IP QoS Mechanisms

QoS Mechanisms

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

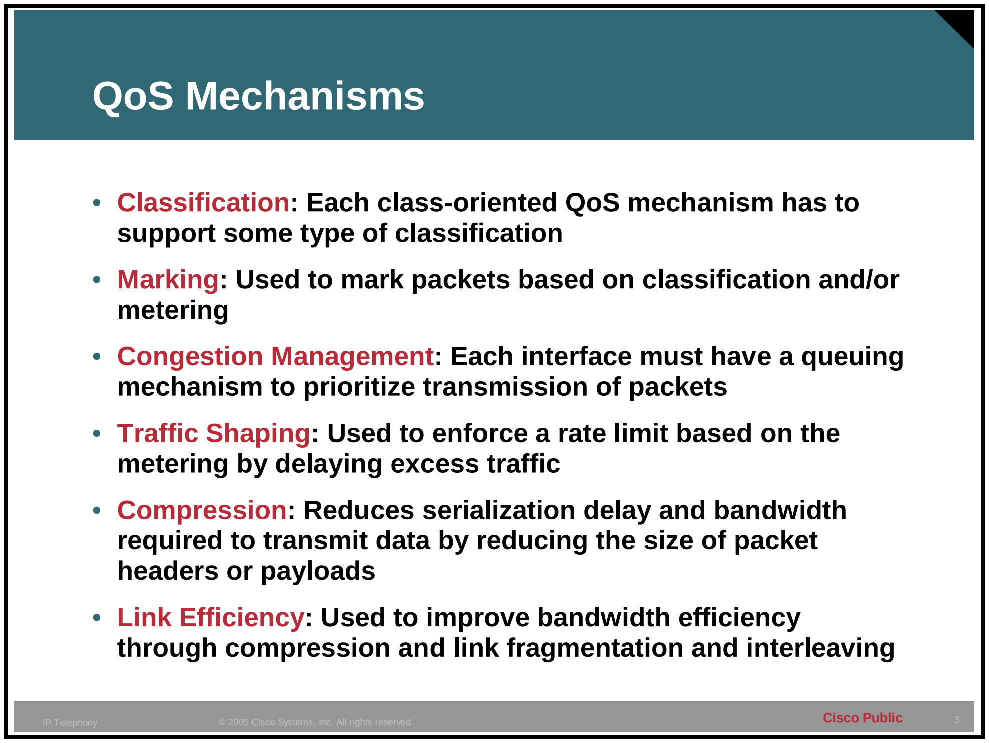

QoS Mechanisms

- **Classification: Each class-oriented QoS mechanism has to support some type of classification**
- **Marking: Used to mark packets based on classification and/or metering**
- **Congestion Management: Each interface must have a queuing mechanism to prioritize transmission of packets**
- **Traffic Shaping: Used to enforce a rate limit based on the metering by delaying excess traffic**
- **Compression: Reduces serialization delay and bandwidth required to transmit data by reducing the size of packet headers or payloads**
- **Link Efficiency: Used to improve bandwidth efficiency through compression and link fragmentation and interleaving**

IP Telephony © 2005 Cisco Systems, Inc. All rights reserved. Cisco Public 3

This slide shows the main categories of QoS tools used in IPTX implementations and describes in layman's 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.

Traffic conditioning mechanisms 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 identify where classification is commonly implemented in a network.

