Bandwidth remaining 25 (%) Max Threshold 64 (packets)
      (pkts matched/bytes matched) 3261/539295
      (depth/total drops/no-buffer drops) 45/1350/0
```

Which traffic class or classes, if any, still have drops? _____

Step 12 Display and verify the outbound service policy on your workgroup WGxR2 router S0/0 interface.

Which traffic class or classes, if any, still have drops? _____

Step 13 Perform an extended ping from WGxR1 to WGxR2 using a repeat count of 100 and a size of 160 bytes and record the results in the table at the end of this lab.

```
WGxR1#ping
Protocol [ip]:
Target IP address: 10.2.x.2
Repeat count [5]: 100
Datagram size [100]: 160
Timeout in seconds [2]:
Extended commands [n]:
Sweep range of sizes [n]:
Type escape sequence to abort.
Sending 100, 160-byte ICMP Echos to 10.2.1.2, timeout is 2
seconds:
```

```
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
Success rate is 100 percent (100/100), round-trip min/avg/max
= 8/17/68 ms
```

Step 14 Repeat the extended ping two more times and record your results in the table at the end of this lab.

Step 15 Repeat Steps 13 and 14, but ping from the WGxR2 to WGxR1 serial 0/0 interface and record the response time results in the table at the end of this lab.

Step 16 Compare the ping results to the results from the previous laboratory exercises (2-2: Baseline QoS Measurement, 3-1: Configuring QoS with AutoQoS).

Comparing all the results, which QoS mechanism provided the best response time for VoIP packets? _____

Step 17 Save your running configurations of the workgroup routers and the workgroup switch to the startup-config in NVRAM.

WGxR1 RTT Measurement Worksheet

Network Situation	Ping 1 RTT (min/ave/max)	Ping 2 RTT (min/ave/max)	Ping 3 RTT (min/ave/max)
Before Pagnet / Callgen traffic (from 2-2: Baseline QoS Measurement)			
After Pagnet/Callgen traffic (from 2-2: Baseline QoS Measurement)			
After enabling Auto-QoS (from 3-1: Configuring QoS with AutoQoS)			
FIFO (from 5-1: Configuring Basic Queuing)			
WFQ (from 5-1: Configuring Basic Queuing)			
LLQ results from this lab			

WGxR2 RTT Measurement Worksheet

Network Situation	Ping 1 RTT (min/ave/max)	Ping 2 RTT (min/ave/max)	Ping 3 RTT (min/ave/max)
Before Pagnet / Callgen traffic (from 2-2: Baseline QoS Measurement)			
After Pagnet/Callgen traffic (from 2-2: Baseline QoS Measurement)			
After enabling Auto-QoS (from 3-1: Configuring QoS with AutoQoS)			
FIFO (from 5-1: Configuring Basic Queuing)			
WFQ (from 5-1: Configuring Basic Queuing)			
LLQ results from this lab			

Exercise Verification

You have completed this exercise when you attain these results:

- LLQ is properly configured on WGxR1 and WGxR2
- The extended pings produced lower round-trip response time with LLQ than WFQ

Lab Exercise 5-3: Queuing on Catalyst Switches

Complete this lab exercise to practice what you learned in the related lesson.

Exercise Objective

In this exercise, you will configure the two queuing methods available on the Catalyst 2950 switch. After completing this exercise, you will be able to meet these objectives:

- Examine queuing configurations resulting from the application of AutoQoS on the Catalyst 2950 switch
- Configure and monitor CoS-to-queue mapping on the Catalyst 2950 switch
- Configure and monitor WRR queuing on the Catalyst 2950 switch

Visual Objective

The figure illustrates what you will accomplish in this exercise.

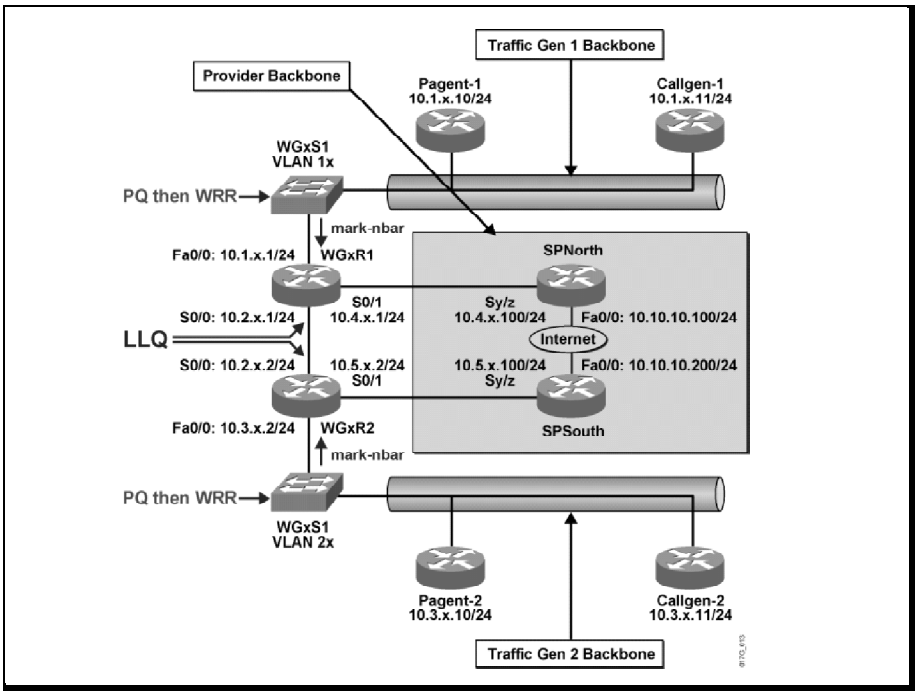

After reviewing the AutoQoS output on the Catalyst 2950 switch, the E-Commerce University IT staff has a few questions about the AutoQoS generated configurations. After meeting with the IT staff and hopefully being able to answer all their questions, the IT staff has decided to change the default Catalyst 2950 queuing from PQ to weighted round robin (WRR) with an expedite queue so that only voice (CoS 5) frames will receive strict priority while giving frames with CoS 0 to 4 a lower weight than frames with CoS 6 to 7.

Required Resources

These are the resources and equipment required to complete this exercise:

- Lab topology configured for QOS
- Student workgroup consisting of two user-controlled Cisco 2610XM routers and one user-controlled Cisco 2950T-24 workgroup switch
- Classroom reference materials as follows:
  - QOS Student Guide
  - QOS Lab Guide
- Student pod workstation with Telnet or console access to workstation pod devices

Command List

The commands used in this exercise are described in the table here.

Queuing on Catalyst Switches Lab Commands

Command	Description
show wrr-queue cos-map	Displays the mapping of the CoS priority queues (PQs)
wrr-queue bandwidth *weight1...weight4*	Assigns weighted round-robin (WRR) weights to the four CoS priority queues
no wrr-queue cos-map	Sets the CoS map to default setting
wrr-queue cos-map *quid cos1...cosn*	Assign CoS values to the CoS priority queues. quid: The queue id of the CoS priority queue cos1...cosn: The CoS values that are mapped to the queue id
mls qos map cos-dscp *dscp1...dscp8*	Defines the CoS-to-DSCP map
mls qos trust [cos | device cisco-phone | dscp]	Configure the port trust state
show wrr-queue bandwidth	Displays the WRR bandwidth allocation for the four CoS priority queues
copy running-config startup-config	Save your entries in the configuration file

Job Aids

These job aids are available to help you complete the laboratory exercise:

- Your assigned workgroup pod number provided by the instructor

Exercise Procedure

Complete these steps:

Step 1 In this lab, you will first examine the default PQ on the Catalyst 2950 workgroup switch. Then you will configure and examine WRR on the Catalyst 2950 workgroup switch.

What is the default queuing mechanism on the Catalyst 2950 switch?

_____

Step 2 Connect to your WGxS1 switch and display the CoS to Queue mapping. Your output should look similar to the following:

```
CoS Value     :  0  1  2  3  4  5  6  7
Priority Queue:  1  1  2  2  3  3  4  4
```

How many output queues per interface are available for the Catalyst 2950? _____

Which queue has the highest priority? _____

Step 3 In Lab Exercise 3-1, Configuring QoS with AutoQoS, the AutoQoS for VoIP feature was enabled on the Catalyst 2950 switch in your workgroup. By default, AutoQoS enables queuing as shown in the following configuration:

```
WGxS1#show auto qos

Initial configuration applied by AutoQoS:
wrr-queue bandwidth 20 1 80 0
no wrr-queue cos-map
wrr-queue cos-map 1 0 1 2 4
wrr-queue cos-map 3 3 6 7
wrr-queue cos-map 4 5
mls qos map cos-dscp 0 8 16 26 32 46 48 56
!
interface FastEthernet0/1
 mls qos trust cos
```

Step 4 What type of queuing is enabled using the above configuration? _____

Step 5 From the "wrr-queue bandwidth 20 1 80 0" command, why is the weight for queue #4 equal to "0"? _____

Step 6 What do the following commands accomplish?

```
wrr-queue cos-map 1 0 1 2 4
wrr-queue cos-map 3 3 6 7
wrr-queue cos-map 4 5
```

_____

Step 7 On the WGxS1 switch, implement WRR queuing as follows:

Configure a PQ that services all frames marked CoS 5.

Configure WRR queue 3 to service all frames marked CoS 6 and CoS 7 and guarantee a bandwidth of 70.

Configure WRR queue 1 to service all frames marked CoS 0 – CoS 4 and guarantee a bandwidth of 30.

WRR queue 2 is unused and should be configured as such.

Step 8 Display and verify the CoS to queue mapping has been correctly configured and the frames marked CoS 5 are services by the PQ.

```
CoS Value     :  0  1  2  3  4  5  6  7
Priority Queue:  1  1  1  1  1  4  3  3
```

Step 9 Display and verify the weight setting for each queue.

