Cisco IPICS Infrastructure Considerations

This chapter contains information about infrastructure issues that you must be aware of when you deploy Cisco IPICS.

To access related documentation, refer to the following URLs:

- IP multicast:

  http://www.cisco.com/univercd/cc/td/doc/product/software/ios124/124cg/himc_c/index.htm

- Quality of Service:

  http://www.cisco.com/univercd/cc/td/doc/product/software/ios124/124cg/hqos_c/index.htm

- Voice over IP (VoIP):

  http://www.cisco.com/univercd/cc/td/doc/product/software/ios124/124tcg/vcl.htm

- Hoot 'n' Holler:

  http://www.cisco.com/en/US/netsol/ns340/ns394/ns165/ns70/networking_solutions_package.html

This chapter includes these topics:

WAN Considerations

To ensure the successful deployment of Cisco IPICS over a WAN, you must carefully plan, design, and implement the WAN. Make sure to consider the following factors:

- Delay—Propagation delay between two sites introduces 6 microseconds per kilometer. Other network delays may also be present.

- Quality of Service—The network infrastructure relies on QoS engineering to provide consistent and predictable end-to-end levels of service for traffic. QoS-enabled bandwidth must be engineered into the network infrastructure.

- Jitter—Varying delay that packets incur through the network as a result of processing, queue, buffer, congestion, or path variation delay. Jitter for the multicast voice traffic must be minimized using QoS features. For related information, see the "Quality of Service" section on page 3-8.

- Packet loss and errors—The network should be engineered to provide sufficient prioritized bandwidth for all voice traffic. Standard QoS mechanisms must be implemented to avoid congestion and packet loss. For related information, see the "Quality of Service" section on page 3-8.

- Bandwidth—Provision the correct amount of bandwidth between each site for the expected call volume. This bandwidth is in addition to bandwidth for other applications and traffic that share the network. The provisioned bandwidth must have QoS enabled to provide prioritization and scheduling for the different classes of traffic. In general, the bandwidth should be over-provisioned and under-subscribed.

- PMCs using remote location—PMCs that choose the remote location use SIP to set up the connection to the RMS. Because SIP is sensitive to latency and excessive delays, this type of connection may experience timeouts and call setup failures. If you experience connectivity or audio quality issues when you use a PMC that is deployed over a WAN, try to isolate the problem by testing a remote PMC connection that is not deployed over a WAN. If you cannot duplicate the problem on the PMC that is not deployed over the WAN, Cisco recommends that you assess the WAN infrastructure to determine if excessive latency or resource limitations in the WAN are causing the problems.

Multicast Routing

Cisco supports the Protocol Independent Multicast (PIM) routing protocol for both sparse mode (SM) and dense mode (DM). However, because of its periodic broadcast and prune mechanism, DM PIM is not recommended for production networks.

Cisco recommends using bidirectional PIM for Cisco IPICS. Bidirectional PIM is an extension of the PIM suite of protocols that implements shared sparse trees with bidirectional data flow. In contrast to PIM-sparse mode, bidirectional PIM avoids keeping source-specific states in a router and allows trees to scale to an arbitrary number of sources while requiring only minimal additional overhead.

The shared trees that are created in PIM SM are unidirectional. Therefore, a source tree must be created to bring a data stream to the rendezvous point (RP), which is the root of the shared tree. Then the data can be forwarded down the branches to receivers. In the unidirectional mode, source data cannot flow up the shared tree toward the RP.

In bidirectional mode, traffic is routed only along a bidirectional shared tree that is rooted at the RP for the group. In bidirectional PIM, the IP address of the RP acts as the key to having all routers establish a loop-free spanning tree topology that is rooted in that IP address. The IP address does not need to be a router. It can be any unassigned IP address on a network that can be reached throughout the PIM domain.

Figure 3-1 shows bidirectional shared tree. In this example, data from the source can flow up the shared tree (*, G) toward the RP and then down the shared tree to the receiver. There is no registration process so source tree (S, G) is created.

Figure 3-1   Bidirectional Shared Tree

Bidirectional PIM is derived from the mechanisms of PIM SM and has many of the same shared tree operations. Bidirectional PIM also has unconditional forwarding of source traffic toward the RP upstream on the shared tree, but no registering process for sources, as provided by PIM SM. These modifications are necessary and sufficient to allow forwarding of traffic to all routers based only on the (*, G) multicast routing entries. Bidirectional PIM eliminates any source-specific state and allows scaling to an arbitrary number of sources.

In a Cisco IPICS deployment, bidirectional PIM provides scalability in the following ways:

- Forwarding traffic based on the shared tree (*, G)—This functionality helps scale the multicast routing table by creating a single routing entry per channel. In SM, a routing entry is created per group and per source. For example, if a channel has 100 participants, it will have 101 multicast routing entries in the routing table. With bidirectional PIM, only a single multicast routing entry in the routing table is created, regardless of the number of participants.

- Basing the Reverse Path Forwarding (RPF) decision on the route to the RP—In SM, RPF decisions about (S, G) entries are based on the source address of the flow, and for bidirectional (*, G), RPF decisions are based on the RP. This functionality eliminate the need to configure hundreds of ip mroute entries to force multicast traffic on the Cisco IPICS Permanent Virtual Circuit (PVC). With bidirection, forcing the multicast traffic on the Cisco IPICS PVC is achieved by tuning the unicast routing protocol to prefer the Cisco IPICS PVC as the best route to reach the RP.

Bandwidth Planning

To ensure sufficient bandwidth for the operation of Cisco IPICS, there are several issues to consider as you plan and deploy your network. These issues include:

- Codec used for VoIP—See the "Codecs" section on page 3-4

- The number of voice streams that will be mixed—See the "Mixing Voice Streams" section on page 3-8

In addition, you should consider the guaranteed bandwidth that is available on the VoIP network. Make sure to take into account both LAN and WAN bandwidth and to consider factors such as Frame Relay, Committed Information Rate (CIR) or Asynchronous Transfer Mode Peak Cell Rate (ATM PCR), Sustained Cell Rate, and burst. For additional information see the "Quality of Service" section on page 3-8.

Codecs

Cisco IPICS uses either the G.711 or G.729a codec. This section provides the following information about codecs:

- Choosing a Codec, page 3-4
- Calculating Codec Bandwidth Use, page 3-5

Choosing a Codec

When choosing a codec for Cisco IPICS, consider the issues that are described in Table 3-1.

Table 3-1   Codec Considerations

	G.711	G.729a
Delay	• Total delay is 25ms less per sample than for G.729a. • Transcoding increases delay.	• Total delay is 25ms greater per sample than for G.711. • Some Cisco IPICS deployments that use G.729a require additional transcoding to convert the G.729a streams to G.711 stream for mixing. This additional DSP function increases delay significantly.
Voice Quality	• Assuming that good VoIP conditions exist, delivers a MOS of 4.1 with a high degree of consistency. • Does tandem well, so no voice quality degradation results from transcoding.	• Assuming that good VoIP conditions exist, typically delivers a Mean Opinion Score (MOS) of 3.7 and can cause more unpredictable results than G.711. • Does not perform as well as G.711 under packet loss conditions. For example, a 3% packet loss rate can have a larger effect on voice quality then a similar packet loss rate under G.711. • Does not tandem as well as G.711. • Transcoding decreases voice quality from a MOS of 3.7 to 3.2.
Bandwidth	• Typically consumes 3 times more bandwidth than G.729a.	• Offers bandwidth savings over G.711. • A Cisco IPICS deployment that connects sites via a WAN may use G.729a to reduce WAN bandwidth, which also may reduce WAN costs.