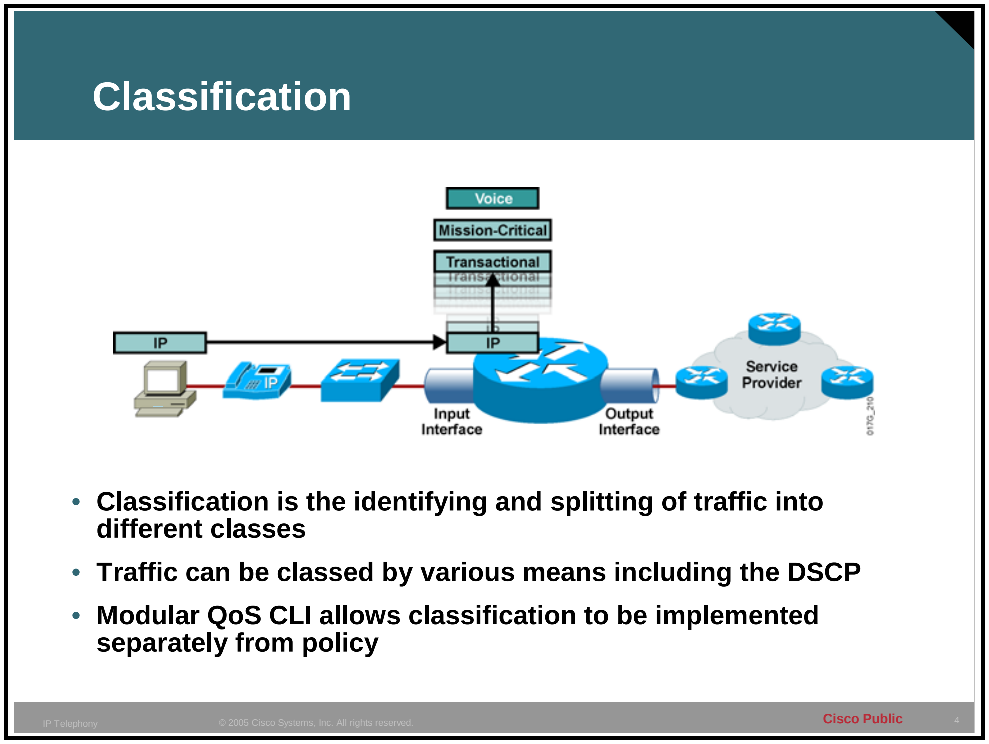

Classification

- Classification is the identifying and splitting of traffic into different classes
- Traffic can be classed by various means including the DSCP
- Modular QoS CLI allows classification to be implemented separately from policy

IP Telephony © 2005 Cisco Systems, Inc. All rights reserved. Cisco Public 4

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. Once 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 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.

Marking

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

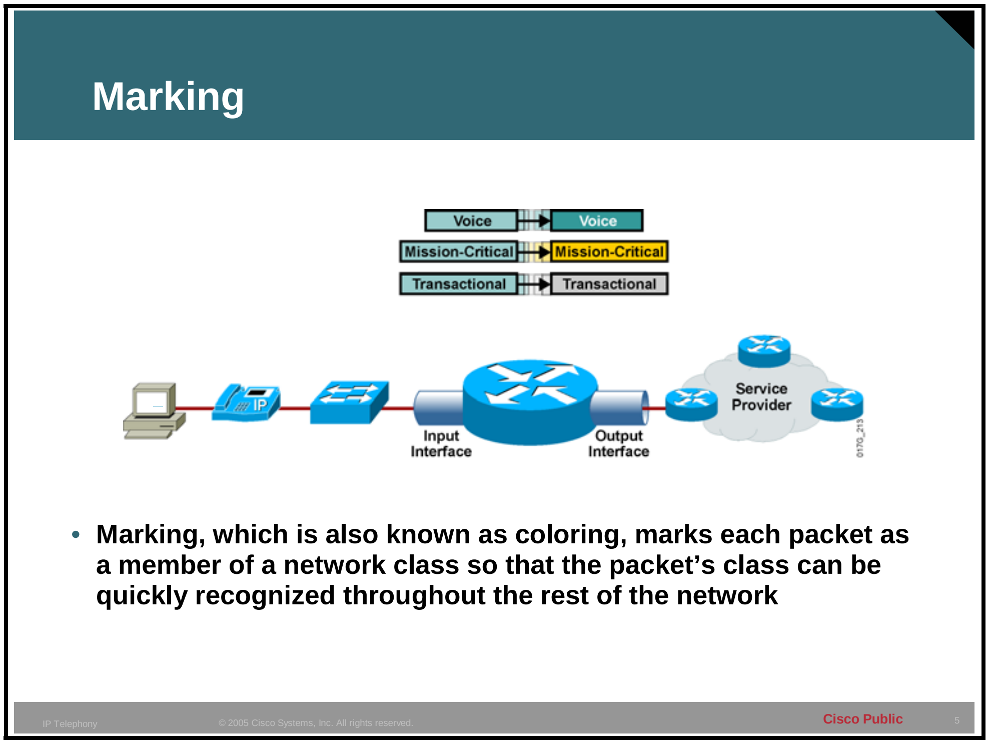

Marking

- **Marking, which is also known as coloring, marks each packet as a member of a network class so that the packet’s class can be quickly recognized throughout the rest of the network**

IP Telephony © 2005 Cisco Systems, Inc. All rights reserved. Cisco Public 5

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’s 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 which 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’s class such as CoS or Frame-Relay Discard Eligibility bit.