```
WRR Queue:    1    2    3    4
bandwidth:    30   1    70   0
```

Step 10 Configure the FastEthernet 0/1 interface on WGxS1 to trust the CoS marking.

Step 11 Ensure that DSCP EF is marked 46 for incoming frames marked CoS 5.

Step 12 Save your running configurations of the workgroup routers and the workgroup switch to the startup-config in NVRAM.

Exercise Verification

You have completed this exercise when you attain these results:

- Correctly described the effects of the resulting AutoQoS configurations on the Catalyst 2950 switch.
- Correctly configure WRR queuing on the Catalyst 2950 switch and enable queue #4 as the expedite queue and map CoS 5 traffic into the expedite queue

Lab Exercise 6-1: Configuring DSCP-Based WRED

Complete this lab exercise to practice what you learned in the related lesson.

Exercise Objective

In this exercise, you will build a weighted random early detection (WRED) traffic profile given a set of parameters and configure DSCP-based WRED with explicit congestion notification (ECN) support to match that traffic profile. After completing this exercise, you will be able to meet these objectives:

- Configure DSCP-Based CB-WRED
- Configure DSCP-Based CB-WRED with ECN
- Monitor DSCP-Based CB-WRED operations

Visual Objective

The figure illustrates what you will accomplish in this exercise.

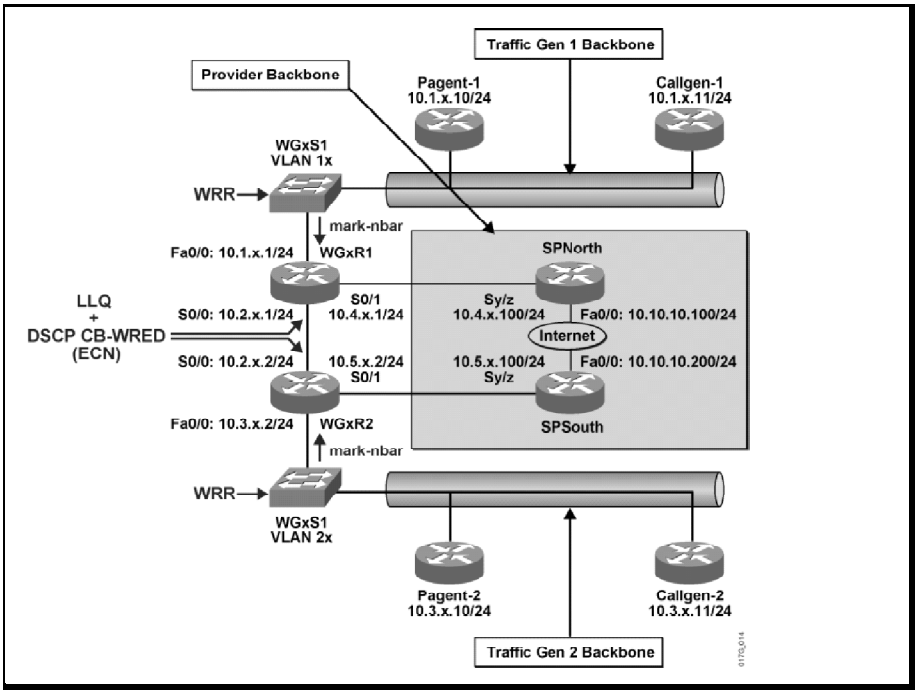

After LLQ was successfully implemented, the voice quality definitely improved. But after monitoring the link utilization on the low-speed 384-kbps link for a week, it has been determined that the average link utilization is low and must be improved.

You recall from the Cisco IOS documentation that global synchronization occurs as waves of congestion crest only to be followed by troughs during which the transmission link is not fully utilized. Global synchronization of TCP hosts can occur because packets are dropped all at once. Global synchronization manifests when multiple TCP hosts reduce their transmission rates in response to packet dropping, and once congestion is reduced, the TCP hosts again

increase their transmission rates. The most important point is that the waves of transmission known as global synchronization result in significant link under-utilization.

In order to reduce TCP global synchronization to improve link utilization, CB-WRED is required to randomly drop packets before the software queue is full. In addition, DSCP-based CB-WRED allows different WRED (drop) profiles for different DSCP values.

Required Resources

These are the resources and equipment required to complete this exercise:

- Lab topology configured for QOS
- Student workgroup consisting of two user-controlled Cisco 2610XM routers and one user-controlled Cisco 2950T-24 workgroup switch
- Classroom reference materials as follows:
  - QOS Student Guide
  - QOS Lab Guide
- Student pod workstation with Telnet or console access to workstation pod devices

Command List

The commands used in this exercise are described in the table here.

DSCP-Based WRED with ECN Lab Commands

Command	Description
policy-map *policy-map-name*	To create or modify a policy map that can be attached to one or more interfaces
class {*class-name* | **class-default**}	To specify the name of the class whose policy you want to create or change or to specify the default class
random-detect dscp *dscpvalue min-threshold max-threshold* **[***mark-probability-denominator***]**	To change the minimum and maximum packet thresholds for the DSCP value
random-detect [*dscp-based* | *prec-based*]	To enable WRED or distributed WRED (DWRED)
random-detect ecn	To enable ECN
show policy-map [*policy-map***]**	To display the configuration of all classes for a specified service policy map or all classes for all existing policy maps
show policy-map interface *interface-name* **[input | output] [class** *class-map-name***]**	To display the packet statistics of all classes that are configured for all service policies either on the specified interface or subinterface
show class-map [*class-map-name***]**	Display QoS class maps
copy running-config startup-config	Save your entries in the configuration file

Job Aids

These job aids are available to help you complete the laboratory exercise:

- Your assigned workgroup pod number provided by the instructor

Exercise Procedure

Complete these steps:

Step 1 Modify the existing llq-policy map on the workgroup WGxR1 router and enable DSCP-Based WRED for the af11-traffic, af21-traffic, af31-traffic, cs1-traffic, and the class-default traffic classes with the following drop thresholds and drop probabilities:

PHB	Minimum Threshold	Maximum Threshold	Mark Probability
af11	26	40	1/10
af21	30	40	1/10
af31	34	40	1/10
cs1	22	40	1/10
Default (BE)	20	40	1/10

Complete the graph of the traffic profile below for all 5 WRED classes. Be sure to indicate each class and the mark probability denominator.

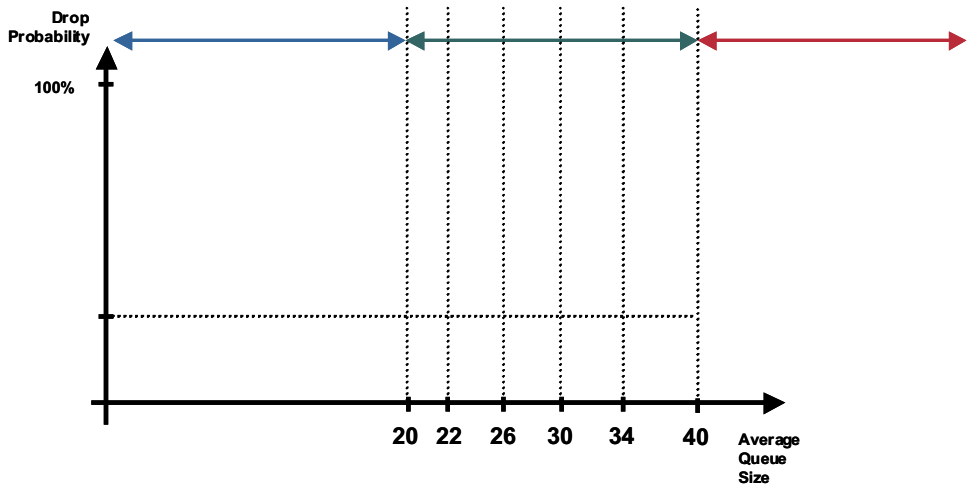

Based on the previous WRED profiles, which traffic class will start dropping packets first?

_____

What does a mark probability of 1/10 means?

_____

Why would you not implement WRED for the ef-traffic class?

_____

Step 2 Display the llq-policy policy map and to verify the WRED configurations.

```
 Policy Map llq-policy
    Class ef-traffic
      Strict Priority
      Bandwidth 168 (kbps) Burst 4200 (Bytes)
    Class af31-traffic
      Bandwidth remaining 40 (%)
            exponential weight 9
            dscp     min-threshold    max-threshold    mark-
probablity
            ------------------------------------------------
-------
            af11      -                -                1/10
            af12      -                -                1/10
            af13      -                -                1/10
            af21      -                -                1/10
            af22      -                -                1/10
            af23      -                -                1/10
            af31      34               40               1/10
            af32      -                -                1/10

            [output omitted]
```

```
    Class af21-traffic
      Bandwidth remaining 20 (%)
            exponential weight 9
            dscp     min-threshold    max-threshold    mark-
probablity
            ------------------------------------------------
-------
            af11      -                -                1/10
            af12      -                -                1/10
            af13      -                -                1/10
            af21      30               40               1/10
            af22      -                -                1/10
            af23      -                -                1/10

            [output omitted]

    Class af11-traffic
      Bandwidth remaining 13 (%)
            exponential weight 9
            dscp     min-threshold    max-threshold    mark-
probablity
            ------------------------------------------------
-------
            af11      26               40               1/10
            af12      -                -                1/10
            af13      -                -                1/10

            [output omitted]

    Class cs1-traffic
      Bandwidth remaining 2 (%)
            exponential weight 9
            dscp     min-threshold    max-threshold    mark-
probablity
            ------------------------------------------------
-------
            af11      -                -                1/10
            af12      -                -                1/10
            af13      -                -                1/10
            af21      -                -                1/10
            af22      -                -                1/10
            af23      -                -                1/10
            af31      -                -                1/10
            af32      -                -                1/10
            af33      -                -                1/10
            af41      -                -                1/10
            af42      -                -                1/10
            af43      -                -                1/10
            cs1       22               40               1/10
            cs2       -                -                1/10

            [output omitted]

    Class class-default
      Bandwidth remaining 25 (%)
            exponential weight 9
            dscp     min-threshold    max-threshold    mark-
probablity
            ------------------------------------------------
-------
```



```
            [output omitted]

            ef          -                   -                   1/10
            rsvp        -                   -                   1/10
            default     20                  40                  1/10
```

Step 3 What is the default exponential weight constant? _____

Are all the drop thresholds and drop probability set correctly? _____

Step 4 Clear the counters on all interfaces on the WGxR1 router.

Step 5 Wait for the interface counters to accumulate traffic statistics for at least 1 minute.

Step 6 Display the output service policy on the S0/0 interface.

```
Serial0/0
  Service-policy output: llq-policy
    Class-map: ef-traffic (match-any)
      2998 packets, 599437 bytes
      5 minute offered rate 4000 bps, drop rate 0 bps
      Match:  dscp ef
        2998 packets, 599437 bytes
        5 minute rate 4000 bps
      Match: access-group 100
        0 packets, 0 bytes
        5 minute rate 0 bps
      Queueing
        Strict Priority
        Output Queue: Conversation 136
        Bandwidth 168 (kbps) Burst 4200 (Bytes)
        (pkts matched/bytes matched) 6/1224
        (total drops/bytes drops) 0/0
    Class-map: af31-traffic (match-all)
      24740 packets, 1097866 bytes
      5 minute offered rate 13000 bps, drop rate 0 bps
      Match:  dscp af31
      Queueing
        Output Queue: Conversation 137
        Bandwidth remaining 40 (%)
        (pkts matched/bytes matched) 278/12232
        (depth/total drops/no-buffer drops) 0/0/0
         exponential weight: 9
         mean queue depth: 0
   dscp    Transmitted      Random drop      Tail drop
Minimum Maximum  Mark             pkts/bytes       pkts/bytes
pkts/bytes    thresh  thresh  prob
   af11       0/0               0/0              0/0
32      40  1/10
   af12       0/0               0/0              0/0
28      40  1/10
   af13       0/0               0/0              0/0
24      40  1/10
   af21       0/0               0/0              0/0
32      40  1/10
   af22       0/0               0/0              0/0
28      40  1/10
   af23       0/0               0/0              0/0
24      40  1/10
   af31   24817/1101254         0/0              0/0
32      40  1/10
```

```
[output omitted]

    Class-map: af21-traffic (match-all)
      27575 packets, 1218592 bytes
      5 minute offered rate 16000 bps, drop rate 0 bps
      Match:  dscp af21
      Queueing
        Output Queue: Conversation 138
        Bandwidth remaining 20 (%)
        (pkts matched/bytes matched) 387/17032
        (depth/total drops/no-buffer drops) 0/0/0
         exponential weight: 9
         mean queue depth: 0
   dscp    Transmitted      Random drop      Tail drop
Minimum Maximum  Mak
              pkts/bytes       pkts/bytes       pkts/bytes
thresh  thresh  prob
   af11       0/0              0/0              0/0
32      40  1/10
   af12       0/0              0/0              0/0
28      40  1/10
   af13       0/0              0/0              0/0
24      40  1/10
   af21   27821/1229416         0/0              0/0
32      40  1/10
   af22       0/0              0/0              0/0
28      40  1/10

[output omitted]

    Class-map: af11-traffic (match-all)
      12326 packets, 543265 bytes
      5 minute offered rate 6000 bps, drop rate 0 bps
      Match:  dscp af11
      Queueing
        Output Queue: Conversation 139
        Bandwidth remaining 13 (%)
        (pkts matched/bytes matched) 232/10208
        (depth/total drops/no-buffer drops) 0/0/0
         exponential weight: 9
         mean queue depth: 0
dscp    Transmitted      Random drop      Tail drop    Minimum
Maximum  Mark
        pkts/bytes       pkts/bytes       pkts/bytes    thresh
thresh  prob
af11   12464/549337          0/0              0/0           32
40  1/10
af12       0/0              0/0              0/0           28
40  1/10
af13       0/0              0/0              0/0           24
40  1/10
af21       0/0              0/0              0/0           32
40  1/10

[output omitted]

    Class-map: cs1-traffic (match-all)
      5714 packets, 272308 bytes
      5 minute offered rate 3000 bps, drop rate 0 bps
      Match:  dscp cs1
          Queueing
```



```
        Output Queue: Conversation 140
        Bandwidth remaining 2 (%)
        (pkts matched/bytes matched) 70/3080
        (depth/total drops/no-buffer drops) 0/0/0
         exponential weight: 9
         mean queue depth: 0
dscp    Transmitted      Random drop      Tail drop    Minimum
Maximum  Mark
        pkts/bytes       pkts/bytes       pkts/bytes    thresh
thresh  prob
af11       0/0              0/0              0/0           32
40  1/10
af12       0/0              0/0              0/0           28
40  1/10
af13       0/0              0/0              0/0           24
40  1/10
af21       0/0              0/0              0/0           32
40  1/10
af22       0/0              0/0              0/0           28
40  1/10
af23       0/0              0/0              0/0           24
40  1/10
af31       0/0              0/0              0/0           32
40  1/10
af32       0/0              0/0              0/0           28
40  1/10
af33       0/0              0/0              0/0           24
40  1/10
af41       0/0              0/0              0/0           32
40  1/10
af42       0/0              0/0              0/0           28
40  1/10
af43       0/0              0/0              0/0           24
40  1/10
 cs1    5822/277060          0/0              0/0           22
40  1/10

[output omitted]

    Class-map: class-default (match-any)
      11939 packets, 2066458 bytes
      5 minute offered rate 68000 bps, drop rate 29000 bps
      Match: any
      Queueing
        Output Queue: Conversation 141
        Bandwidth remaining 25 (%)
        (pkts matched/bytes matched) 14950/2588882
        (depth/total drops/no-buffer drops) 58/10269/0
         exponential weight: 9
         mean queue depth: 43
dscp    Transmitted      Random drop      Tail drop    Minimum
Maximum Mark
        pkts/bytes       pkts/bytes       pkts/bytes    thresh
thresh  prob

[output omitted]

 cs5       0/0              0/0              0/0           30
40  1/10
 cs6       6/504            0/0              0/0           32
40  1/10
```

```
 cs7        0/0                0/0                0/0             34
40  1/10
  ef        0/0                0/0                0/0             36
40  1/10
rsvp        0/0                0/0                0/0             36
40  1/10
default     4872/1481315        254/75741        10434/1136558   20
40  1/10
```

Do you see any drops from any of the traffic classes? _____

If so, which one or ones?

_____

Note:	You may see drops on only one of the workgroup routers due to the varying traffic rate from the Pagent routers.

How many af31 packets have been transmitted within the af31-traffic class?

_____

How many af21 packets have been transmitted within the af21-traffic class?

_____

How many af11 packets have been transmitted within the af11-traffic class?

_____

How many cs1 packets have been transmitted within the cs1-traffic class?

_____

The class-default traffic class has transmitted packets marked with which DSCP setting? _____

What types of packets are marked with CS 6 by the Cisco IOS software?

_____

Step 7 Repeat Steps 1 through 7 for the WGxR2 router.

The default class has many dropped packets in it. Although WRED congestion avoidance has been applied and is randomly dropping packets in this class, it may be dropping packets unnecessarily. Ideally, the router should send traffic without dropping as the average queue size increases, but allow the end station to signal it to slow down.

Step 8 Enable WRED ECN for the class-default traffic class on the WGxR1 and WGxR2 routers.

Step 9 On WGxR1 and WGxR2, display the llq-policy policy map and verify the ECN settings for the class-default traffic class.