Calculating Codec Bandwidth Use

This section explains how to calculate bandwidth use for codecs.

By default, Cisco IOS sends all VoIP traffic (that is, media traffic that uses RTP) at a rate of 50 packets/second. In addition to the voice sample, each packets includes an IP, UDP, and Real-time Transport Protocol (RTP) header, which adds 40 bytes to the packet. Layer 2 headers (such as Frame Relay, Point-to-Point Protocol, Ethernet) also add bytes to each packet.

The amount of bandwidth that is consumed by a VoIP call depends on the codec that is used, and can be calculated as follows. Make sure to also add the appropriate number of bytes for the Layer 2 header to determine the actual bandwidth that is consumed.

G.729a (8K CS-ACELP)

50 packets/second

20ms samples / packet = 20 bytes

AP/UDP/RTP headers/packet = 40 bytes

(20 bytes [payload] + 40 bytes [headers]) * 50 packets/second = 3,000 bytes * 8 bits = 24 kbps

G.711 (64K PCM)

50 packets/second

20ms samples / packet = 160 bytes

AP/UDP/RTP headers/packet = 40 bytes

(160 bytes [payload] + 40 bytes [headers]) * 50 packets/second = 10,000 bytes * 8 bits = 80 kbps

Table 3-2 shows sample bandwidth consumption. In this table,

- The examples assume a payload size (bytes) of 20 ms samples per packet with 50 packets per second.
- The value n is equal to the number of voice streams in a session.
- The encompassed bandwidth includes IP/UDP/RTP headers (40 bytes) in the bandwidth calculation.
- Compressed RTP (cRTP) reduces the IP/UDP/RTP headers to between 2 and 4 bytes per packet. The calculation of compressed bandwidth uses 4 bytes for a compressed IP/UDP/RTP headers per packet.
- Make sure to add the appropriate number of bytes for the Layer 2 header to determine the actual bandwidth consumed.

Table 3-2   Sample Bandwidth Usage

Codec	Payload Size (bytes)	Bandwidth/Voice Stream (kbps)		RTCP Bandwidth per Cisco IPICS Session (kbps)	Example: 1 Voice Stream in a Session (kbps)	
		Uncompressed	Compressed		Uncompressed	Compressed
G.729a	20	24	9.6	3.6	27.6	13.2
G.711	160	80	65.6	12.0	92.0	77.6

According to RFC 1889 (*RTP: A Transport Protocol for Real-Time Applications*), the RTCP traffic for any RTP stream is limited to a maximum of 5% of the voice stream (RTP + RTCP). This limitation applies to the three streams that participate in a Cisco IPICS session. Therefore, the RTCP Bandwidth per Cisco IPICS Session is calculated by multiplying the bandwidth per voice stream by 3 and then multiplying that product by 0.05.

When you design a Cisco IPICS network within a campus network, you should not run into any bandwidth-related issues because IP multicast is used to replicate a voice stream and map it to an IP multicast group, in which RMS resources are not used. When remote users connect over a WAN that is not multicast enabled, the RMS converts a multicast stream to an IP unicast stream, which conserves bandwidth on the WAN. When the IP unicast voice stream arrives at the RMO, the RMS converts the IP unicast stream to an multicast stream. When the voice streams traverse a WAN, the RMS resources are used. In this scenario the RMS gateways are configured for M1:U12:M2 unicast connection trunks. For additional information, see the "M1:U12:M2 Connection Trunks" section on page 5-13.

As an example of the bandwidth affect of PMCs that are deployed across a WAN, assume that 40 PMC users are communicating over a WAN. Although each multicast voice stream is converted to an IP unicast voice stream so that it can traverse as a IP unicast stream, this scenario could still require substantial bandwidth, depending on the number of channels and the codec type that is used by each PMC. In this example, the bandwidth requirements are as follows:

G.729a

40 PMC Users x 8 Channels per User x (24 kbps + 3.6kbps) = 8.832 kbps

G.711

40 PMC Users x 8 Channels per User x (80 kbps + 12 kbps) = 29,440 kbps

cRTP, Variable-Payload Sizes and Aggressive VAD

There are several methods that you can use to modify the bandwidth that a call consumes. These methods include the following:

- RTP Header Compression, page 3-6
- Adjustable Byte Size of the Voice Payload, page 3-7
- Aggressive Voice Activity Detection, page 3-7

RTP Header Compression

As described in the "Codecs" section on page 3-4, IP/UDP/RTP headers add 40 bytes to each packet. However, a packet header is typically unchanged throughout a call. You can enable cRTP for VoIP calls, which reduces the size of IP/UDP/RTP headers to 2 to 4 bytes per packet.

For detailed information about cRTP, refer to *Understanding Compression (Including cRTP) and Quality of Service*, which is available at this URL:

http://www.cisco.com/en/US/tech/tk543/tk762/technologies_tech_note09186a0080108e2c.shtml

Adjustable Byte Size of the Voice Payload

You can control the size of the voice payload that is included in each Cisco IPICS voice packet. To do so, use the bytes parameter in a VoIP dial peer. For example:

```
dial-peer voice 1 voip
destination-pattern 4085551234
codec g729r8 bytes 40
session protocol multicast
session target ipv4:239.192.1.1:21000
```

Modifying the number of bytes per packet changes the number of packets that are sent per second. You can calculate the number of packets that are sent per second as shown in these examples:

G.729a codec, with default 20byte payload/packet

Codec rate: 8,000 bits/second * 8 bits = 1,000 byes/second

Sampling interval: 10 ms

Default payload size: 20 bytes/packet (2 samples/packet)

1,000 bytes/sec/20bytes/pkt = 50 packets/sec

G.729a codec, with 40 bytes defined in VoIP dial-peer

Codec rate: 8,000 bits/second * 8 bits = 1,000 byes/second

Sampling interval: 10 ms

Payload size: 40 bytes/packet

1,000 bytes/sec/40 bytes/pkt = 25 packets/sec

Note Increasing payload size increases the delay per sample by the same amount. For example, increasing payload size from 20 ms to 40 ms increases the delay per sample by 20 ms.

Aggressive Voice Activity Detection

Voice Activation Detection (VAD) is a mechanism that allows a digital signal processor (DSP) to dynamically sense pauses in conversation. When such pauses occur, no VoIP packets are sent into the network. VAD can reduce the amount of bandwidth used for a VoIP call by up to 50%.