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.

Trust Boundaries

This topic describes concept of trust boundaries and how they are used with classification and marking.

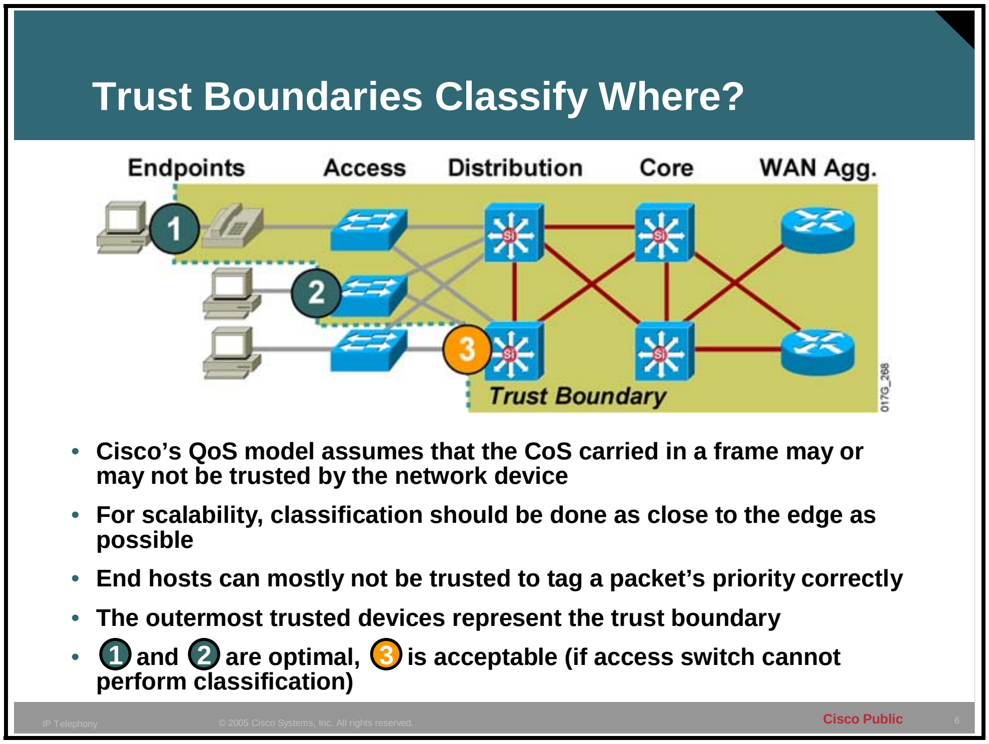

Trust Boundaries Classify Where?

- Cisco's QoS model assumes that the CoS carried in a frame may or may not be trusted by the network device
- For scalability, classification should be done as close to the edge as possible
- End hosts can mostly not be trusted to tag a packet's priority correctly
- The outermost trusted devices represent the trust boundary
- ① and ② are optimal, ③ is acceptable (if access switch cannot perform classification)

IP Telephony © 2005 Cisco Systems, Inc. All rights reserved. Cisco Public 6

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.