```
<output omitted>

    Class class-default
      Bandwidth remaining 25 (%)
            exponential weight 9
            explicit congestion notification
            dscp     min-threshold    max-threshold    mark-
 probablity
            ------------------------------------------------
 --------
            af11      -                -                1/10
            af12      -                -                1/10
            af13      -                -                1/10
            af21      -                -                1/10
            af22      -                -                1/10
            af23      -                -                1/10
            af31      -                -                1/10
            af32      -                -                1/10
            af33      -                -                1/10
            af41      -                -                1/10
            af42      -                -                1/10
            af43      -                -                1/10
            cs1       -                -                1/10
            cs2       -                -                1/10
            cs3       -                -                1/10
            cs4       -                -                1/10
            cs5       -                -                1/10
            cs6       -                -                1/10
            cs7       -                -                1/10
            ef        -                -                1/10
            rsvp      -                -                1/10
            default   20               40               1/10
```

Step 10 Clear the interface counters on both your workgroup routers using the clear counters command.

Step 11 Wait for the interface counters to accumulate traffic statistics for at least 1 minute.

Step 12 Display the output service policy on the serial 0/0 interface on both WGxR1 and WGxR2.

```
 Serial0/0

  Service-policy output: llq-policy

 <output omitted>

    Class-map: class-default (match-any)
      13451 packets, 2328036 bytes
      5 minute offered rate 68000 bps, drop rate 33000 bps
      Match: any
      Queueing
        Output Queue: Conversation 141
        Bandwidth remaining 25 (%)
        (pkts matched/bytes matched) 16611/2876499
        (depth/total drops/no-buffer drops) 0/11398/0
         exponential weight: 9
```

explicit congestion notification
mean queue depth: 0

dscp	Transmitted	**Random drop**	Tail drop	Minimum	Maximum	Mark
	pkts/bytes	**pkts/bytes**	pkts/bytes	thresh	thresh	prob
af11	0/0	0/0	0/0	32	40	1/10
af12	0/0	0/0	0/0	28	40	1/10
af13	0/0	0/0	0/0	24	40	1/10
af21	0/0	0/0	0/0	32	40	1/10
af22	0/0	0/0	0/0	28	40	1/10
af23	0/0	0/0	0/0	24	40	1/10
af31	0/0	0/0	0/0	32	40	1/10
af32	0/0	0/0	0/0	28	40	1/10
af33	0/0	0/0	0/0	24	40	1/10
af41	0/0	0/0	0/0	32	40	1/10
af42	0/0	0/0	0/0	28	40	1/10
af43	0/0	0/0	0/0	24	40	1/10
cs1	0/0	0/0	0/0	22	40	1/10
cs2	0/0	0/0	0/0	24	40	1/10
cs3	0/0	0/0	0/0	26	40	1/10
cs4	0/0	0/0	0/0	28	40	1/10
cs5	0/0	0/0	0/0	30	40	1/10
cs6	7/588	0/0	0/0	32	40	1/10
cs7	0/0	0/0	0/0	34	40	1/10
ef	0/0	0/0	0/0	36	40	1/10
rsvp	0/0	0/0	0/0	36	40	1/10
default	5417/1627803	**343/110564**	11462/1242888	20	40	1/10

dscp	ECN Mark
	pkts/bytes
af11	0/0
af12	0/0
af13	0/0
af21	0/0
af22	0/0
af23	0/0


```
   af31        0/0
   af32        0/0
   af33        0/0
   af41        0/0
   af42        0/0
   af43        0/0
    cs1        0/0
    cs2        0/0
    cs3        0/0
    cs4        0/0
    cs5        0/0
    cs6        0/0
    cs7        0/0
     ef        0/0
   rsvp        0/0
default        0/0
```

Is ECN enabled for the class-default traffic class? _____

What is the mean queue depth for the class-default traffic class? _____

How many ECN marked packets are there for the class-default traffic class if any? _____

No packets are marked ECN, yet there are many random WRED drops in the default class. Explain. _____

Step 13 On WGxR1, display the class map for the ef-traffic service class.

```
Class Map match-any ef-traffic (id 10)
  Match dscp ef
  Match access-group 100
```

Step 14 On WGxR1, remove ICMP packets from the EF service class by removing the match access-group 100 from the ef-traffic class map.

Step 15 On WGxR1, display the class map for the ef-traffic service class and verify ICMP (access-group 100) has been removed.

```
Class Map match-any ef-traffic (id 10)
  Match   dscp ef
```

Step 16 Which traffic class will the ICMP traffic belong to now that it has been removed from the EF service class? _____

Step 17 Repeat Steps 13 through 15 for WGxR2.

Step 18 Clear the interface counters on both your workgroup routers using the clear counters command.

Step 19 From the WGxR1 workgroup router, perform an extended ping to the WGxR2 router serial 0/0 interface then record the ping response time in the table at the end of the lab. For the extended ping, use a repeat count of 50, a datagram size of 1500, and use extended commands to set the ToS to 0x02.

```
WGxR1#ping
Protocol [ip]:
Target IP address:10.2.x.2
Repeat count [5]: 50
Datagram size [100]: 1500
Timeout in seconds [2]:
Extended commands [n]: y
Source address or interface:
Type of service [0]: 0x02
Set DF bit in IP header? [no]:
Validate reply data? [no]:
Data pattern [0xABCD]:
Loose, Strict, Record, Timestamp, Verbose[none]:
Sweep range of sizes [n]:
Type escape sequence to abort.
Sending 50, 1500-byte ICMP Echos to 10.2.1.2, timeout is 2
seconds:
.!.!..!!.!.!!.!!.!..!..!!.!!.!.!..!!!.!!...!!..!.!
Success rate is 52 percent (26/50), round-trip min/avg/max =
84/397/802 ms
```

What does setting the ToS byte to 0x02 achieve? _____

Step 20 Display the output service policy on the serial 0/0 interface for the class-default service class only.

```
 Serial0/0

  Service-policy output: llq-policy

    Class-map: class-default (match-any)
      13805 packets, 2561463 bytes
      5 minute offered rate 77000 bps, drop rate 2000 bps
      Match: any
      Queueing
        Output Queue: Conversation 141
        Bandwidth remaining 25 (%)
        (pkts matched/bytes matched) 13788/2560575
        (depth/total drops/no-buffer drops) 2/483/0
         exponential weight: 9
         explicit congestion notification
         mean queue depth: 10
```

dscp	Transmitted	**Random drop**	Tail drop	Minimum	Maximum	Mark
	pkts/bytes	**pkts/bytes**	pkts/bytes	thresh	thresh	prob
af11	0/0	0/0	0/0	32	40	1/10
af12	0/0	0/0	0/0	28	40	1/10
af13	0/0	0/0	0/0	24	40	1/10
af21	0/0	0/0	0/0	32	40	1/10


```
   af22       0/0                0/0                0/0            28      40
1/10
   af23       0/0                0/0                0/0            24      40
1/10
   af31       0/0                0/0                0/0            32      40
1/10
   af32       0/0                0/0                0/0            28      40
1/10
   af33       0/0                0/0                0/0            24      40
1/10
   af41       0/0                0/0                0/0            32      40
1/10
   af42       0/0                0/0                0/0            28      40
1/10
   af43       0/0                0/0                0/0            24      40
1/10
    cs1       0/0                0/0                0/0            22      40
1/10
    cs2       0/0                0/0                0/0            24      40
1/10
    cs3       0/0                0/0                0/0            26      40
1/10
    cs4       0/0                0/0                0/0            28      40
1/10
    cs5       0/0                0/0                0/0            30      40
1/10
    cs6       9/756              0/0                0/0            32      40
1/10
    cs7       0/0                0/0                0/0            34      40
1/10
     ef       0/0                0/0                0/0            36      40
1/10
   rsvp       0/0                0/0                0/0            36      40
1/10
default   13741/2556848        308/50344        175/27697         20      40
1/10
```

```
        dscp      ECN Mark
                  pkts/bytes
        af11         0/0
        af12         0/0
        af13         0/0
        af21         0/0
        af22         0/0
        af23         0/0
        af31         0/0
        af32         0/0
        af33         0/0
        af41         0/0
        af42         0/0
        af43         0/0
         cs1         0/0
         cs2         0/0
         cs3         0/0
         cs4         0/0
         cs5         0/0
         cs6         0/0
         cs7         0/0
          ef         0/0
        rsvp         0/0
```

```
default         6/9024
```

Do you see any ECN marked packets for the class-default now?

_____

Note:	You may only see ECN marked packets on one of the two workgroup routers.

Step 21 Return ICMP traffic (match access-group 100) to the EF service class.

Step 22 Display the class map for the ef-traffic service class and verify that ICMP traffic is now a member of the EF service class.

```
Class Map match-any ef-traffic (id 10)
  Match dscp ef
  Match access-group 100
```

Step 23 Save your running configurations of the workgroup routers and the workgroup switch to the startup-config in NVRAM.

Exercise Verification

You have completed this exercise when you attain these results:

- You have successfully configured DSCP-based CB-WRED on WGxR1 and WGxR2 in all service classes except for the EF service class
- You have successfully configured WRED with ECN on the default class on WGxR1 and WGxR2
- You have successfully verified the operation of DSCP-Based CB-WRED with ECN

Lab Exercise 7-1: Configuring Class-Based Policing

Complete this lab exercise to practice what you learned in the related lesson.

Exercise Objective

In this exercise, you will configure class-based policing to rate limit incoming packets on an interface. After completing this exercise, you will be able to meet these objectives:

- Configure class-based policing
- Monitor the operation of class-based policing

Visual Objective

The figure illustrates what you will accomplish in this exercise.

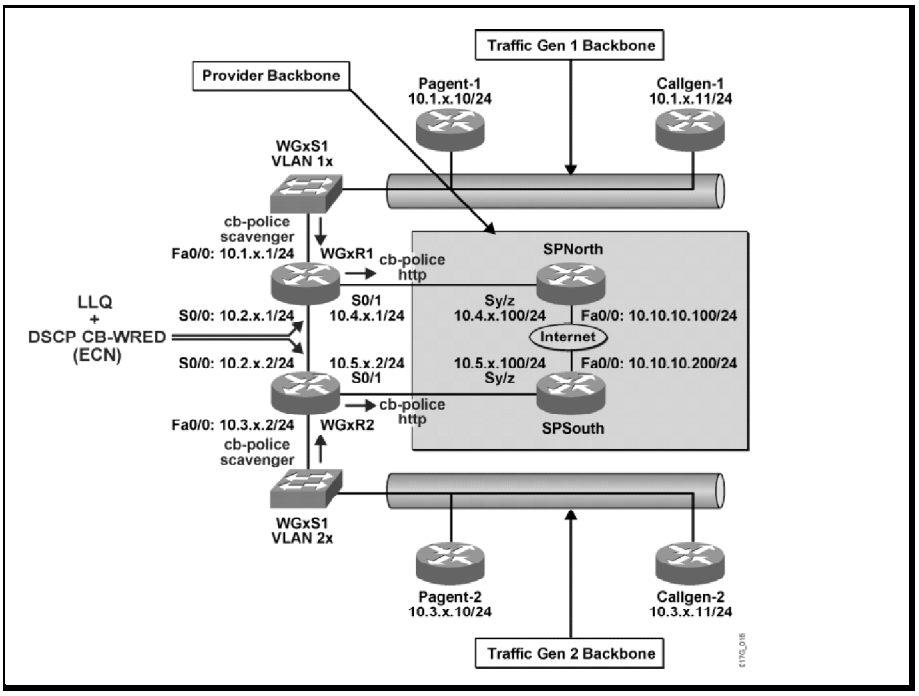

Because the wireless Internet access for the students and faculties have been implemented, you also noticed that the peer-to-peer file sharing traffic (particularly, Napster and Kazaa) are constantly increasing. Therefore, it is now required to police the Napster and Kazaa traffic using class-based policing inbound to the Fa0/0 interface on the workgroup router.

For the Internet connection, most of the traffic from the E-Commerce University is http (web) traffic out to the Internet. The service provider is providing sub-rate access and is implementing an input service policy to police E-Commerce University inbound traffic.

Policing of the http traffic using class-based policing outbound to the S0/1 interface on workgroup router will be implemented to conserve bandwidth on the E-Commerce University Internet connection.

This http policing policy in not placed inbound to the Fa0/0 interface on the workgroup router because intranet http traffic will still be required to flow between the two E-Commerce University campuses across their 384-kbps lease line connection (S0/0).

Required Resources

These are the resources and equipment required to complete this exercise:

- Lab topology configured for QOS
- Student workgroup consisting of two user-controlled Cisco 2610XM routers and one user-controlled Cisco 2950T-24 workgroup switch
- Classroom reference materials as follows:
  - QOS Student Guide
  - QOS Lab Guide
- Student pod workstation with Telnet or console access to workstation pod devices

Command List

The commands used in this exercise are described in the table here.

Configuring Class-Based Policing Lab Commands

Command	Description
policy-map *policy-map-name*	To create or modify a policy map that can be attached to one or more interfaces
class {*class-name* | **class-default**}	To specify the name of the class whose policy you want to create or change or to specify the default class
police *bps* [*burst-normal*] [*burst-max*] **conform-action** *action* **exceed-action** *action* [**violate-action** *action*]	To configure traffic policing
clear counters	To clear the interface counters
shutdown	To disable an interface
class-map *class-map-name*	To create a class map to be used for matching packets to a specified class
match protocol *protocol-name*	To configure the match criteria for a class map on the basis of the specified protocol
police cir percent *percent* [**bc** *conform-burst-in-msec*] [**pir percent** *percent*] [**be** *peak-burst-in-msec*]	To configure traffic policing on the basis of a percentage of bandwidth available on an interface
show policy-map [*policy-map*]	To display the configuration of all classes for a specified service policy map or all classes for all existing policy maps
show policy-map interface *interface-name* [**input** | **output**] [**class** *class-map-name*]	To display the packet statistics of all classes that are configured for all service policies either on the specified interface or subinterface
show policy-map *policy-map* **class** *class-name*	To display the configuration for the specified class of the specified policy map
show class-map [*class-map-name*]	To display all class maps and their matching criteria
copy running-config startup-config	Save your entries in the configuration file

Job Aids

These job aids are available to help you complete the laboratory exercise:

- Your assigned workgroup pod number provided by the instructor

Exercise Procedure

Complete these steps:

Step 1 Modify the existing input service policy **mark-nbar** on the FastEthernet 0/0 interface of the WGxR1 and WGxR2 routers to police the scavenger traffic class to a maximum rate limit of 8 kbps. All conforming traffic should be sent (transmitted), and all exceeding traffic should be dropped.

In this case, do you need to implement a single or dual token bucket?

_____

Will this be a single- or dual-rate policing implementation?

_____

Step 2 Display the mark-nbar policy map and verify the correct policing configuration.