Although VAD conserves bandwidth in VoIP, it disrupts and marginalizes Cisco IPICS signaling, which is used for LMR and PTT packet streams. Be aware of this issue if you use VAD in a Cisco IPICS deployment.

When configuring LMR gateway ports, VAD should not be used if the radio supports Carrier Operated Relay (COR) or Carrier Operated Squelch (COS). Radios that support COR/COS signaling can provide hardwired signaling to the LMR port to start generating packets. Using COR/COS gating is an efficient way to control the audio input and to avoid the possibility of dropping short burst of voice data that may fall below the VAD activation values.

Each voice port has different environmental noises and different users, which can cause a wide variation in noise and speech levels. Conventional VAD can handle these variations, but it is designed for unicast. Conventional VAD usually prefers over-detection to under-detection, as good voice quality is typically given precedence over bandwidth conservation. But in a multicast environment, over-detection and under-detection are not desirable because they degrade voice quality.

Aggressive VAD can be used in a multicast environment to avoid over-detection. With aggressive VAD, when a DSP detects signals with an unknown signal-to-noise ratio (SNR), the DSP does not transmit spurious packets. With conventional VAD, when the DSP detects signals with an unknown SNR, the DSP continues to transmit packets, which can cause unwanted traffic to take over all slots that are available for voice streams.

You can enable aggressive VAD by enabling the `vad aggressive` configuration setting under a dial peer as follows:

```
dial-peer voice 10 voip
destination-pattern 111
session protocol multicast
session target ipv4:239.192.1.1:21000
vad aggressive
```

Mixing Voice Streams

As described in the "Virtual Talk Groups" section on page 2-10, the DSPs in a Cisco IPICS deployment can mix up to three voice streams. However, the DSPs do not perform a summation function. For example, if three G.729a streams (24K each with headers) are received by a router or gateway, the mixed stream would consume 72K bandwidth. Even though each user in a VTG or a channel in the VTG receives a single mixed audio stream, the DSP does not send a single 24 K stream.

It is important to consider this issue when you plan bandwidth in a Cisco IPICS network. It is especially important when planning WAN bandwidth, which can be more expensive and less available then LAN bandwidth.

Because the Cisco Hoot 'n' Holler feature mixes up to three voice streams at a time, you do not need to provision voice bandwidth for more than three times the per-call bandwidth for each WAN site that includes routers with the Cisco Hoot 'n' Holler feature.

Note   An audio channel that is mixed through a VTG experiences an additional 60 ms of delay.

Quality of Service

There are several QoS features that should be enabled so that a Cisco IPICS deployment can deliver toll-quality VoIP QoS. This section provides an overview of these feature for Point-to-Point Protocol (PPP) and Frame Relay WAN topologies and for deployments on LAN media.

This section includes these topics:

QoS Overview

QoS provides consistent voice latency and minimal packet loss. The following recommendations apply to QoS in campus LAN and WAN environments:

- Classify voice RTP streams as expedited forwarding (EF) or IP precedence 5 and place them into a priority queue on all network elements

- Classify voice control traffic as assured forwarding 31 (AF31) or IP precedence 3 and place it into a second queue on all network elements

As you design a VoIP network to deploy real-time applications such as Cisco IPICS, consider the following issues, which can affect voice quality:

- Packet loss—Causes voice clipping and skips. The industry-standard codec algorithms used in DSPs can correct for up to 30 ms of lost voice. Cisco VoIP technology uses 20 ms samples of voice payload per VoIP packet. Therefore, for the codec correction algorithms to be effective, only a single packet can be lost during any time. Packet loss can be a significant problem for real-time applications because they are not designed to retransmit packets.

- Delay—Causes either voice quality degradation due to the end-to-end voice latency or packet loss if the delay is variable. If the delay is variable, such as queue delay in bursty data environments, there is a risk of jitter buffer overruns at the receiving end. Longer delays can cause buffer overflow and underflow, and unnatural pauses in human conversations. Because Cisco IPICS supports a PTT service, the typical one-way delay requirement of 150 ms as recommended in the International Telecommunication Union (ITU) G.114 specification does not directly apply. PTT users are aware of radio protocol, so a more re reasonable delay is 400 ms as outlined in the ITU G.173 specification.

- Jitter—Variable delay. While some delay is acceptable, delay that constantly changes can cause inconsistent and inefficient DSP buffering. It also can cause inconsistent voice quality.

- Ability to Prioritize VoIP traffic—Involves the use of queuing techniques, such as IP RTP Priority and Low-Latency Queuing, that are available in Cisco IOS.

- Ability to make VoIP traffic best fit the LAN or WAN network—Involves making sure that small VoIP packets do not get delayed behind large data packets (an event called *serialization*).

If network are designed and built to provide low delay, limited jitter, and limited packet loss, real-time applications such as Cisco IPICS solution can be successful.

IOS Queuing Techniques

Cisco IOS provides a wide variety of QoS features. The following features are particularly useful for a Cisco IPICS deployment:

For more detailed documentation about IP RTP Priority, refer to the "Congestion Management Overview" chapter in *Cisco IOS Quality of Service Solutions Configuration Guide*, which is available at this URL:

http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122cgcr

IP RTP Priority

IP RTP Priority can be applied to point-to-point links and to Frame Relay PVCs. It allows you to provision a fixed amount of bandwidth (in Kb) that is always available for Cisco IPICS packets. If there are no Cisco IPICS packets present in the network (that is, nobody is speaking), the bandwidth is available to other data applications. This predefined amount of bandwidth is serviced as a strict priority-queue within the overall structure of Weighted-Fair Queuing (WFQ). The entrance criteria to this priority queue is a range of UDP ports that are used by Cisco IPICS to send IP packets.

Cisco IPICS uses the UDP port that is selected on the VoIP dial peer, and the next sequential port. The ports can range from 21000 through 65534. The first port must be an even number within this range.

The following example shows the UDP port (24100) defined in the VoIP dial-peer, so the range for the IP RTP Priority is 24100-24101:

```
dial-peer voice 1 voip
destination-pattern 1111
session protocol multicast
codec g711ulaw
session target ipv4:239.10.0.100:24100
!
interface serial 0/0
ip address 10.1.1.1
ip rtp priority 24100 2 64
```

Low Latency Queuing

Low-Latency Queuing (LLQ) applies to point-to-point links and to Frame Relay PVCs. LLQ creates a strict priority queue, as does IP RTP Priority, but LLQ applies the strict priority queue as a service-class within Class-Based Weighted Fair Queueing (CBWFQ). The functionality of fixed allocation but dynamic usage is again similar to IP RTP Priority.

A primary difference between IP RTP Priority and LLQ is that LLQ allows the use of access control lists (ACLs) as the entrance criteria to the priority queue. This capability provides you with flexibility in determining what types of traffic are allowed into the priority queue.