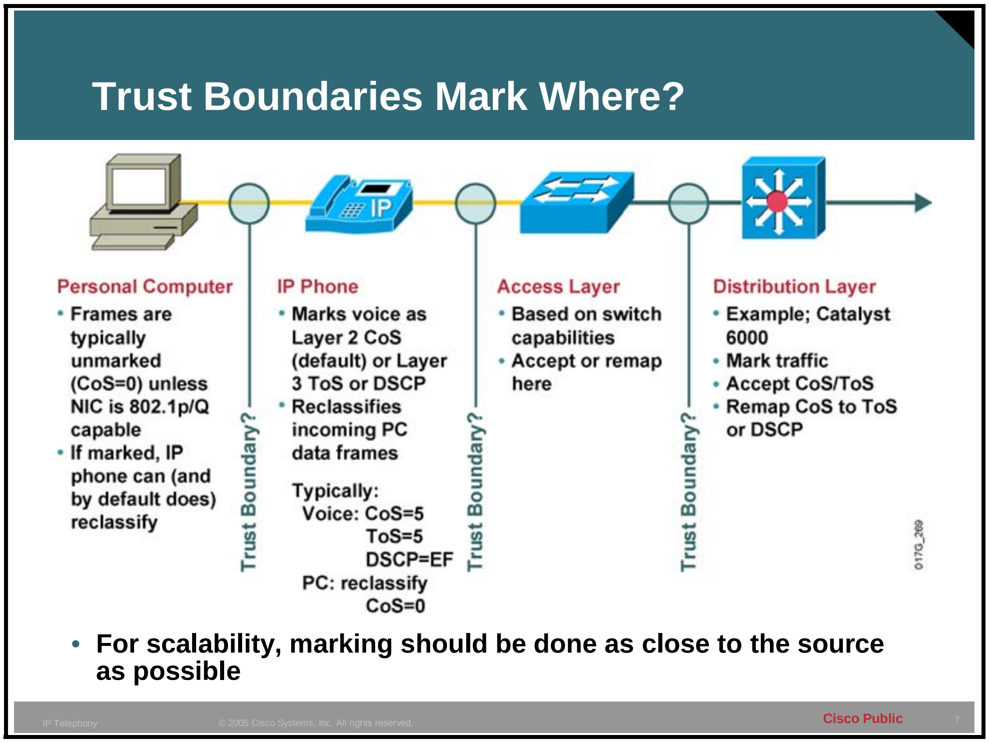

Trust Boundaries Mark Where?

Personal Computer

- Frames are typically unmarked (CoS=0) unless NIC is 802.1p/Q capable
- If marked, IP phone can (and by default does) reclassify

IP Phone

- Marks voice as Layer 2 CoS (default) or Layer 3 ToS or DSCP
- Reclassifies incoming PC data frames

Typically:
Voice: CoS=5
ToS=5
DSCP=EF
PC: reclassify
CoS=0

Access Layer

- Based on switch capabilities
- Accept or remap here

Distribution Layer

- Example; Catalyst 6000
- Mark traffic
- Accept CoS/ToS
- Remap CoS to ToS or DSCP

- **For scalability, marking should be done as close to the source as possible**

IP Telephony © 2005 Cisco Systems, Inc. All rights reserved. Cisco Public 7

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.

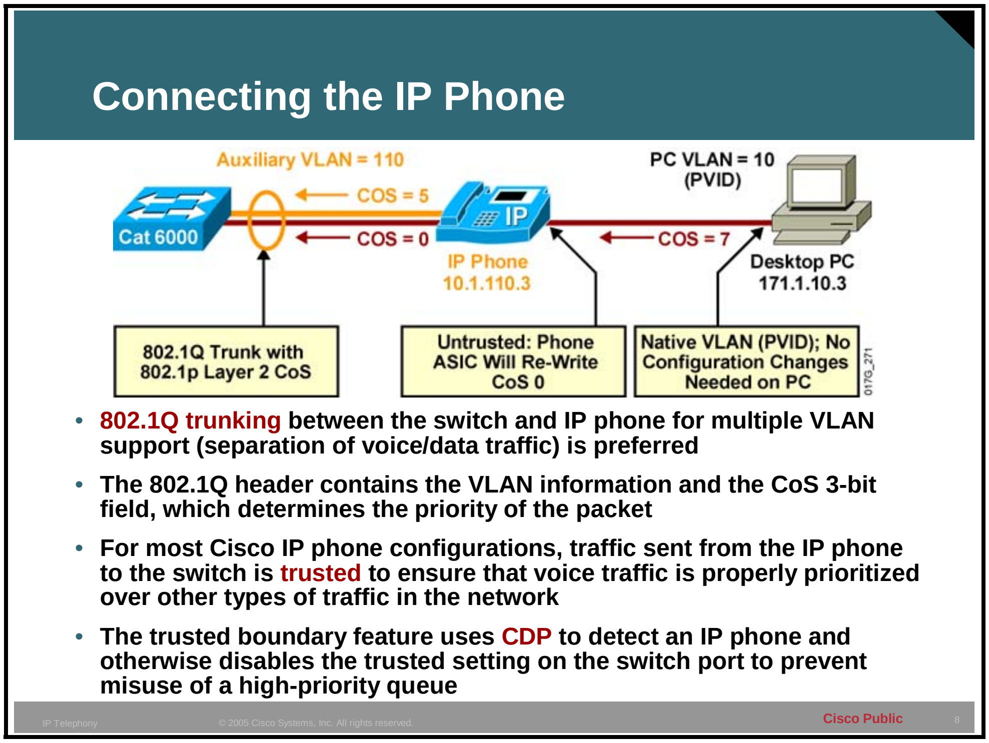
Auxiliary VLAN = 110
PC VLAN = 10
(PVID)
COS = 5
COS = 0
COS = 7
Cat 6000
IP
IP Phone
10.1.110.3
Desktop PC
171.1.10.3
802.1Q Trunk with 802.1p Layer 2 CoS
Untrusted: Phone ASIC Will Re-Write CoS 0
Native VLAN (PVID); No Configuration Changes Needed on PC
017G_271