```
Policy Map mark-nbar
  Class real-time
    set dscp ef
  Class mission-critical
    set dscp af31
  Class interactive
    set dscp af21
  Class bulk
    set dscp af11
  Class scavenger
    set dscp cs1
   police cir 8000 bc 1500
     conform-action transmit
     exceed-action drop
```

What is the default value of Bc in bytes and in bits? _____

How is the default value of Bc calculated by the Cisco IOS software?

_____

Based on the default value of Bc, what is the value of Tc? _____

Step 3 Clear the interface counters on both your workgroup routers using the clear counters command.

Step 4 Wait for the interface counters to accumulate traffic statistics for at least 1 minute.

Step 5 Display the input service-policy on the FastEthernet 0/0 interface of your workgroup routers for the scavenger class only.

```
FastEthernet0/0

 Service-policy input: mark-nbar

   Class-map: scavenger (match-any)
     4587 packets, 1227926 bytes
     5 minute offered rate 41000 bps, drop rate 35000 bps
     Match: protocol kazaa2
       1529 packets, 624944 bytes
       5 minute rate 22000 bps
     Match: protocol napster
       3058 packets, 602982 bytes
```



```
    5 minute rate 20000 bps
  QoS Set
    dscp cs1
      Packets marked 4719
  police:
      cir 8000 bps, bc 1500 bytes
    conformed 1350 packets, 83499 bytes; actions:
      transmit
    exceeded 3369 packets, 1179763 bytes; actions:
      drop
    conformed 8000 bps, exceed 36000 bps

Class-map: class-default (match-any)
  20994 packets, 3892053 bytes
  5 minute offered rate 107000 bps, drop rate 0 bps
  Match: any
```

How many packets have been dropped in the scavenger traffic class?

_____

What is the conformed bit rate for the scavenger traffic?

_____

What is the exceed bit rate for the scavenger traffic?

_____

Step 6 On the WGxR1 and WGxR2 routers, configure a new class map called **web-outbound** and use NBAR to classify all HTTP traffic into that traffic class.

Step 7 Display the newly configured class map and verify its configuration.

```
Class Map match-all web-outbound (id 1)
  Match protocol http
```

Step 8 On the WGxR1 and WGxR2 routers, configure a new policy map called **http-police** to police the web (http) traffic to a CIR of 50 percent of the link bandwidth (use the Cisco IOS software default for Bc and Be). All conforming traffic should be transmitted (sent), and all exceeding traffic should be remarked to CS1 then transmitted (sent). All violating traffic should be dropped.

In this case, do you need to implement a single or dual token bucket?

_____

Step 9 Display the newly configured policy map and verify its configuration.

```
Policy Map http-police
  Class web-outbound
   police cir percent 50 be 0
     conform-action transmit
     exceed-action set-dscp-transmit cs1
     violate-action drop
```

Step 10 Apply the **http-police** policy map to the S0/1 interface of the WGxR1 and WGxR2 routers in the outbound direction. The http-police policy map is applied to the S0/1 interface and not the tunnel interface (from the QoS Pre-Classify lab) because the tunnel interface is only used as a backup for the leased line connection between the two E-Commerce University campuses.

Step 11 To test this lab, it will be necessary to administratively disable (shutdown) the S0/0 interface on the WGxR1 and WGxR2 routers, to force the Pagent http traffic to flow via the S0/1 link. Administratively disable the serial 0/0 interface and clear the interface counters on both the WGxR1 and WGxR2 routers.

Step 12 Wait for the interface counters to accumulate traffic statistics for at least 1 minute.

Step 13 Display the outbound service policy on the Serial 0/1 interface.

```
Class-map: web-outbound (match-all)
     754 packets, 409886 bytes
     5 minute offered rate 16000 bps, drop rate 4000 bps
     Match: protocol http
     police:
         cir 50 %
         cir 384000 bps, bc 12000 bytes, be 1500 bytes
       conformed 664 packets, 332894 bytes; actions:
         transmit
       exceeded 40 packets, 1792 bytes; actions:
         set-dscp-transmit cs1
       violated 50 packets, 75200 bytes; actions:
         drop
       conformed 14000 bps, exceed 0 bps, violate 4000 bps

Class-map: class-default (match-any)
     7253 packets, 860150 bytes
     5 minute offered rate 29000 bps, drop rate 8000 bps
     Match: any
```

Are there any violating HTTP packets being dropped?

_____

Are there any exceeding HTTP packets being remarked to CS1 then sent?

_____

Are there any conforming HTTP packets being sent?

_____

Based on a CIR of 50 percent, what is the CIR (in bps), Bc and Be (in bytes) computed by the Cisco IOS software?

_____

Step 14 Re-enable the serial 0/0 interface on both the WGxR1 and WGxR2 routers.

Step 15 Save your running configurations of the workgroup routers and the workgroup switch to the startup-config in NVRAM.

Exercise Verification

You have completed this exercise when you attain these results:

- You have correctly configured single-rate policing on the FastEthernet 0/0 interface of the WGxR1 and WGxR2 routers.
- You have correctly configured percentage-based policing on the Serial 0/1 interface of the WGxR1 and WGxR2 routers.
- You have verified the successful operation of your class-based policing configuration.

Lab Exercise 7-2: Configuring Class-Based Shaping

Complete this lab exercise to practice what you learned in the related lesson.

Exercise Objective

In this exercise, you will configure class-based shaping to rate-limit outgoing packets on an interface. After completing this exercise, you will be able to meet these objectives:

- Configure class-based shaping
- Monitor the operation of class-based shaping

Visual Objective

The figure illustrates what you will accomplish in this exercise.

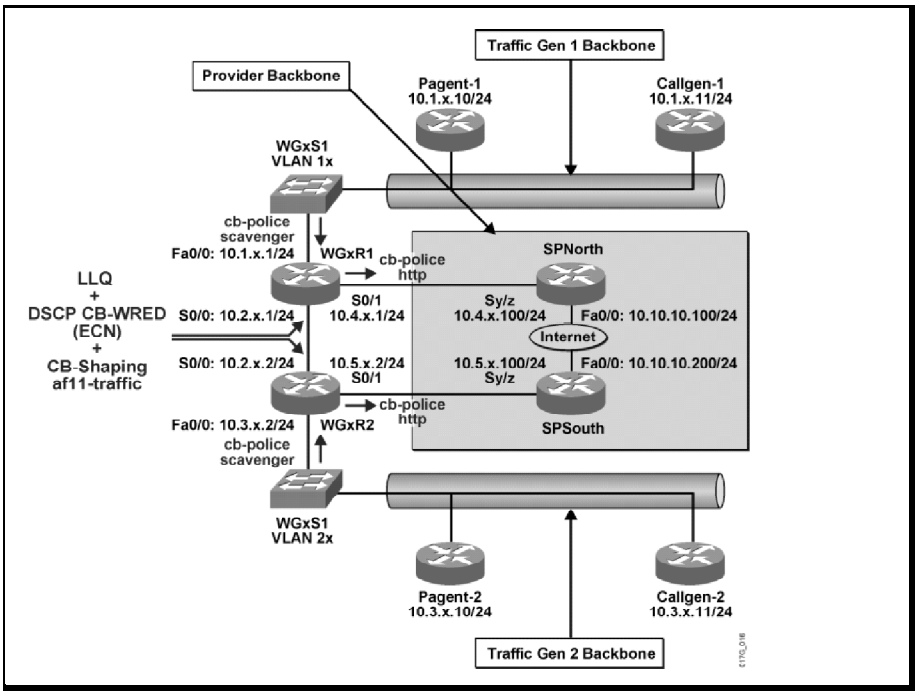

After successfully policing the Napster and Kazaa traffic, it was decided by the E-Commerce University IT staff to also shape the FTP traffic across the 384-kbps lease line since the FTP traffic is also consuming a lot of bandwidth at times. Class-based shaping will be used to apply an upper bandwidth limit to the FTP traffic class.

Required Resources

These are the resources and equipment required to complete this exercise:

- Lab topology configured for QOS
- Student workgroup consisting of two user-controlled Cisco 2610XM routers and one user-controlled Cisco 2950T-24 workgroup switch
- Classroom reference materials as follows:
  - QOS Student Guide
  - QOS Lab Guide
- Student pod workstation with Telnet or console access to workstation pod devices

Command List

The commands used in this exercise are described in the table here.

Configuring Class-Based Shaping Lab Commands

Command	Description
show policy-map *policy-map* **class** *class-name*	To display the configuration for the specified class of the specified policy map
policy-map *policy-map-name*	To create or modify a policy map that can be attached to one or more interfaces
class {*class-name* | **class-default**}	To specify the name of the class whose policy you want to create or change or to specify the default class
shape {**average** | **peak**} *cir* [*bc*] [*be*]	To specify average or peak rate traffic shaping
show policy-map [*policy-map*]	To display the configuration of all classes for a specified service policy map or all classes for all existing policy maps
show policy-map interface *interface-name* [**input** | **output**] [**class** *class-map-name*]	To display the packet statistics of all classes that are configured for all service policies either on the specified interface or subinterface
copy running-config startup-config	Save your entries in the configuration file.

Job Aids

These job aids are available to help you complete the laboratory exercise:

- Your assigned workgroup pod number provided by the instructor

Exercise Procedure

Complete these steps:

Step 1 Verify the existing policy map named llq-policy on the WGxR1 and WGxR2 routers, by displaying only the af11-traffic class. Recall from the classification and marking lab exercise, all FTP traffic is marked with AF11.

```
Class af11-traffic
Bandwidth remaining 13 (%)
      exponential weight 9
      explicit congestion notification
      dscp     min-threshold    max-threshold    mark-probablity
      ------------------------------------------------------------
      af11      26               40               1/10
      af12      -                -                1/10
      af13      -                      -                1/10
      af21      -                      -                1/10
      af22      -                -                1/10
      af23      -                -                1/10
      af31      -                      -                1/10
      af32      -                -                1/10
      af33      -                      -                1/10
      af41      -                      -                1/10
      af42      -                      -                1/10
      af43      -                      -                1/10
      cs1       -                      -                1/10
      cs2       -                      -                1/10
      cs3       -                      -                1/10
      cs4       -                      -                1/10
      cs5       -                      -                1/10
      cs6       -                      -                1/10
      cs7       -                      -                1/10
      ef        -                      -                1/10
      rsvp      -                      -                1/10
      default   -                      -                1/10
```

Step 2 On WGxR1 and WGxR2, modify the existing llq-policy map to shape the af11-traffic class to an average rate of 8 kbps. Allow the Cisco IOS software to automatically calculate the Bc and Be value.

Step 3 Verify the existing llq-policy map and the shaping configuration on your workgroup WGxR1 and WGxR2 routers by displaying only the af11-traffic class.

```
 Class af11-traffic
      Bandwidth remaining 13 (%)
            exponential weight 9
            explicit congestion notification
            dscp     min-threshold    max-threshold    mark-
probablity
            ------------------------------------------------------------
--
            af11      26               40               1/10
            af12      -                -                1/10
            af13      -                -                1/10

           [output omitted]
```

```
                ef          -                   -                   1/10
                rsvp        -                   -                   1/10
                default     -                   -                   1/10

          Traffic Shaping
             Average Rate Traffic Shaping
                     CIR 8000 (bps) Max. Buffers Limit 1000 (Packets)
```

Step 4 Verify the shaping configuration on the existing output service policy map on the Serail 0/0 interface of the WGxR1 and WGxR2 routers, by displaying only the af11-traffic class.