The following example shows how LLQ is used to prioritize Cisco IPICS traffic:

```
access-list102 permit udp host 10.1.1.1 host 239.10.0.100 range 24100 24101
!
class-map voice
match access-group 102
!
policy-map policy1
class voice
priority 50
!
multilink virtual-template 1
!
interface virtual-template 1
ip address 172.17.254.161 255.255.255.248
no ip directed-broadcast
no ip mroute-cache
service-policy output policy1
ppp multilink
ppp multilink fragment-delay 20
ppp multilink interleave
!
interface serial 2/0
bandwidth 256
```

```
no ip address
no ip directed-broadcast
encapsulation ppp
no fair-queue
clockrate 256000
ppp multilink
multilink-group 1
```

QoS with Frame Relay

If you deploy Cisco IPICS in a Frame Relay network, be aware that Frame Relay does not inherently provide QoS. Frame Relay is a best-effort service that expects upper-layer applications to handle retransmissions that occur because of packet loss in the Frame Relay cloud.

Frame Relay typically provides the following parameters:

- Committed Information Rate (CIR)—Amount of bandwidth that the Frame Relay carrier guarantees to be available at all times for a particular PVC. The carrier does not make any guarantees for packets sent above CIR.

- Burst—Maximum amount of data that the Frame Relay carrier allows to be sent on a particular PVC.

To offer the QoS over Frame service, carriers use a technique called *over-provisioning bandwidth*, in which they sell more bandwidth then they can provide at a particular time. This technique works because not all Frame Relay customers require all available bandwidth at one time.

Some Frame Relay carriers also guarantee a Frame Relay network that is always available and that will not drop any customer packets.

A Frame Relay carrier employs a variety of methods to offer a CIR + Burst service, including the following:

- Marking packets with discard eligible (DE) or dropping the packets—Because real-time applications such as VoIP use UDP for transport, there is no mechanism for packets to be retransmitted. This situation is not a problem for VoIP because users would not want to hear a dropped word later in a sentence. Packet loss is generally not acceptable for real-time VoIP applications because it can result in choppy audio and garbled speech.

- Buffering all packets above the CIR—Eliminates lost packets, but can introduce jitter and delay because of the depth or rate at which the Frame Relay switches empty buffers.

Table 3-3 summarizes key recommendations when deploying Cisco IPICS on a network with Frame Relay.

Table 3-3   Recommendations when Deploying Cisco IPICS with Frame Relay

Recommendation	Technique	Comments
To avoid introducing packet loss or jitter into a Cisco IPICS network, make sure that traffic that exceeds the CIR is not sent into a Frame Relay network.	Use the IOS Frame Relay Traffic Shaping (FRTS) feature.	Allows a router to police traffic on a per-PVC basis so that it does not send any traffic above the CIR.
In a Frame Relay environment, make sure that packets that are sent across a WAN link do not exceed the Committed Information Rate (CIR)	Enable the FRF.12 feature in the Frame network.	FRF.12 is a Frame-Relay-Forum Implementation Agreement that specifies how to fragment and reassemble packets on a Frame Relay network at Layer 2 of the OSI model. By fragmenting large data packets, the smaller Cisco IPICS packets will not be delayed or subject to serialization, which helps to eliminate delay and jitter of the Cisco IPICS packets. Because the fragmentation and reassembly is done at Layer 2 of the OSI model, it does not adversely effect any upper-layer protocols (such as IPX or Appletalk or IP with DNF bits set) that do not handle fragmentation.
Implement a queuing technique that provides strict priority to Cisco IPICS packets.	Use a technique such as Low-Latency Queuing (LLQ)	The LLQ feature brings strict priority queuing to the Class-Based Weighted Fair Queuing (CBWFQ) method. 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.

Example

Consider a Cisco IPICS Frame Relay network with the following characteristics:

- Three routers connected through 64 K Frame Relay PVCs in a hub and spoke topology, with Router-1 being the hub.

- All routers configured to traffic-shape their data and voice on the WAN to CIR, and all routers using IP RTP Priority to guarantee QoS for the Cisco IPICS packets.

- Frame Relay broadcast-queue enabled on the serial interfaces.

- One Cisco IPICS Channel configured.

Because the broadcast queue is only 40 packets deep by default and Cisco IPICS components (PMC, Cisco Unified IP Phone, RMS) transmits packets at 50 packets/second, the broadcast-queue must be set to prevent voice packets from dropping and to maintain voice quality. The recommended setting for the broadcast-queue is 64 8000 25 (64 queue size, 8,000 bytes per second (64,000 bps), and 25 packets per second).

Frame Relay Broadcast Queue

Broadcast queue is a feature that is used in medium and large IP or IPX networks where routing and service access point (SAP) broadcasts must flow across a Frame Relay network. The broadcast queue is managed independently of the normal interface queue, has its own buffers, and has a configurable size and data rate.

To enable broadcast queue, use this interface command:

frame-relay *broadcast-queue size byte-rate packet-rate*

A broadcast queue is given a maximum transmission rate (throughput) limit, which is measured in bytes per second and packets per second. The queue is serviced to ensure that only this maximum rate is provided. Because the broadcast queue has priority when transmitting at a rate below the configured maximum, it has a guaranteed minimum bandwidth allocation. The two transmission rate limits are intended to avoid flooding the interface with broadcasts. The actual limit in any second is the first rate limit that is reached. Given the transmission rate restriction, additional buffering is required to store broadcast packets.

The broadcast queue can be configured to store a large number of broadcast packets. You should set the queue size to a value that avoids loss of broadcast routing update packets. The exact size depends on the protocol being used and the number of packets required for each update. To be safe, the queue size should be set so that one complete routing update from each protocol and for each data-link connection identifier (DLCI) can be stored. As a general rule, start with 20 packets per DLCI. The byte rate should be less than both of the following:

- n/4 times the minimum remote access rate (measured in bytes per second), where n is the number of DLCIs to which the broadcast must be replicated

- 1/4 the local access rate (measured in bytes per second)

The packet rate is not critical if the byte rate is set conservatively. In general, the packet rate should be set assuming 250-byte packets. The frame-relay broadcast-queue command defaults are as follows:

- Size—64 packets

- Byte-rate—256000 bytes per second

- Packet-rate—36 packets per second

The following configuration is an example of a Frame Relay connection with an ear and mouth (E&M) port:

```
Router-1 (Hub Router)

hostname FR-1
!
ip multicast-routing
!
voice class permanent 1
 signal timing oos timeout disabled
 signal keepalive disabled
 signal sequence oos no-action
!
interface Vif1
ip address 1.1.1.1 255.255.255.0
ip pim sparse-dense-mode
!
router rip
network 1.1.1.0
network 5.5.5.0
network 5.5.6.0
!
```

```
interface Serial0/0
no frame-relay broadcast-queue
encapsulation frame-relay
frame-relay traffic-shaping
frame-relay broadcast-queue 64 8000 250
!
interface Serial0/0.1 point-to-point
ip address 5.5.5.1 255.255.255.0
ip pim sparse-dense-mode
frame-relay class ipics
frame-relay interface-dlci 100
frame-relay ip rtp header-compression
!
interface Serial0/0.1 point-to-point
ip address 5.5.6.1 255.255.255.0
ip pim sparse-dense-mode
frame-relay class ipics
frame-relay interface-dlci 100
frame-relay ip rtp header-compression
!
map-class frame-relay ipics
frame-relay cir 128000
frame-relay bc 1280
frame-relay mincir 128000
no frame-relay adaptive-shaping
frame-relay fair-queue
frame-relay fragment 160
frame-relay ip rtp priority 16384 16384 128
!
voice-port 1/0/0
 connection trunk 111
 operation 4-wire
!
dial-peer voice 1 voip
 destination-pattern 111
 voice class permanent 1
 session protocol multicast
 session target ipv4:239.111.0.0:21000
 ip precedence 5
!

Router-2 (Spoke Router)

hostname FR-2
!
ip multicast-routing
!
voice class permanent 1
 signal timing oos timeout disabled
 signal keepalive disabled
 signal sequence oos no-action
!
interface Vif1
ip address 1.1.2.1 255.255.255.0
ip pim sparse-dense-mode
!
router rip
network 1.1.2.0
network 5.5.5.0
!
interface Serial0/0
no frame-relay broadcast-queue
encapsulation frame-relay
frame-relay traffic-shaping
```

```
frame-relay broadcast-queue 64 8000 250
!
interface Serial0/0.1 point-to-point
ip address 5.5.5.2 255.255.255.0
ip pim sparse-dense-mode
frame-relay class ipics
frame-relay interface-dlci 100
frame-relay ip rtp header-compression
!
map-class frame-relay ipics
frame-relay cir 128000
frame-relay bc 1280
frame-relay mincir 128000
no frame-relay adaptive-shaping
frame-relay fair-queue
frame-relay fragment 160
frame-relay ip rtp priority 16384 16384 128
!
voice-port 1/0/0
 connection trunk 111
 operation 4-wire
!
dial-peer voice 1 voip
 destination-pattern 111
 voice class permanent 1
 session protocol multicast
 session target ipv4:239.111.0.0:21000
 ip precedence 5
!


Router-3 (Spoke Router)

hostname FR-3
!
ip multicast-routing
!
voice class permanent 1
 signal timing oos timeout disabled
 signal keepalive disabled
 signal sequence oos no-action
!
interface Vif1
ip address 1.1.3.1 255.255.255.0
ip pim sparse-dense-mode
!
router rip
network 1.1.3.0
network 5.5.6.0
!
interface Serial0/0
no frame-relay broadcast-queue
encapsulation frame-relay
frame-relay traffic-shaping
frame-relay broadcast-queue 64 8000 250
!
interface Serial0/0.1 point-to-point
ip address 5.5.6.2 255.255.255.0
ip pim sparse-dense-mode
frame-relay class ipics
frame-relay interface-dlci 100
frame-relay ip rtp header-compression
!
map-class frame-relay ipics
frame-relay cir 128000
```

```
frame-relay bc 1280
frame-relay mincir 128000
no frame-relay adaptive-shaping
frame-relay fair-queue
frame-relay fragment 160
frame-relay ip rtp priority 16384 16384 128
!
voice-port 1/0/0
!connection trunk 111
!operation 4-wire
!
dial-peer voice 1 voip
!destination-pattern 111
!voice class permanent 1
!session protocol multicast
!session target ipv4:239.111.0.0:21000
!ip precedence 5
!
end
```

Configuration with Bidirectional PIM Multicast

Bidirectional PIM multicast is preferred over unidirectional multicast when two PVCs, one dedicated to channel traffic and the other to data traffic, are used. It helps to reduce the number of ip mroute entries that are needed in the router to route multicast traffic. Bidirectional PIM requires one router in the network to act as the rendezvous point (RP).

In the following configuration example, the RP is the loopback interface of Router-1. (The RP can be any interface on any router in the network, as long as it is reachable.)

```
Router-1 (RP node)

hostname bidir-rp
!
ip multicast-routing
!
voice class permanent 1
 signal timing oos timeout disabled
 signal keepalive disabled
 signal sequence oos no-action
!
voice class permanent 2
 signal timing oos timeout disabled
 signal keepalive disabled
 signal sequence oos no-action]
!
interface Loopback1
ip address 10.10.2.1 255.255.255.0
ip pim sparse-mode
!
interface Vif1
ip address 10.1.2.1 255.255.255.0
ip pim sparse-mode
load-interval 30
!
router rip
network 10.1.2.0
network 10.100.0.0
network 10.101.0.0
!
interface Serial0/0
no ip address
```

```
encapsulation frame-relay IETF
load-interval 30
no fair-queue
frame-relay traffic-shaping
frame-relay lmi-type cisco
!
interface Serial0/0.1 point-to-point
description channel pvc
bandwidth 256
ip address 10.100.100.1 255.255.255.0
ip pim sparse-mode
frame-relay interface-dlci 100
class channel
!
interface Serial0/0.2 point-to-point
description data pvc
ip address 10.101.101.1 255.255.255.0
frame-relay interface-dlci 200
class data
!
ip classless
ip pim bidir-enable
ip pim rp-address 10.10.2.1 10 override bidir
!
map-class frame-relay channel
frame-relay cir 128000
frame-relay bc 1000
frame-relay be 0
no frame-relay adaptive-shaping
!
map-class frame-relay data
frame-relay cir 768000
frame-relay mincir 128000
frame-relay adaptive-shaping becn
!
voice-port 1/0/0
 voice class permanent 1
 timeouts wait-release 3
 timing dialout-delay 70
 connection trunk 111
 operation 4-wire
 signal lmr
!
dial-peer voice 1 voip
destination-pattern 111
session protocol multicast
session target ipv4:239.111.0.0:21000
ip precedence 5
!
end

Router-2 (non-RP node)

hostname bidir-2
!
ip multicast-routing
!
voice class permanent 1
 signal timing oos timeout disabled
 signal keepalive disabled
 signal sequence oos no-action
!
voice class permanent 2
 signal timing oos timeout disabled
```

```
 signal keepalive disabled
 signal sequence oos no-action
!
interface Loopback1
ip address 10.10.3.1 255.255.255.0
ip pim sparse-mode
!
interface Vif1
ip address 10.1.3.1 255.255.255.0
ip pim sparse-mode
load-interval 30
!
router rip
network 10.1.3.0
network 10.100.0.0
network 10.101.0.0
!
interface Serial0/0
no ip address
encapsulation frame-relay IETF
load-interval 30
no fair-queue
frame-relay traffic-shaping
frame-relay lmi-type cisco
!
interface Serial0/0.1 point-to-point
description channel pvc
bandwidth 256
ip address 10.100.100.2 255.255.255.0
ip pim sparse-mode
frame-relay interface-dlci 100
class channel
!
interface Serial0/0.2 point-to-point
description data pvc
ip address 10.101.101.2 255.255.255.0
frame-relay interface-dlci 200
class data
!
ip classless
ip route 10.10.2.1 255.255.255.255 Serial0/0.1
ip pim bidir-enable
ip pim rp-address 10.10.2.1 10 bidir
!
map-class frame-relay channel
frame-relay cir 128000
frame-relay bc 1000
frame-relay be 0
no frame-relay adaptive-shaping
!
map-class frame-relay data
frame-relay cir 768000
frame-relay mincir 128000
frame-relay adaptive-shaping becn
!
voice-port 1/0/0
 voice class permanent 1
 playout-delay nominal 100
 playout-delay minimum high
 playout-delay mode adaptive
 playout-delay maximum 250
 timeouts wait-release 3
 timing dialout-delay 70
 connection trunk 111
```

```
 operation 4-wire
 signal lmr
!
dial-peer voice 1 voip
 destination-pattern 111
 session protocol multicast
 session target ipv4:239.111.0.0:21000
 ip precedence 5
```