Connecting the IP Phone

- **802.1Q trunking between the switch and IP phone for multiple VLAN support (separation of voice/data traffic) is preferred**
- **The 802.1Q header contains the VLAN information and the CoS 3-bit field, which determines the priority of the packet**
- **For most Cisco IP phone configurations, traffic sent from the IP phone to the switch is trusted to ensure that voice traffic is properly prioritized over other types of traffic in the network**
- **The trusted boundary feature uses CDP to detect an IP phone and otherwise disables the trusted setting on the switch port to prevent misuse of a high-priority queue**

IP Telephony © 2005 Cisco Systems, Inc. All rights reserved. Cisco Public 8

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.

Congestion Management

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

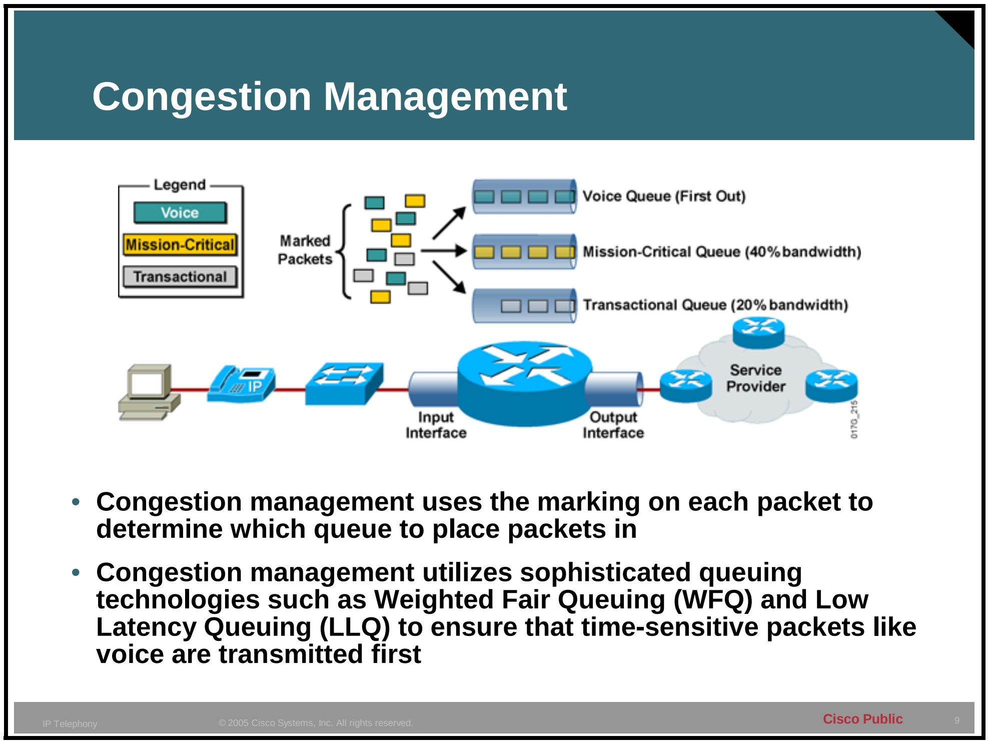

Congestion Management

- Congestion management uses the marking on each packet to determine which queue to place packets in
- Congestion management utilizes sophisticated queuing technologies such as Weighted Fair Queuing (WFQ) and Low Latency Queuing (LLQ) to ensure that time-sensitive packets like voice are transmitted first

IP Telephony © 2005 Cisco Systems, Inc. All rights reserved. Cisco Public 9

Congestion management mechanisms (queuing algorithms) use the marking on each packet to determine which queue to place packets in. 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 (first-in, first-out)
- PQ (priority queuing)
- CQ (custom queuing)
- WFQ (weighted fair queuing)
- CB-WFQ (class-based WFQ)
- LLQ (low latency queuing)

LLQ (low latency queuing) is now the preferred method. It is a hybrid (Priority Queuing and Class Based-Weighted Fair Queuing) queuing method developed specifically to meet the requirements of real time traffic such as voice.