```
      Serial0/0
       Service-policy output: llq-policy
         Class-map: af11-traffic (match-all)
           222923 packets, 10047151 bytes
           5 minute offered rate 4000 bps, drop rate 0 bps
           Match:  dscp af11
           Queueing
             Output Queue: Conversation 139
             Bandwidth remaining 13 (%)
             (pkts matched/bytes matched) 13758/827672
             (depth/total drops/no-buffer drops) 0/0/0
              exponential weight: 9
              mean queue depth: 0
dscp    Transmitted      Random drop      Tail drop    Minimum Maximum
Mark
        pkts/bytes       pkts/bytes       pkts/bytes    thresh  thresh
prob
af11  222923/10047151        0/0              0/0         26      40  1/10
af12       0/0               0/0              0/0         28      40  1/10

[output omitted]

default    0/0               0/0              0/0         20      40  1/10

[output omitted]

      Traffic Shaping
           Target/Average   Byte   Sustain   Excess    Interval  Increment
             Rate           Limit  bits/int  bits/int  (ms)      (bytes)
             8000/8000      2000   8000      8000      1000      1000

        Adapt  Queue     Packets   Bytes     Packets   Bytes     Shaping
        Active Depth                         Delayed   Delayed   Active
        -      0         2682      118008    0         0         no
```

What is the Bc and Be value which is automatically determined by the Cisco IOS software?

_____

What is the committed time window (Tc) (time interval)?

_____

Step 5 Save your running configurations of the workgroup routers and the workgroup switch to the startup-config in NVRAM.

Exercise Verification

You have completed this exercise when you attain these results:

- You have successfully configures average-rate shaping on the serial 0/0 interface of WGxR1 and WGxR2 for the AF11 traffic class.
- You have verified the successful operation of your class-based shaping configuration.

Lab Exercise 8-1: Configuring Class-Based Header Compression

Complete this lab exercise to practice what you learned in the related lesson.

Exercise Objective

In this exercise, you will configure and monitor class-based RTP header compression on a PPP/FR link. After completing this exercise, you will be able to meet these objectives:

- Configure class-based RTP header compression
- Monitor the operation of class-based RTP header compression

Visual Objective

The figure illustrates what you will accomplish in this exercise.

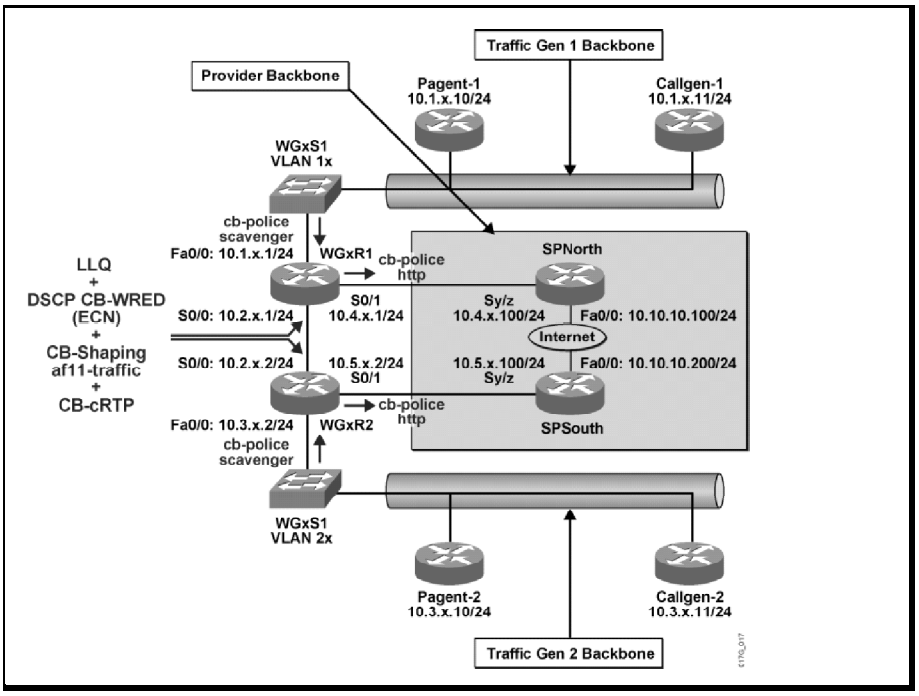

One of the E-Commerce University IT staff was reading a Voice over IP Cisco Press book and realized that the IP/UDP/RTP header overhead for voice packet is very high. Therefore, the IT staff has decided to implement class-based RTP header compression to reduce the size of the packet headers and the associated overhead on the ef-traffic class.

Required Resources

These are the resources and equipment required to complete this exercise:

- Lab topology configured for QOS
- Student workgroup consisting of two user-controlled Cisco 2610XM routers and one user-controlled Cisco 2950T-24 workgroup switch
- Classroom reference materials as follows:
  - QOS Student Guide
  - QOS Lab Guide
- Student pod workstation with Telnet or console access to workstation pod devices

Command List

The commands used in this exercise are described in the table here.

Configuring Class-Based Header Compression Lab Commands

Command	Description
policy-map *policy-map-name*	To create or modify a policy map that can be attached to one or more interfaces
class {*class-name* | **class-default**}	To specify the name of the class whose policy you want to create or change or to specify the default class
compression header ip [**rtp** | **tcp**]	To configure Real-Time Transport Protocol (RTP) or TCP IP header compression for a specific class
show policy-map [*policy-map*]	To display the configuration of all classes for a specified service policy map or all classes for all existing policy maps
show policy-map interface *interface-name* [**input** | **output**] [**class** *class-map-name*]	To display the packet statistics of all classes that are configured for all service policies either on the specified interface or subinterface
clear counters	To clear the interface counters
copy running-config startup-config	Save your entries in the configuration file

Job Aids

These job aids are available to help you complete the laboratory exercise:

- Your assigned workgroup pod number provided by the instructor

Exercise Procedure

Complete these steps:

Step 1 Modify the existing llq-policy map on the WGxR1 and WGxR2 routers to enable class-based RTP header compression on the ef-traffic class.

Step 2 Display the llq-policy map (just display the ef-traffic class) to verify the class-based RTP header compression configuration on the ef-traffic class.

```
Class ef-traffic
  Strict Priority
  Bandwidth 168 (kbps) Burst 4200 (Bytes)
     compress:
         header ip rtp
```

Step 3 Clear the interface counters on both your workgroup routers using the clear counters command.

Step 4 Wait for the interface counters to accumulate traffic statistics for at least 1 minute.

Step 5 Display the output service policy on the workgroup routers S0/0 interface (only show the ef-traffic class).

```
Serial0/0

  Service-policy output: llq-policy

    Class-map: ef-traffic (match-any)
      325 packets, 65932 bytes
      5 minute offered rate 2000 bps, drop rate 0 bps
      Match:  dscp ef
        325 packets, 65932 bytes
        5 minute rate 2000 bps
      Match: access-group 100
        0 packets, 0 bytes
        5 minute rate 0 bps
      Queueing
        Strict Priority
        Output Queue: Conversation 136
        Bandwidth 168 (kbps) Burst 4200 (Bytes)
        (pkts matched/bytes matched) 5/830
        (total drops/bytes drops) 0/0
      compress:
       header ip rtp
       UDP/RTP compression:
       Sent: 324 total, 321 compressed,
             12135 bytes saved, 52441 bytes sent
             1.23 efficiency improvement factor
             99% hit ratio, five minute miss rate 0
misses/sec, 0 max
             rate 2000 bps
```

How many RTP packets were sent? _____

How many RTP packets were compressed? _____

How many bytes were saved because of RTP header compression? _____

Step 6 Save your running configurations of the workgroup routers and the workgroup switch to the startup-config in NVRAM.

Exercise Verification

You have completed this exercise when you attain these results:

- You have successfully configured class-based RTP header compression on the serial 0/0 interface of WGxR1 and WGxR2 for the ef-traffic service class.
- You have successfully verified the operation of the class-based RTP header compression operation.

Lab Exercise 8-2: Configuring LFI

Complete this lab exercise to practice what you learned in the related lesson.

Exercise Objective

In this exercise, you will configure and monitor Multilink PPP (MLP) with interleaving on a PPP link. After completing this exercise, you will be able to meet these objectives:

- Configure link fragmentation and interleaving on PPP WAN Links
- Monitor the operation of link fragmentation and interleaving

Visual Objective

The figure illustrates what you will accomplish in this exercise.

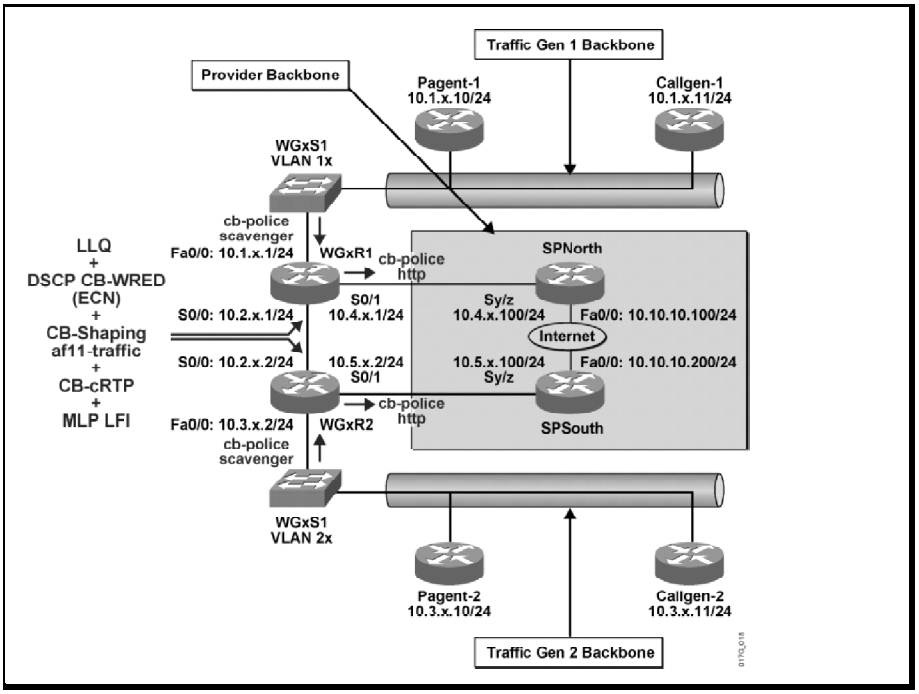

After measuring voice packet delay and jitter over their slow WAN link, the E-Commerce IT staff is concerned that the jitter is still too high. The IT staff asks you why there is too much jitter and how it can be reduced. Being well educated by Cisco cell loss priorities (CLPs) in QoS, you answer, “The LLQ mechanism prioritized voice traffic in the software queue, but the hardware queue (Tx ring) always uses a FIFO scheduling mechanism. Therefore, after packets of different applications leave the software queue, they will mix with other packets in the hardware transmit queue (TxQ), even if their software queue processing was expedited. Thus, a voice packet may be immediately sent to the hardware tx-queue, where two large FTP packets may still be waiting for transmission. The voice packet must wait until the FTP packets are transmitted, thus producing an unacceptable delay in the voice path. Because links are variably utilized, this delay varies with time and may produce unacceptable jitter in jitter-sensitive applications such as voice.”

Knowing that there will be no budget to upgrade the slow WAN link speed any time soon, you offered to improve voice delay and jitter by implementing MLP with interleaving on the slow WAN link.

Required Resources

These are the resources and equipment required to complete this exercise:

- Lab topology configured for QOS
- Student workgroup consisting of two user-controlled Cisco 2610XM routers and one user-controlled Cisco 2950T-24 workgroup switch
- Classroom reference materials as follows:
  - QOS Student Guide
  - QOS Lab Guide
- Student pod workstation with Telnet or console access to workstation pod devices

Command List

The commands used in this exercise are described in the table here.

Configuring LFI Lab Commands

Command	Description
interface *interface-id*	Enter interface configuration mode and the physical interface identification
encapsulation *encapsulation-type*	To set the encapsulation method used by the interface
[no] ip address ip-address mask	To set a primary or secondary IP address for an interface
bandwidth *kbps*	To set and communicate to higher-level protocols the current bandwidth value for an interface
interface multilink *multilink-bundle-number*	To create a multilink bundle and enter multilink interface configuration mode
ppp multilink	To enable MLP on an interface
ppp multilink interleave	Enables real-time packet interleaving
ppp multilink fragment-delay	(Optional) Configures a maximum fragment delay. If, for example, you want a voice stream to have a maximum bound on delay of 20 ms and you specify 20 ms using this command, MLP will choose a fragment size based on the configured value.
ppp multilink group *group-number*	To restrict a physical link to joining only a designated multilink-group interface
[no] service-policy {input | output} *policy-map-name*	To attach a policy map to an input interface or VC, or an output interface or VC

Job Aids

These job aids are available to help you complete the laboratory exercise:

- Your assigned workgroup pod number provided by the instructor

Exercise Procedure

Complete these steps:

Caution: The following steps must be performed in the order shown.

Step 1 Remove the IP address from the serial 0/0 interface of the WGxR1 router.

Step 2 Create a multilink virtual interface (multilink 1) on the workgroup WGxR1 router.

Step 3 Set the bandwidth and IP address on the workgroup WGxR1 router multilink 1 interface as follows:

Parameter	Value
IP Address	10.2.x.1
Subnet Mask	255.255.255.0
bandwidth	384kbps

Step 4 Place the workgroup WGxR1 router S0/0 interface into multilink-group 1.

Step 5 Repeat Steps 1 through 4 for the workgroup WGxR2 router, using the following table for Step 3:

Parameter	Value
IP Address	10.2.x.2
Subnet Mask	255.255.255.0
bandwidth	384kbps

Step 6 Display the running-config of the S0/0 and the multilink 1 interface to verify the MLP configuration.