QoS with Point-to-Point Connections

This section provides information for WANs that have point-to-point connections that include any of these encapsulations:

- Point-to-Point Protocol (PPP)
- Multilink Point-to-Point Protocol (MLPPP)
- High-Level Data Link Control (HDLC)

Guaranteed bandwidth is not an issue on point-to-point (or leased) lines, but you do need to consider connection speed and queuing in these situations. As described in the "QoS with Frame Relay" section on page 3-11, links below 768 K require that larger data packets be fragmented to avoid serialization. In addition, you should use a queuing technique that provides strict priority to Cisco IPICS packets, such as IP RTP Priority, or Low-Latency Queuing.

The FRF.12 fragmentation and reassembly technique that is discussed in the "QoS with Frame Relay" section on page 3-11 does not apply to point-to-point links. For point-to-point links below 768 K, use Multilink PPP (MLPPP) for encapsulation. MLPPP provides feature called Link Fragmentation and Interleaving (LFI). LFI is similar in operation to FRF.12 in that it handles fragmentation at Layer 2.

LFI is not required for networks with link speeds above 768 K because 1,500 bytes packet do not cause more than approximately 10 ms of transport delay. This delay should be acceptable for most delay budgets, so for these networks, HDLC or PPP encapsulation are acceptable.

The following example shows configuring MLPPP with LFI:

```
interface Serial0
bandwidth 64
no ip address
no ip directed-broadcast
encapsulation ppp
no ip route-cache
no ip mroute-cache
no fair-queue
ppp multilink
multilink-group 1
!
interface Multilink 1
ip address 10.1.1.1 255.255.255.252
no ip directed-broadcast
no ip route-cache
ip rtp header-compression iphc-format
ip tcp header-compression iphc-format
no ip mroute-cache
fair-queue 64 256 1000
ppp multilink
ppp multilink fragment-delay 10
ppp multilink interleave
multilink-group 1
ip rtp priority 16384 16383 30
```

QoS for a LAN

When you deploy QoS in a LAN, classify and mark applications as close to their sources as possible. For example, implement QoS in a Cisco Catalyst switch for PMCs or Cisco Unified IP Phones that connect to the Cisco IPICS server via multicast. For LMRs, implement QoS in the dial peer that is configured for the E&M port that connects to the radios

To classify and mark applications, follow these recommendations:

- Use Differentiated Services Code Point (DSCP) markings whenever possible.
- Follow standards based DSCP per-hop behaviors (Fibs) to ensure interoperation and provide for future expansion. These standards include:
  - RFC 2474 Class Selector Codepoints
  - RFC 2597 Assured Forwarding Classes
  - RFC 3246 Expedited Forwarding

> **Note** Because SIP-based PMCs are in the data VLAN instead of in the voice VLAN, Cisco IPICS automatically pushes QoS configuration for PMCs to the RMS.

QoS at the WAN Edge

QoS should be configured at the WAN edge so that QoS settings are forwarded to the next-hop router. When you configure QoS at the WAN edge, follow these recommendations:

- If the combined WAN circuit-rate is significantly below 100 Mbps, enable egress shaping on the Cisco Catalyst switches (when supported)
- If the combined WAN circuit-rate is significantly below 100 Mbps and the Catalyst switch does not support shaping, enable egress policing (when supported)

Policing

Policing is configured so that traffic of a certain class that exceeds the allocated bandwidth is marked as discard eligible (DE) or is dropped, so it prevents denial of service (DoS) or a virus attacks. When you configure policing, follow theses recommendations.

- Police traffic flows as close to their sources as possible.
- Perform markdown according to standards-based rules, whenever supported.
- RFC 2597 specifies how Assured Forwarding traffic classes should be marked down (AF11 > AF12 > AF13). You should follow this specification when DSCP-Based WRED is supported on egress queues.
- Cisco Catalyst platforms do not support DSCP-Based WRED. Scavenger-class remarking is a viable alternative.
- Non-AF classes do not have a markdown scheme defined in standards, so Scavenger-class remarking is a viable option.
- Profile applications to determine what constitutes "normal" or "abnormal" flows (within a 95% confidence interval).
- Deploy campus access-edge policers to remark abnormal traffic to Scavenger.

- Deploy a second-line of defense at the distribution-layer via per-user microflow policing.

- Provision end-to-end "less-than-best-Effort" scavenger-class queuing policies.

Queuing

Queuing is a method of buffering traffic so that the traffic does not overflow the allocated bandwidth on a WAN. To provide service guarantees, enable queuing at any node that has the potential for congestion.

When you enable queuing, follow these recommendations:

- Reserve at least 25% of the bandwidth of a link for the default best effort class

- Limit the amount of strict-priority queuing to 33% of the capacity of a link

- Whenever a Scavenger queuing class is enabled, assigned to it a minimal amount of bandwidth

- To ensure consistent per-hop behavior (PHB), configure consistent queuing policies in the campus, WAN, and VPN, according to platform capabilities

- Enable WRED on all TCP flows, if supported. DSCP-based WRED is recommended

Trust Boundaries

The Cisco IPICS QoS infrastructure is defined by using a trust boundary. For detailed information about trust boundary concepts, refer to *Cisco Unified Communications SRND Based on Cisco Unified CallManager 4.x*, which is available at this URL:

http://www.cisco.com/en/US/products/sw/voicesw/ps556/
products_implementation_design_guide_book09186a00806e8a79.htmlf

A trust boundary can include PMCs (local and remote), LMRs, and Cisco Unified IP Phones. IP precedence should be marked for PMCs and Cisco Unified IP Phones, with a suggested value of 5 for voice traffic (such as RTP) and 3 for voice signaling (such as SIP or SCCP).

For a LMR PTT client, an LMR gateway marks the traffic coming from E&M ports to IP precedence 5 as follows:

```
voice-port 1/0/0
 voice class permanent 1
 connection trunk 111
 operation 4-wire
!
dial-peer voice 111 voip
 destination-pattern 111
 session protocol multicast
 session target ipv4:239.111.0.111:21000
 ip precedence 5
!
```

For a PMC using the remote location, an RMS marks the IP precedence value of 5. The Cisco IPICS server provides the QoS and other necessary configurations to the RMS server for the PMC.

Cisco IPICS traffic that comes from a PMC, LMR, or Cisco Unified IP Phone aggregates on an access switch, and QoS configuration is applied on this switch. Once marked, these values for IP precedence are honored through out the network.