Traffic Shaping

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

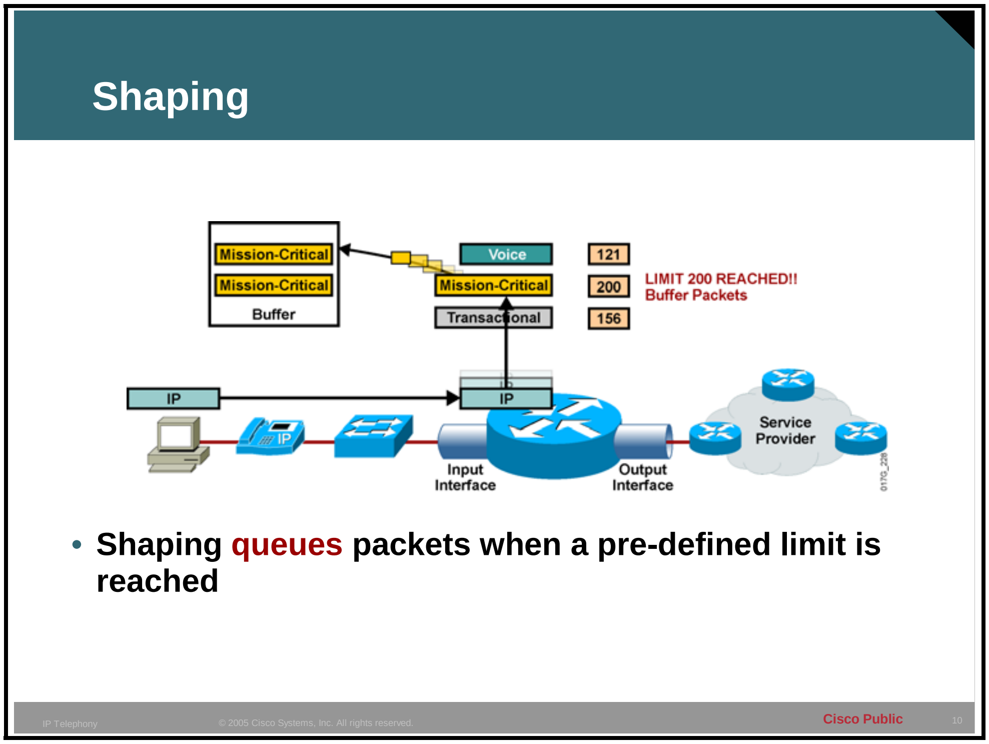

Shaping

- Shaping **queues** packets when a pre-defined limit is reached

IP Telephony © 2005 Cisco Systems, Inc. All rights reserved. Cisco Public 10

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.

Cisco’s 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).