```
Current configuration : 166 bytes
!
interface Serial0/0
 description to wgxr1
 bandwidth 384
 no ip address
 service-policy output llq-policy
 encapsulation ppp
 ppp multilink
 multilink-group 1
end

Current configuration : 112 bytes
!
interface Multilink1
 bandwidth 384
```



```
 ip address 10.2.x.1 255.255.255.0
 ppp multilink
 multilink-group 1
end
```

MLP is enabled now but is Interleaving also enabled based on the above configuration? _____

Step 7 Use the show ip interface brief command to verify that the multilink-group 1 interface is UP, has the proper IP address and that the S0/0 interface is UP.

```
WGxR1#show ip interface brief

Interface        IP-Address OK? Method Status
      Protocol
FastEthernet0/0  10.1.1.1   YES NVRAM  up                    up
Serial0/0        unassigned YES manual up                    up
Serial0/1        10.4.1.1   YES NVRAM  administratively
downdown
Multilink1       10.2.1.1   YES manual up                    up
Virtual-Access1  unassigned YES unset  up                    up
```

Step 8 Ping the multilink 1interface of WGxR2 from WGxR1.

```
WGxR1>ping 10.2.x.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.2.x.2, timeout is 2
seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max =
4/7/8 ms
```

Step 9 Ping the multilink 1interface of WGxR1 from WGxR2.

```
WGxR2>ping 10.2.x.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.2.x.1, timeout is 2
seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max =
4/7/8 ms
```

Step 10 Enter the **show ip ospf neighbor** command to verify the OSPF neighbor relationship is now formed over the multilink interface.

```
WGxR2#show ip ospf neighbor

Neighbor ID  Pri  State   Dead Time  Address    Interface
10.2.1.1      0   FULL/ - 00:00:30   10.2.1.1   Multilink1
```

Step 11 Enable PPP multilink interleaving on the multilink 1 interface on both WGxR1 and WGxR2. Use a fragment-delay of 10.

What is the unit of the fragment delay? _____

Step 12 Disable the output service policy on the S0/0 interface of the WGxR1 and WGxR2 routers.

Step 13 Enable the llq-policy on the multilink1 interface in the outbound direction on WGxR1 and WGxR2.

Note: Ensure that the OSPF neighbor state is established. You can use the **show ip ospf neighbor** command to verify this.

Step 14 Clear the interface counters on both your workgroup routers using the clear counters command.

Step 15 Wait for the interface counters to accumulate traffic statistics for at least 1 minute.

Step 16 What type of queuing is the S0/0 interface using now? _____.

```
WGxR1#show interface serial0/0
Serial0/0 is up, line protocol is up
  Hardware is PowerQUICC Serial
  Description: to wg1r2
  MTU 1500 bytes, BW 384 Kbit, DLY 20000 usec,
     reliability 255/255, txload 16/255, rxload 16/255
  Encapsulation PPP, LCP Open, multilink Open, loopback not
set
  Last input 00:00:02, output 00:00:02, output hang never
  Last clearing of "show interface" counters 00:05:38
  Input queue: 0/75/0/0 (size/max/drops/flushes); Total output
drops: 0
  Queueing strategy: weighted fair  [suspended, using FIFO]
  FIFO output queue 0/40, 0 drops
  5 minute input rate 25000 bits/sec, 9 packets/sec
  5 minute output rate 25000 bits/sec, 5 packets/sec
     3219 packets input, 1124012 bytes, 0 no buffer
     Received 0 broadcasts, 0 runts, 0 giants, 0 throttles
     0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored, 0
abort
     1641 packets output, 1132534 bytes, 0 underruns
     0 output errors, 0 collisions, 0 interface resets
     0 output buffer failures, 0 output buffers swapped out
     0 carrier transitions
     DCD=up  DSR=up  DTR=up  RTS=up  CTS=up
```

Step 17 Display the multilink interface and examine how many packets have been interleaved?

```
Multilink1 is up, line protocol is up
  Hardware is multilink group interface
  Internet address is 10.2.x.1/24
  MTU 1500 bytes, BW 384 Kbit, DLY 100000 usec,
     reliability 255/255, txload 145/255, rxload 32/255
  Encapsulation PPP, LCP Open, multilink Open
  Open: IPCP, loopback not set
  DTR is pulsed for 2 seconds on reset
```



```
  Last input 00:00:07, output never, output hang never
  Last clearing of "show interface" counters 00:01:24
  Input queue: 0/75/0/0 (size/max/drops/flushes); Total output
drops: 2927
  Queueing strategy: weighted fair
  Output queue:
91/1000/64/2927/3030(size/maxtotal/threshold/drops/interleaves
)
     Conversations  5/6/128 (active/max active/max total)
     Reserved Conversations 5/5 (allocated/max allocated)
     Available Bandwidth 32 kilobits/sec
  5 minute input rate 49000 bits/sec, 112 packets/sec
  5 minute output rate 219000 bits/sec, 183 packets/sec
     13826 packets input, 608664 bytes, 0 no buffer
     Received 0 broadcasts, 0 runts, 0 giants, 0 throttles
     0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored, 0
abort
     24456 packets output, 3740947 bytes, 0 underruns
     0 output errors, 0 collisions, 0 interface resets
     0 output buffer failures, 0 output buffers swapped out
     0 carrier transitions
```

Step 18 Use the **show ppp multilink** command to examine the fragment size in bytes (calculated by the Cisco IOS software based on the fragment delay of 10 ms).

```
WGxR1#show ppp multilink

Multilink1, bundle name is wg1r2
  Bundle up for 16:55:24, 183/255 load
  Receive buffer limit 12192 bytes, frag timeout 1000 ms
    0/0 fragments/bytes in reassembly list
    0 lost fragments, 0 reordered
    0/0 discarded fragments/bytes, 0 lost received
    0x47C9F received sequence, 0x73936 sent sequence
  Member links: 1 active, 0 inactive (max not set, min not
set)
    Se0/0, since 03:54:35, 480 weight, 472 frag size
```

Step 19 Save your running configurations of the workgroup routers and the workgroup switch to the startup-config in NVRAM.

Exercise Verification

You have completed this exercise when you attain these results:

- You have successfully created a MLP interface on WGxR1 and WGxR2.
- You have successfully configured link fragmentation and interleaving (LMI) over the multilink 1 interface on WGxR1 and WGxR2.
- You have successfully verified the operation of your LFI configuration.

Lab Exercise Answer Key

Lab Exercise 2-1: QoS Lab Setup and Initialization

When you complete this lab exercise, your router configuration will be similar to the following, with differences that are specific to your pod.

WG1R1

```
hostname WG1R1
!
enable secret 5 $1$n4//$vbCjudYcBR3yNPJqI.1tT0
!
ip subnet-zero
!
interface FastEthernet0/0
 ip address 10.1.1.1 255.255.255.0
 duplex auto
 speed auto
!
interface Serial0/0
 bandwidth 384
 ip address 10.2.1.1 255.255.255.0
 encapsulation ppp
 clockrate 384000
 no fair-queue
!
interface Serial0/1
 bandwidth 768
 ip address 10.4.1.1 255.255.255.0
 encapsulation ppp
!
router ospf 1
 log-adjacency-changes
 network 10.0.0.0 0.255.255.255 area 0
!
ip http server
ip classless
!
line con 0
line aux 0
line vty 0 4
 password cisco
 login
!
end
```

WG1R2

```
hostname WG1R2
!
enable secret 5 $1$07qt$nKIz/sUIIRYMZ7urfJPtp1
!
ip subnet-zero
!
interface FastEthernet0/0
 ip address 10.3.1.2 255.255.255.0
```



```
 duplex auto
 speed auto
!
interface Serial0/0
 bandwidth 384
 ip address 10.2.1.2 255.255.255.0
 encapsulation ppp
 no fair-queue
!
interface Serial0/1
 bandwidth 768
 ip address 10.5.1.2 255.255.255.0
 encapsulation ppp
!
router ospf 1
 log-adjacency-changes
 network 10.0.0.0 0.255.255.255 area 0
!
line con 0
line aux 0
line vty 0 4
 password cisco
 login
!
end
```

WG1S1

```
hostname WG1S1
!
enable secret 5 $1$Yq48$E3tAlJjcYAP9qJpdmr0nu.
!
vlan 11
 name vlan11
!
vlan 21
 name vlan21
ip subnet-zero
vtp domain qos
vtp mode transparent
!
interface FastEthernet0/1
 switchport trunk allowed vlan 11,21
 switchport mode trunk
 no ip address
!
interface FastEthernet0/2
 switchport access vlan 11
 switchport mode access
 no ip address
!
interface FastEthernet0/3
 switchport access vlan 21
 switchport mode access
 no ip address
!
interface Vlan1
 no ip address
 no ip route-cache
 shutdown
```

```
!
!
line con 0
line vty 5 15
!
end
```

Lab Exercise 2-2: Baseline QoS Measurement

When you complete this lab exercise, your router configuration will be similar to the following, with differences that are specific to your pod.

Sample WGxR1 to WGxR2 QoS Baseline ping Results

Packet Size	Without Pagent/Callgen	With Pagent/Callgen
160 bytes	**min/avg/max**	**min/avg/max**
Extended ping 1	8/8/12	8/43/120
Extended ping 2	8/8/12	16/42/116
Extended ping 3	8/9/16	12/43/112
	success rate %	**success rate %**
Extended ping 1	100%	100%
Extended ping 2	100%	100%
Extended ping 3	100%	100%

Sample WGxR1 QoS Baseline show interfaces Results

	Without Pagent/Callgen	With Pagent/Callgen
Queuing Strategy	WFQ (fair queue)	WFQ (fair queue)
Reliability, Txload, Rxload	255, 1, 1	255, 209, 39
Total Output Drops	0	1975
Output Queue: size/max total	0/1000	61/1000
Output Queue: threshold/drops	64/0	64/1975
Packets Output	3001	131760
Drop % (Calculated by you as: Total Output Drop / Packets Output)	0/3001 = 0	1975/131760 = 1.5%

Lab Exercise 3-1: Configuring QoS with AutoQoS

When you complete this lab exercise, your router configuration will be similar to the following, with differences that are specific to your pod.

Task 1: The following configuration is entered to enable AutoQoS on the WGxR1 and WGxR2 routers:

```
interface serial 0/0
auto qos voip

ip access-list extended AutoQoS-VoIP-RTCP
 permit udp any any range 16384 32767
 permit icmp any any echo
 permit icmp any any echo-reply
```

Task 2: The following configuration is entered to enable AutoQoS on the WGxS1 switch:

```
interface fastethernet 0/1
auto qos voip trust
```

Task 3: The following are the answers to the questions in this lab exercise:

Is the ping maximum response time shorter than before AutoQoS was enabled? Explain.

Yes, because ping was moved into the expedited forwarding class with VoIP traffic. The minimum time may be higher because the traffic generation tools are generating traffic at different data rates.

Is the drop rate higher or lower or about the same than before AutoQoS was enabled? Explain.

The drop rate is about the same because AutoQoS only affects the voice over IP traffic and not the data traffic. The VoIP traffic load in the lab is minimal compared to the data.

The following are sample QoS Lab results:

Sample WGxR1 to WGxR2 AutoQoS ping Results

Packet Size	Without AutoQoS (From Lab 2-2)	With AutoQoS (This Lab)
160 bytes	**min/avg/max**	**min/avg/max**
Extended ping 1	8/43/120	32/54/88
Extended ping 2	16/42/116	32/53/88
Extended ping 3	12/43/112	32/54/88
	success rate %	**success rate %**
Extended ping 1	100%	100%
Extended ping 2	100%	100%
Extended ping 3	100%	100%

Sample WGxR1 AutoQoS show interfaces Results

	Without AutoQoS (From Lab 2-2) show interface s0/0	With AutoQoS (This Lab) show interface multilink
Queuing Strategy	WFQ (fair queue)	WFQ (fair queue)
Reliability, Txload, Rxload	255, 209, 39	255, 245, 41
Total Output Drops	1975	216
Output Queue: size/max total	61/1000	161/1000
Output Queue: threshold/drop/interleaves	64/1975	64/216/8534
Packets Output	131760	20530
Drop % (Calculated by you as: Total Output Drop / Packets Output)	1975/131760 = 1.5%	216/20530=1.1%

Lab Exercise 4-1: Classification and Marking Using MQC

When you complete this lab exercise, your router configuration will be similar to the following, with differences that are specific to your pod.

```
access-list 101 permit tcp any any eq ftp
access-list 101 permit tcp any any eq ftp-data
access-list 102 permit tcp any any eq www


 class-map match-all match-www
  match access-group 102
 class-map match-all match-ftp
  match access-group 101

 policy-map mark-apps
  class match-ftp
   set dscp af11
```



```
  class match-www
   set dscp default

interface FastEthernet0/0
 service-policy input mark-apps
```

Lab Exercise 4-2: Classification Using NBAR

When you complete this lab exercise, your router configuration will be similar to the following, with differences that are specific to your pod.