VoIP bearer traffic is placed in a strict priority queue, when possible. The boundary nodes police at the ingress level to rate-limit the VoIP traffic to avoid potential bandwidth exhaustion and the possibility of DoS attack through priority queues.

Figure 3-2 shows a trust boundary.

Figure 3-2   ***Trust Boundary***
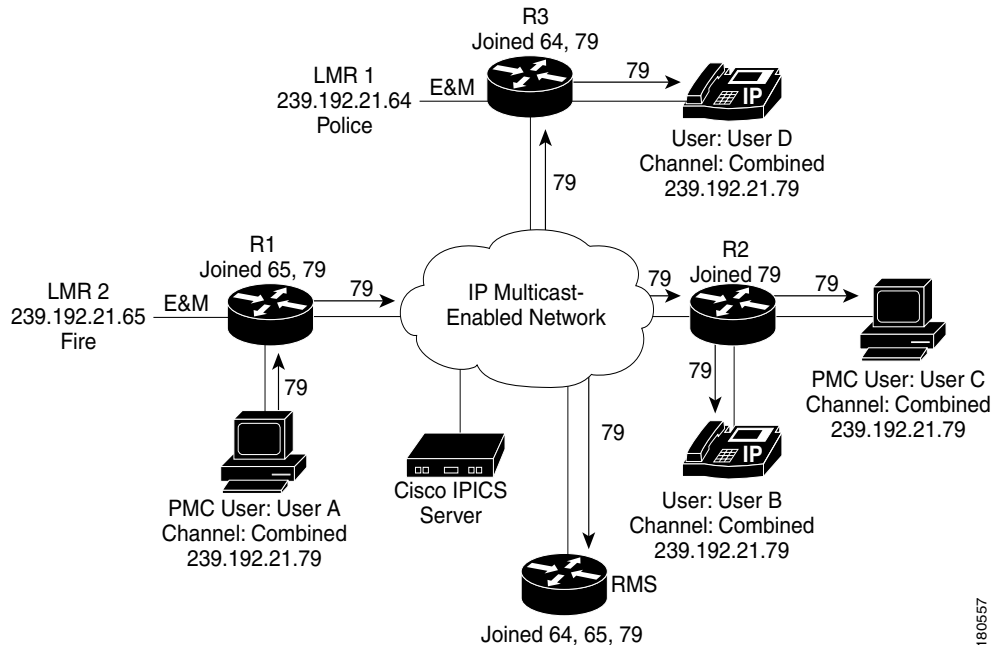
The following example shows access layer QoS configuration for a Cisco Catalyst 3550:

```
CAT3550(config)#mls qos map policed-dscp 0 24 46 to 8
! Excess traffic marked 0 or CS3 or EF will be remarked to CS1
CAT3550(config)#
CAT3550(config)#class-map match-all IPICS-VOICE
CAT3550(config-cmap)# match access-group name IPICS-VOICE
CAT3550(config)#policy-map IPICS-PTTC
CAT3550(config-pmap)#class IPICS-VOICE
CAT3550(config-pmap-c)# set ip dscp 46
! VoIP is marked to DSCP EF
CAT3550(config-pmap-c)# police 128000 8000 exceed-action policed-dscp-transmit
! Out-of-profile IPICS VoIP (G711) is marked down to Scavenger (CS1)
CAT3550(config-pmap-c)#class IPICS-SIGNALING
CAT3550(config-pmap-c)# set ip dscp 24
! Signalling is marked to DSCP CS3
CAT3550(config-pmap-c)# police 32000 8000 exceed-action policed-dscp-transmit
! Out-of-profile Signalling is marked down to Scavenger (CS1)
CAT3550(config-pmap-c)#class class-default
CAT3550(config-pmap-c)# set ip dscp 0
CAT3550(config-pmap-c)# police 5000000 8000 exceed-action policed-dscp-transmit
! Out-of-profile data traffic is marked down to Scavenger (CS1) 50000 (Depends on per
customer design or IPICS Engineering team might have some recommendations)
CAT3550(config-pmap-c)# exit
CAT3550(config-pmap)#exit
CAT3550(config)#
CAT3550(config)#interface range FastEthernet0/1 - 48
CAT3550(config-if)# service-policy input IPICS-PTTC
! Attaching the policy map IPICS-PTTC to the interface range
CAT3550(config-if)#exit
CAT3550(config)#
CAT3550(config)#ip access-list extended IPICS-VOICE
```

```
                   Modified this list since the multicast addresses recommend above are incorrect. Just need
                   the last entry, unless also using GLOP then keep 233
                   ! Extended ACL for the IPICS Address/Port ranges
                   CAT3550(config-ext-nacl)#
                    permit udp 233.0.0.0 0.255.255.255 233.0.0.0 0.255.255.255 range 21000 65534
                    permit udp 233.0.0.0 0.255.255.255 239.0.0.0 0.255.255.255 range 21000 65534
                    permit udp 239.0.0.0 0.255.255.255 233.0.0.0 0.255.255.255 range 21000 65534
                    permit udp 239.0.0.0 0.255.255.255 239.0.0.0 0.255.255.255 range 21000 65534
                   CAT3550(config-ext-nacl)#ip access-list extended IPICS-SIGNALING
                   ! Extended ACL for the remote PMC clients SIP signaling
                   CAT3550(config-ext-nacl)# permit udp <RMS IP Address> <Any > eq 5060
                   ! <Any> is the address on Remote PMC SIP Client
                   CAT3550(config-ext-nacl)#end
                   CAT3550#
```

Port Utilization

This section describes the ports that can be used in a Cisco IPICS deployment. You can use this information to determine how best to define the QOS or firewall settings at a port level, if required. Information about how to facilitate modifications to the port ranges are included.

Table 3-4 describes the default ports that are used by Cisco IPICS components.

Table 3-4   Default Ports used by Cisco IPICS Components

Protocol	Device	Destination Port	Remote Device
HTTP	PMC	TCP 80	Cisco IPICS server
	Cisco IPICS Administration Console	TCP 80	Cisco IPICS server
	Cisco Unified IP Phone	TCP 80	Cisco Unified CallManager, Cisco Unified CallManager Express
HTTPS	PMC	TCP 443	Cisco IPICS server
	Cisco IPICS Administration Console	TCP 443	Cisco IPICS server
SIP	PMC / Policy Engine	UDP 5060	RMS / Policy Engine SIP provider **Note** Used for remote PMC to RMS and for Policy Engine to SIP provider

Table 3-4   Default Ports used by Cisco IPICS Components (continued)

Protocol	Device	Destination Port	Remote Device
RTP/RTCP	PMC	UDP 16384-32766	RMS **Note** The session target for voice multicasting dial peers is a multicast address in the range 224.0.1.0 through 239.255.255.255. This session target must be the same for all routers in a session. The audio RTP port is an even number in the range 16384 through 32767, and must also be the same for all routers in a session. An odd-numbered port (UDP port number + 1) is used for the RTCP traffic for that session.
	Policy Engine	UDP 32768-61000	Cisco Unified CallManager, Cisco Unified CallManager Express
ICMP (PING)	PMC	ICMP	Cisco IPICS server
IGMP	PMC	ICMP	Multicast group
SSH	Cisco IPICS server	TCP 22	RMS

Guidelines for Using IP Multicast Addresses with Cisco IPICS

When you use multicast communications with Cisco IPICS be aware of the following guidelines:

- Cisco strongly recommends that you configure IP multicast addresses that are only in the 239.192.0.0 to 239.251.255.255 range.