Compression

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

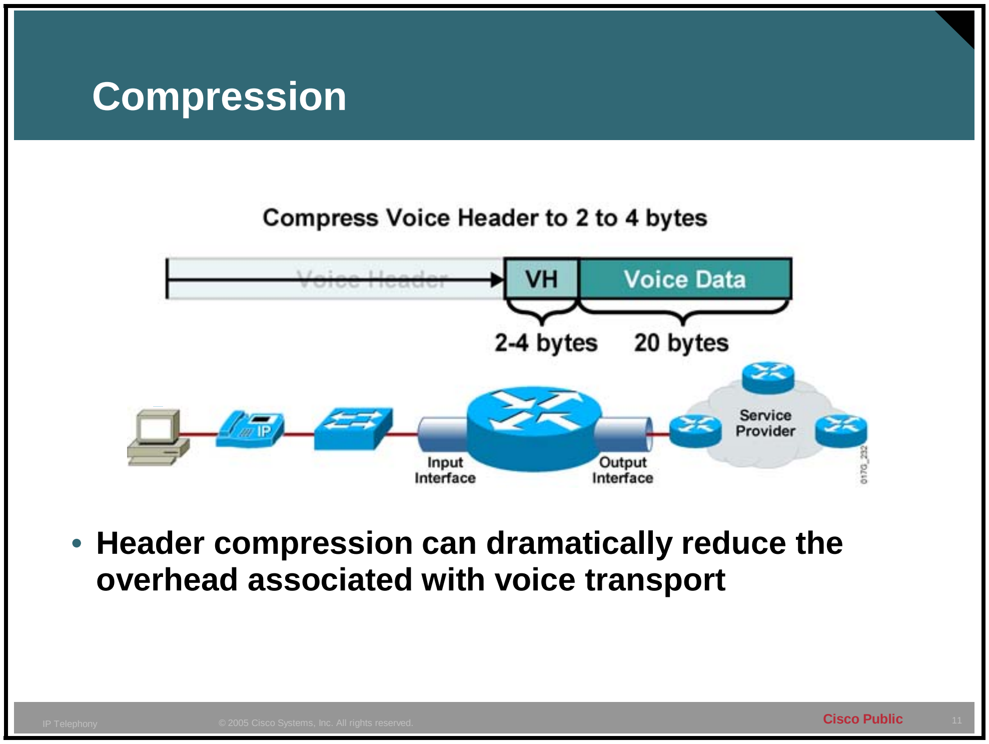

Compression

- **Header compression can dramatically reduce the overhead associated with voice transport**

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 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 transmitting 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 UDP header, and a 12-byte RTP header. By using cRTP, as shown in the graphic above, 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.

Link Fragmentation and Interleaving

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

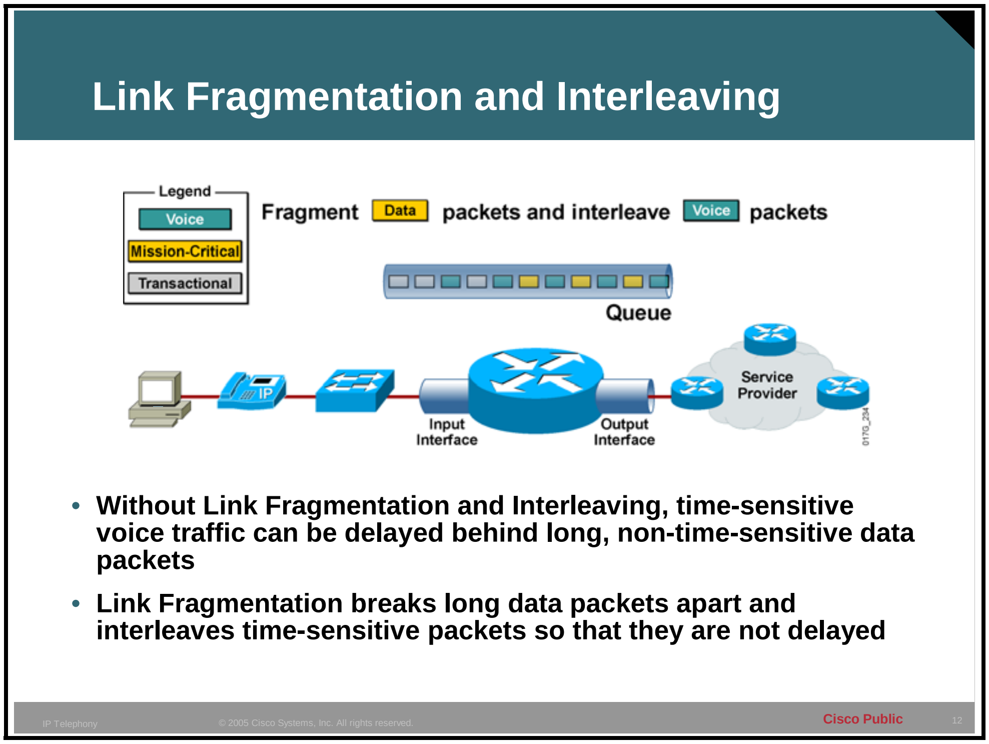

Link Fragmentation and Interleaving

- Without Link Fragmentation and Interleaving, time-sensitive voice traffic can be delayed behind long, non-time-sensitive data packets
- Link Fragmentation breaks long data packets apart and interleaves time-sensitive packets so that they are not delayed

IP Telephony © 2005 Cisco Systems, Inc. All rights reserved. Cisco Public 12

Interactive traffic, such as Telnet and Voice over IP, 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.

Link Fragmentation and Interleaving (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.