```
WG1R1(config)#int fa 0/0
WG1R1(config-if)#no service-policy input mark-apps
WG1R1(config-if)#ip nbar protocol-discovery
```

NBAR should have discovered the following protocols:

HTTP	citrix	sqlnet
napster	netbios	FTP
exchange	kazaa2	LDAP
RTP	Unknown	

```
ip access-list extended VoIP-RTCP
 permit udp any any range 16384 32767
!
ip access-list extended Voice-Control
 permit tcp any any eq 1720
 permit tcp any any range 11000 11999
 permit udp any any eq 2427
 permit tcp any any eq 2428
 permit tcp any any range 2000 2002
 permit udp any any eq 1719
 permit udp any any eq 5060
!
 class-map match-all bulk
  match protocol ftp
 class-map match-any real-time
  match protocol rtp
  match protocol icmp
  match access-group name VoIP-RTCP
 class-map match-any mission-critical
  match protocol sqlnet
  match access-group name Voice-Control
 class-map match-all interactive
  match protocol citrix
 class-map match-all default
  match any
 class-map match-any scavenger
  match protocol kazaa2
  match protocol napster
!
 policy-map mark-nbar
  class real-time
   set dscp ef
  class mission-critical
   set dscp af31
```

```
  class interactive
   set dscp af21
  class bulk
   set dscp af11
  class scavenger
   set dscp cs1
  class class-default
   set dscp default
!
interface FastEthernet0/0
 service-policy input mark-nbar
```

Lab Exercise 4-3: Configuring QoS Pre-Classify

When you complete this lab exercise, your router configuration will be similar to the following, with differences that are specific to your pod.

WGxR1:

```
interface Tunnel0
 ip unnumbered FastEthernet0/0
 qos pre-classify
 tunnel source Serial0/1
 tunnel destination 10.5.x.2
!
ip route 10.3.x.0 255.255.255.0 Tunnel0
```

WGxR2:

```
interface Tunnel0
 ip unnumbered FastEthernet0/0
 qos pre-classify
 tunnel source Serial0/1
 tunnel destination 10.4.x.1
!
ip route 10.1.x.0 255.255.255.0 Tunnel0
```

The following are the answers to the questions in this lab exercise:

What happens when QoS pre-classify is configured on the tunnel interface?

The original packet headers are visible to the tunnel 0 interface and hence useable for QoS manipulation.

Can WFQ now distinguish between different application flows? Or does WFQ still only see one flow (protocol 47=GRE)?

Multiple flows are now visible

Lab Exercise 4-4: LAN-Based Packet Classification and Marking

When you complete this lab exercise, your router configuration will be similar to the following, with differences that are specific to your pod.

```
class-map match-all callgen2
  match access-group 2
```



```
class-map match-all callgen1
  match access-group 1
!
!
policy-map mark-callgen
  class callgen1
    set ip dscp 46
  class callgen2
    set ip dscp 46
!
mls qos map cos-dscp 0 8 16 24 32 46 48 56
!
interface FastEthernet0/1
  service-policy input mark-callgen
 mls qos trust device cisco-phone
 mls qos trust cos
!
access-list 1 permit 10.1.1.11
access-list 2 permit 10.3.1.11
```

Lab Exercise 5-1: Configuring Basic Queuing

When you complete this lab exercise, your router configuration will be similar to the following, with differences that are specific to your pod.

Sample WGxR1 RTT Measurement Worksheet

Network Situation	Ping 1 RTT (min/ave/max)	Ping 2 RTT (min/ave/max)	Ping 3 RTT (min/ave/max)
Before Pagnet / Callgen traffic (from 2-2: Baseline QoS Measurement)	8/8/12	8/8/12	8/9/16
After Pagnet/Callgen traffic (from 2-2: Baseline QoS Measurement)	8/43/120	16/42/116	12/43/112
After enabling Auto-QoS (from 3-1: Configuring QoS with AutoQoS)	32/54/88	32/53/88	32/54/88
FIFO (from 5-1: Configuring Basic Queuing)	28/122/265	44/129/228	28/131/249
WFQ (from 5-1: Configuring Basic Queuing)	12/46/128	16/48/188	12/43/121

The following are the answers to the questions in this lab exercise:

Knowing that the S0/0 interface on WGxR1 is 384 kbps, what will be the default queuing mechanism?

Weighted Fair Queuing

Which queuing method is S0/0 using now after WFQ was disabled?

First In First Out (FIFO)

What are some disadvantages of with using FIFO queuing for the voice and mission-critical traffic?

Smaller packets can suffer excessive delay and variable delay waiting in a FIFO queue behind larger data packets. Aggressive flows can also starve fragile flow types like voice and interactive traffic.

What is the default Congestion Discard Threshold (CDT)?

64

What is the default number of maximum conversations on the S0/0 interface?

128

What is a benefit of increasing the maximum conversations?

The probability that two distinct flows will be classified into the same dynamic queue is reduced.

What is the available bandwidth and how is it calculated?

Available bandwidth in this case is 288 kbps. Available bandwidth is calculated as 75% of the configured interface bandwidth.

What is the significance of the weight in these conversations?

Weight is used by WFQ to reduce the finish time of queued packets, making them appear smaller than they are to WFQ. This results in faster dispatching of lower weighted packets.

What factor(s) can influence the weight?

IP Precedence or DSCP Class Selector Markings can influence WFQ weight.

Lab Exercise 5-2: Configuring LLQ

When you complete this lab exercise, your router configuration will be similar to the following, with differences that are specific to your pod.

```
access-list 100 permit icmp any any echo
access-list 100 permit icmp any any echo-reply

 class-map match-any ef-traffic
  match  dscp ef
  match access-group 100
 class-map match-all af21-traffic
  match  dscp af21
 class-map match-all af31-traffic
  match  dscp af31
 class-map match-all af11-traffic
  match  dscp af11
 class-map match-all cs1-traffic
  match  dscp cs1
 class-map match-all class-default
  match any
!
 policy-map llq-policy
  class ef-traffic
   priority 168
  class af31-traffic
```



```
   bandwidth remaining percent 40
  class af21-traffic
   bandwidth remaining percent 20
  class af11-traffic
   bandwidth remaining percent 13
  class cs1-traffic
   bandwidth remaining percent 2
  class class-default
   bandwidth remaining percent 25
!
interface Serial0/0
 service-policy output llq-policy

interface FastEthernet0/0
 service-policy input mark-nbar
```

Sample WGxR1 RTT Measurement Worksheet

Network Situation	Ping 1 RTT (min/ave/max)	Ping 2 RTT (min/ave/max)	Ping 3 RTT (min/ave/max)
Before Pagnet / Callgen traffic (from 2-2: Baseline QoS Measurement)	8/8/12	8/8/12	8/9/16
After Pagnet/Callgen traffic (from 2-2: Baseline QoS Measurement)	8/43/120	16/42/116	12/43/112
After enabling Auto-QoS (from 3-1: Configuring QoS with AutoQoS)	32/54/88	32/53/88	32/54/88
FIFO (from 5-1: Configuring Basic Queuing)	28/122/265	44/129/228	28/131/249
WFQ (from 5-1: Configuring Basic Queuing)	12/46/128	16/48/188	12/43/121
LLQ results from this lab	8/17/72	8/17/56	8/19/77

The following are the answers to the questions in this lab exercise:

Comparing all the results, which QoS mechanism provided the best response time for VoIP packets?

Low Latency Queuing (LLQ)

Lab Exercise 5-3: Queuing on Catalyst Switches

When you complete this lab exercise, your router configuration will be similar to the following, with differences that are specific to your pod.

```
wrr-queue bandwidth 30 1 70 0
wrr-queue cos-map 1 0 1 2 3 4
wrr-queue cos-map 3 6 7
wrr-queue cos-map 4 5
```

The following are the answers to the questions in this lab exercise:

What is the default queuing mechanism on the Catalyst 2950 switch?

Priority Queuing

How many output queues per interface are available for the Catalyst 2950?

Four

Which queue has the highest priority?

Queue 4

What type of queuing is enabled using the following configuration?

```
WGxS1#show auto qos

Initial configuration applied by AutoQoS:
wrr-queue bandwidth 20 1 80 0
no wrr-queue cos-map
wrr-queue cos-map 1 0 1 2 4
wrr-queue cos-map 3 3 6 7
wrr-queue cos-map 4 5
mls qos map cos-dscp 0 8 16 26 32 46 48 56
!
interface FastEthernet0/1
 mls qos trust cos
```

WRR with an Expedite Queue

From the "wrr-queue bandwidth 20 1 80 0" command, why is the weight for queue #4 equal to "0"?

Setting the weight of queue 4 to 0 configures queue 4 as a strict priority queue, also known as an expedite queue.

What do the following commands accomplish?

```
wrr-queue cos-map 1 0 1 2 4
wrr-queue cos-map 3 3 6 7
wrr-queue cos-map 4 5
```

These commands map frames with CoS values to specific output queues.

Lab Exercise 6-1: Configuring DSCP-Based WRED

When you complete this lab exercise, your router configuration will be similar to the following, with differences that are specific to your pod.

Configuration through Step 7:

```
policy-map llq-policy
 class ef-traffic
  priority 168
 class af31-traffic
  bandwidth remaining percent 40
  random-detect dscp-based
  random-detect dscp af31   34    40    10
 class af21-traffic
  bandwidth remaining percent 20
  random-detect dscp-based
  random-detect dscp af21   30    40    10
 class af11-traffic
```



```
 bandwidth remaining percent 13
 random-detect dscp-based
 random-detect dscp af11   26    40    10
class cs1-traffic
 bandwidth remaining percent 2
 random-detect dscp-based
 random-detect dscp cs1   22    40    10
class class-default
 bandwidth remaining percent 25
 random-detect dscp-based
 random-detect dscp 0   20    40    10
```

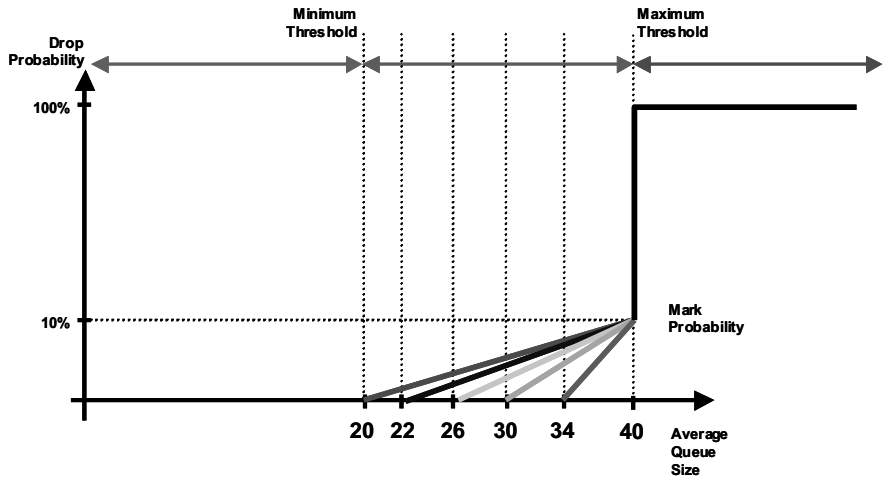

Based on the previous WRED profiles, which traffic class will start dropping packets first?

Default (BE) class

What does a mark probability of 1/10 means?

When the average queue size is at the maximum threshold, the router will be dropping one out of every ten packets.

Why would you not implement WRED for the ef-traffic class?

It is preferred that no packet dropping occurs in the EF traffic class.

What is the default exponential weight constant?

9

What types of packets are marked with CS 6 by the Cisco IOS software?

Router control and routing protocol traffic

Configuration for remainder of the lab:

```
class class-default
 bandwidth remaining percent 25
 random-detect dscp-based
 random-detect ecn
 random-detect dscp 0   20    40    10
```

No packets are marked ECN, yet there are many random WRED drops in the default class. Explain.

The router will only set ECN if the transiting packets are marked such that the endpoints are ECN capable.

Which traffic class will the ICMP traffic belong to now that it has been removed from the EF service class?

Default (BE) service class

What does setting the ToS byte to 0x02 achieve?

It enables ECN in the packet by indicating that the endpoint is ECN capable.

Lab Exercise 7-1: Configuring Class-Based Policing

When you complete this lab exercise, your router configuration will be similar to the following, with differences that are specific to your pod.

Configuration in Step 2:

```
policy-map mark-nbar
 class real-time
  set dscp ef
 class mission-critical
  set dscp af31
 class interactive
  set dscp af21
 class bulk
  set dscp af11
 class scavenger
  set dscp cs1
  police cir 8000
    conform-action transmit
    exceed-action drop
```

Configuration for the remainder of the lab:

```
class-map match-all web-outbound
 match protocol http

policy-map http-police
 class web-outbound
  police cir percent 50
    conform-action transmit
    exceed-action set-dscp-transmit cs1
    violate-action drop
```

The following are the answers to the questions in this lab exercise:

In this case, do you need to implement a single or dual token bucket?

Single Token Bucket

Will this be a single or dual-rate policing implementation?

Single Rate Policing

What is the default value of Bc in bytes and in bits?

1500 bytes or 12,000 bits

How is the default value of Bc calculated by the Cisco IOS software?

Cir / 32 or 1500 bytes, whichever is larger

Based on the default value of Bc, what is the value of Tc?

Tc=Bc / CIR = 12000/8000 = 1/5s

In this case, do you need to implement a single or dual token bucket?

Dual Token Bucket

Based on a CIR of 50 percent, what is the CIR (in bps), Bc and Be(in bytes) calculated by the Cisco IOS software?

CIR=384000 bps, Bc=12,000 bytes, Be=1500 bytes

Lab Exercise 7-2: Configuring Class-Based Shaping

When you complete this lab exercise, your router configuration will be similar to the following, with differences that are specific to your pod.