    - This address range is part of the Administratively Scoped Block, as specified by RFC 3171, and is intended for use in a local domain. As such, this address range is less likely to cause an addressing conflict in an existing multicast domain.

    - For more information, refer to RFC 3171, *Internet Assigned Numbers Authority (IANA) Guidelines for IPv4 Multicast Address Assignment* and to RFC 2365, *Administratively Scoped IP Multicast*.

- Cisco IPICS permits the use of IP multicast addresses that span the range 224.0.0.0 through 239.255.255.255, where the first octet contains 224, 232, 233, 238 or 239 and subsequent octets contain 0 through 255.

    - Although these well-known group address ranges can be used with Cisco IPICS, they should not be used because they can cause undesirable results if they are used improperly.

    - Refer to RFC 3171 for more information and use guidelines for these IP multicast addresses.

For additional information about the use of IP multicast addressing, refer to the following URL:

http://www.cisco.com/en/US/tech/tk828/tsd_technology_support_protocol_home.html

Multicast and Unicast

When the IPICS solution is configured, channels are assigned static multicast IP addresses and port pairs. Ports that are assigned to a PMC must be within the acceptable port range for a PMC. In addition, the multicast addresses that are defined in the multicast address pool for dynamic allocation must be within the acceptable port range for a PMC.

When a PMC connects using the remote location, it establishes a unicast media connection to an RMS. The UDP port that is assigned for the media connection will be allocated from within the supported range.

QOS Policy Considerations

When defining QOS policies that will be assigned to a UDP port range, using a Source Host and Destination Host addresses of ANY allows the QOS policy to be properly set based on the PMC UDP port range. In this case, UDP ports that are assigned by the RMS are not considered, which helps to simplify the QOS policies.

Securing the Cisco IPICS Infrastructure

The following sections provide information about providing system security for Cisco IPICS:

Secure Socket Layer

Cisco IPICS uses Secure Socket Layer (SSL) to encrypt communications between a PMC and the Cisco IPICS server. The browser with which you access the Cisco IPICS Administration Console uses HTTPS. To enforce SSL, you must install a certificate on the Cisco IPICS server. You can use a self-signed certificate or, to impose additional security, you can purchase and set up a digitally-signed certificate. In addition, the RMS control uses SSH as a client.

For additional information, refer to the "Installing Third Party Certificates on the Cisco IPICS server section in *Release Notes for Cisco IPICS, Release 1.0(2)*.

Cisco Security Agent

The Cisco Security Agent (CSA), which is optionally available separately from Cisco IPICS, provides intrusion detection and prevention for the Cisco IPICS server and the for the PMC. Instead of focusing on attacks, CSA focuses on preventing malicious and undesired activities on the host. CSA detects and blocks the damaging activities. CSA ships with pre-defined policies that prevent most types of malicious activity from occurring. Since malicious activity is always undesired, security at this level is very inexpensive to deploy. Little or no environment tuning is required. In rare cases, critical business applications can be affected if CSA prevents your web server from adding a new user account

Firewalls and Access Control Lists

Use a firewall and access control lists (ACLs) in front of the IPICS server and other Cisco IPICS components to add an extra layer of security. For example, you can use a firewall or an ACL to allow only call control and management packets to reach the Cisco IPICS server, and to block unnecessary traffic such as Telnet or TFTP traffic. You can use ACLs to allow only the source addresses that are supposed to access your network.

When you use a firewall, it must support state-full inspection of voice signaling protocol. Cisco IPICS uses UDP ports 21000 through 65534, and a firewall must only open the ports that are needed to support this application. In addition, make sure that the firewall supports application layer gateway (ALG) capabilities. ALG inspects signaling packets to discover what UDP port an RTP stream is going to use and dynamically opens a pinhole for that UDP port.

Other Security Recommendations

For additional security in a Cisco IPICS network, follow these recommendations:

- Use TACACS+ and RADIUS to provide highly secure access in your network
- Do not rely only on VLANs for separation; also provide layer 3 filtering at the access layer of your network
- Use VLANs and IP filters between your voice and data network
- Use out of band management switches and routers with SSH, HTTPS, OOB, permit lists, and so on to control who is accessing your network devices
- Disable unused switch ports on the LAN switches and place them in an unused VLAN so that they are not misused
- Use spanning tree (STP) attack mitigation tools such as BPDU Guard and Root Guard
- Disperse critical resources to provide redundancy
- Provide limited and controlled access to power switches
- Use IDS Host software on the Cisco IPICS server and other network servers to ensure security of voice applications

Cisco IPICS Network Management System

When you plan for managing and monitoring a Cisco IPICS network, define the parameters that can be monitored in the Cisco IPICS environment. You can use the outputs from these parameters to establish a set of alarms for spontaneous problems and to establish a proactive early warning system.

As you develop a management and monitoring policy for your network, take these actions:

- For each component in the network, define the parameters that must monitored on the component
- Select the network management and monitoring tools that are appropriate for monitoring the parameters that you defined

Managing the Overall Network

The Cisco Multicast Manager (CMM) is a web-based network management application that is designed to aid in the monitoring and troubleshooting of multicast networks. Cisco Multicast Manager includes the following features and benefits:

- Early warning of problems in multicast networks

- In-depth troubleshooting and analysis capabilities

- On demand, real time and historical reporting capabilities

- Optimization of network utilization and enhancement of services delivery over multicast enabled networks

CMM can monitor all multicast-capable devices that are running Cisco IOS Software, including Layer 2 switches. For more detailed information about CMM, refer to this URL:

http://www.cisco.com/en/US/products/ps6337/index.html

If you are using Cisco Unified IP Phones as PTT clients in your Cisco IPICS network, you can use various IP Telephony (IPT) management tools to manage these devices. For example, you can use Enterprise IPT management solution, which uses OpenView Gateway Statistics Utility (GSU) Reporting Solution and CiscoWorks IP Telephony Environment Monitor (ITEM) solution to provide real-time, detailed fault analysis specifically designed for Cisco IPT devices. This tool evaluates the health of IPT implementations and provides alerting and notification of problems and areas that should be addressed to help minimize IPT service interruption. IPT management solution also identifies underutilized or imbalanced gateway resources, and provides historical trending and forecasting of capacity requirements

Other items to monitor in a Cisco IPICS network in include the following:

- Cisco IPICS server health

- Cisco IPICS services health

- IP gateway health

- Cisco Unified CallManager functionality

- QoS monitoring

- L2/L3 switches and applications