```
policy-map llq-policy
class ef-traffic
 priority 168
class af31-traffic
 bandwidth remaining percent 40
 random-detect dscp-based
 random-detect dscp 10   34    40   10
class af21-traffic
 bandwidth remaining percent 20
 random-detect dscp-based
 random-detect dscp 10   30    40   10
class af11-traffic
 bandwidth remaining percent 13
 random-detect dscp-based
 random-detect dscp 10   26    40   10
 shape average 8000
class cs1-traffic
 bandwidth remaining percent 2
 random-detect dscp-based
 random-detect dscp 10   22    40   10
class class-default
 bandwidth remaining percent 25
 random-detect dscp-based
 random-detect ecn
 random-detect dscp 0   20    40    10
```

The following are the answers to the questions in this lab exercise:

What is the Bc and Be value which is automatically determined by the Cisco IOS software?

Bc=8000 bps, Be=8000 bps

What is the Tc (time interval)?

1 second

Lab Exercise 8-1: Configuring Class-Based Header Compression

When you complete this lab exercise, your router configuration will be similar to the following, with differences that are specific to your pod.

```
policy-map llq-policy
 class ef-traffic
  priority 168
  compress header ip rtp
 class af31-traffic
  bandwidth remaining percent 40
  random-detect dscp-based
  random-detect dscp 10   34    40    10
 class af21-traffic
  bandwidth remaining percent 20
  random-detect dscp-based
  random-detect dscp 10   30    40    10
 class af11-traffic
  bandwidth remaining percent 13
  random-detect dscp-based
  random-detect dscp 10   26    40    10
  shape average 8000
 class cs1-traffic
  bandwidth remaining percent 2
  random-detect dscp-based
  random-detect dscp 10   22    40    10
 class class-default
  bandwidth remaining percent 25
  random-detect dscp-based
  random-detect ecn
  random-detect dscp 0   20    40    10
```

Lab Exercise 8-2: Configuring LFI

When you complete this lab exercise, your router configuration will be similar to the following, with differences that are specific to your pod.

WGxR1:

```
interface Multilink1
 bandwidth 384
 ip address 10.2.x.1 255.255.255.0
 service-policy output llq-policy
 ppp multilink
 ppp multilink fragment-delay 10
 ppp multilink interleave
 multilink-group 1
!
interface Serial0/0
 bandwidth 384
 no ip address
 encapsulation ppp
 clockrate 384000
 ppp multilink
 multilink-group 1
```

WGxR2:

```
interface Multilink1
 bandwidth 384
 ip address 10.2.x.2 255.255.255.0
 service-policy output llq-policy
 ppp multilink
 ppp multilink fragment-delay 10
 ppp multilink interleave
 multilink-group 1
!
interface Serial0/0
 bandwidth 384
 no ip address
 encapsulation ppp
 ppp multilink
 multilink-group 1
```

The following are the answers to the questions in this lab exercise:

What is the unit of the fragment delay?

milliseconds (ms)

What type of queuing is the S0/0 interface using now?

first-in, first-out (FIFO)

Final Workgroup Device Configurations

When you complete all lab exercises in this course, your workgroup devices will have the following configurations, with differences that are specific to your pod.

WG1R1:

```
hostname WG1R1
!
enable secret 5 $1$n4//$vbCjudYcBR3yNPJqI.1tT0
!
ip subnet-zero
ip cef
!
 class-map match-all bulk
  match protocol ftp
 class-map match-any real-time
  match protocol rtp
  match protocol icmp
  match access-group name VoIP-RTCP
 class-map match-all match-www
  match access-group 102
 class-map match-all match-ftp
  match access-group 101
 class-map match-all web-outbound
  match protocol http
 class-map match-any ef-traffic
  match  dscp ef
  match access-group 100
 class-map match-all af21-traffic
  match  dscp af21
```



```
 class-map match-all af31-traffic
  match  dscp af31
 class-map match-all af11-traffic
  match  dscp af11
 class-map match-any mission-critical
  match protocol sqlnet
  match access-group name Voice-Control
 class-map match-all cs1-traffic
  match  dscp cs1
 class-map match-all interactive
  match protocol citrix
 class-map match-all default
  match any
 class-map match-all match-sw-be
  match  dscp default
 class-map match-all match-sw-ef
  match  dscp ef
 class-map match-any scavenger
  match protocol kazaa2
  match protocol napster
!
 policy-map mark-nbar
  class real-time
   set dscp ef
  class mission-critical
   set dscp af31
  class interactive
   set dscp af21
  class bulk
   set dscp af11
  class scavenger
   set dscp cs1
   police cir 8000
     conform-action transmit
     exceed-action drop
  class class-default
   set dscp default
 policy-map http-police
  class web-outbound
   police cir percent 50
     conform-action transmit
     exceed-action set-dscp-transmit cs1
     violate-action drop
 policy-map mark-apps
  class match-ftp
   set dscp af11
  class match-www
   set dscp default
 policy-map verify-mark
  class match-sw-ef
  class match-sw-be
 policy-map llq-policy
  class ef-traffic
   priority 168
   compress header ip rtp
  class af31-traffic
   bandwidth remaining percent 40
   random-detect dscp-based
   random-detect dscp 10   34    40    10
  class af21-traffic
   bandwidth remaining percent 20
```

```
   random-detect dscp-based
   random-detect dscp 10   30    40    10
  class af11-traffic
   bandwidth remaining percent 13
   random-detect dscp-based
   random-detect dscp 10   26    40    10
   shape average 8000
  class cs1-traffic
   bandwidth remaining percent 2
   random-detect dscp-based
   random-detect dscp 10   22    40    10
  class class-default
   bandwidth remaining percent 25
   random-detect dscp-based
   random-detect ecn
   random-detect dscp 0   20    40    10
!
interface Multilink1
 bandwidth 384
 ip address 10.2.1.1 255.255.255.0
 service-policy output llq-policy
 ppp multilink
 ppp multilink fragment-delay 10
 ppp multilink interleave
 multilink-group 1
!
interface Tunnel0
 ip unnumbered FastEthernet0/0
 qos pre-classify
 tunnel source Serial0/1
 tunnel destination 10.5.1.2
!
interface FastEthernet0/0
 ip address 10.1.1.1 255.255.255.0
 service-policy input mark-nbar
 duplex auto
 speed auto
!
interface Serial0/0
 bandwidth 384
 no ip address
 encapsulation ppp
 clockrate 384000
 ppp multilink
 multilink-group 1
!
interface Serial0/1
 bandwidth 768
 ip address 10.4.1.1 255.255.255.0
 service-policy output http-police
 encapsulation ppp
!
router ospf 1
 log-adjacency-changes
 network 10.0.0.0 0.255.255.255 area 0
!
ip http server
ip classless
!
ip access-list extended VoIP-RTCP
 permit udp any any range 16384 32767
```



```
!
ip access-list extended Voice-Control
 permit tcp any any eq 1720
 permit tcp any any range 11000 11999
 permit udp any any eq 2427
 permit tcp any any eq 2428
 permit tcp any any range 2000 2002
 permit udp any any eq 1719
 permit udp any any eq 5060
!
access-list 100 permit icmp any any echo
access-list 100 permit icmp any any echo-reply
access-list 101 permit tcp any any eq ftp
access-list 101 permit tcp any any eq ftp-data
access-list 102 permit tcp any any eq www
!
line con 0
line aux 0
line vty 0 4
 password cisco
 login
!
!
end
```

WG1R2:

```
hostname WG1R2
!
enable secret 5 $1$07qt$nKIz/sUIIRYMZ7urfJPtp1
!
ip subnet-zero
ip cef
!
 class-map match-all bulk
  match protocol ftp
 class-map match-any real-time
  match protocol rtp
  match protocol icmp
  match access-group name VoIP-RTCP
 class-map match-all match-www
  match access-group 102
 class-map match-all match-ftp
  match access-group 101
 class-map match-all web-outbound
  match protocol http
 class-map match-any ef-traffic
  match  dscp ef
  match access-group 100
 class-map match-all af21-traffic
  match  dscp af21
 class-map match-all af31-traffic
  match  dscp af31
 class-map match-all af11-traffic
  match  dscp af11
 class-map match-any mission-critical
  match protocol sqlnet
  match access-group name Voice-Control
 class-map match-all cs1-traffic
  match  dscp cs1
 class-map match-all interactive
```

```
 match protocol citrix
class-map match-all default
 match any
class-map match-all match-sw-be
 match  dscp default
class-map match-all match-sw-ef
 match  dscp ef
class-map match-any scavenger
 match protocol kazaa2
 match protocol napster
!
policy-map mark-nbar
 class real-time
  set dscp ef
 class mission-critical
  set dscp af31
 class interactive
  set dscp af21
 class bulk
  set dscp af11
 class scavenger
  set dscp cs1
  police cir 8000
    conform-action transmit
    exceed-action drop
 class class-default
  set dscp default
policy-map http-police
 class web-outbound
  police cir percent 50
    conform-action transmit
    exceed-action set-dscp-transmit cs1
    violate-action drop
policy-map mark-apps
 class match-ftp
  set dscp af11
 class match-www
  set dscp default
policy-map verify-mark
 class match-sw-ef
 class match-sw-be
policy-map llq-policy
 class ef-traffic
  priority 168
  compress header ip rtp
 class af31-traffic
  bandwidth remaining percent 40
  random-detect dscp-based
  random-detect dscp 10   34    40    10
 class af21-traffic
  bandwidth remaining percent 20
  random-detect dscp-based
  random-detect dscp 10   30    40    10
 class af11-traffic
  bandwidth remaining percent 13
  random-detect dscp-based
  random-detect dscp 10   26    40    10
  shape average 8000
 class cs1-traffic
  bandwidth remaining percent 2
  random-detect dscp-based
```



```
   random-detect dscp 10   22    40    10
  class class-default
   bandwidth remaining percent 25
   random-detect dscp-based
   random-detect ecn
   random-detect dscp 0   20    40    10
!
interface Multilink1
 bandwidth 384
 ip address 10.2.1.2 255.255.255.0
 service-policy output llq-policy
 ppp multilink
 ppp multilink fragment-delay 10
 ppp multilink interleave
 multilink-group 1
!
interface Tunnel0
 ip unnumbered FastEthernet0/0
 qos pre-classify
 tunnel source Serial0/1
 tunnel destination 10.4.1.1
!
interface FastEthernet0/0
 ip address 10.3.1.2 255.255.255.0
 service-policy input mark-nbar
 duplex auto
 speed auto
!
interface Serial0/0
 bandwidth 384
 no ip address
 encapsulation ppp
 ppp multilink
 multilink-group 1
!
interface Serial0/1
 bandwidth 768
 ip address 10.5.1.2 255.255.255.0
 service-policy output http-police
 encapsulation ppp
!
router ospf 1
 log-adjacency-changes
 network 10.0.0.0 0.255.255.255 area 0
!
ip http server
ip classless
!
ip access-list extended VoIP-RTCP
 permit udp any any range 16384 32767
!
ip access-list extended Voice-Control
 permit tcp any any eq 1720
 permit tcp any any range 11000 11999
 permit udp any any eq 2427
 permit tcp any any eq 2428
 permit tcp any any range 2000 2002
 permit udp any any eq 1719
 permit udp any any eq 5060
!
access-list 100 permit icmp any any echo
```

```
access-list 100 permit icmp any any echo-reply
access-list 101 permit tcp any any eq ftp
access-list 101 permit tcp any any eq ftp-data
access-list 102 permit tcp any any eq www
!
line con 0
line aux 0
line vty 0 4
 password cisco
 login
!
!
end
```

WG1S1:

```
hostname WG1S1
!
enable secret 5 $1$Yq48$E3tAlJjcYAP9qJpdmr0nu.
!
vlan 11
 name vlan11
!
vlan 21
 name vlan21
wrr-queue bandwidth 30 1 70 0
wrr-queue cos-map 1 0 1 2 3 4
wrr-queue cos-map 3 6 7
wrr-queue cos-map 4 5
!
class-map match-all callgen2
  match access-group 2
class-map match-all callgen1
  match access-group 1
!
policy-map mark-callgen
  class callgen1
    set ip dscp 46
  class callgen2
    set ip dscp 46
!
mls qos map cos-dscp 0 8 16 24 32 46 48 56
ip subnet-zero
vtp domain qos
vtp mode transparent
!
interface FastEthernet0/1
 switchport trunk allowed vlan 11,21
 switchport mode trunk
 no ip address
 service-policy input mark-callgen
 mls qos trust device cisco-phone
 mls qos trust cos
!
interface FastEthernet0/2
 switchport access vlan 11
 switchport mode access
 no ip address
!
interface FastEthernet0/3
```



```
 switchport access vlan 21
 switchport mode access
 no ip address
!
interface Vlan1
 no ip address
 no ip route-cache
 shutdown
!
access-list 1 permit 10.1.1.11
access-list 2 permit 10.3.1.11
!
line con 0
line vty 5 15
!
end
```

