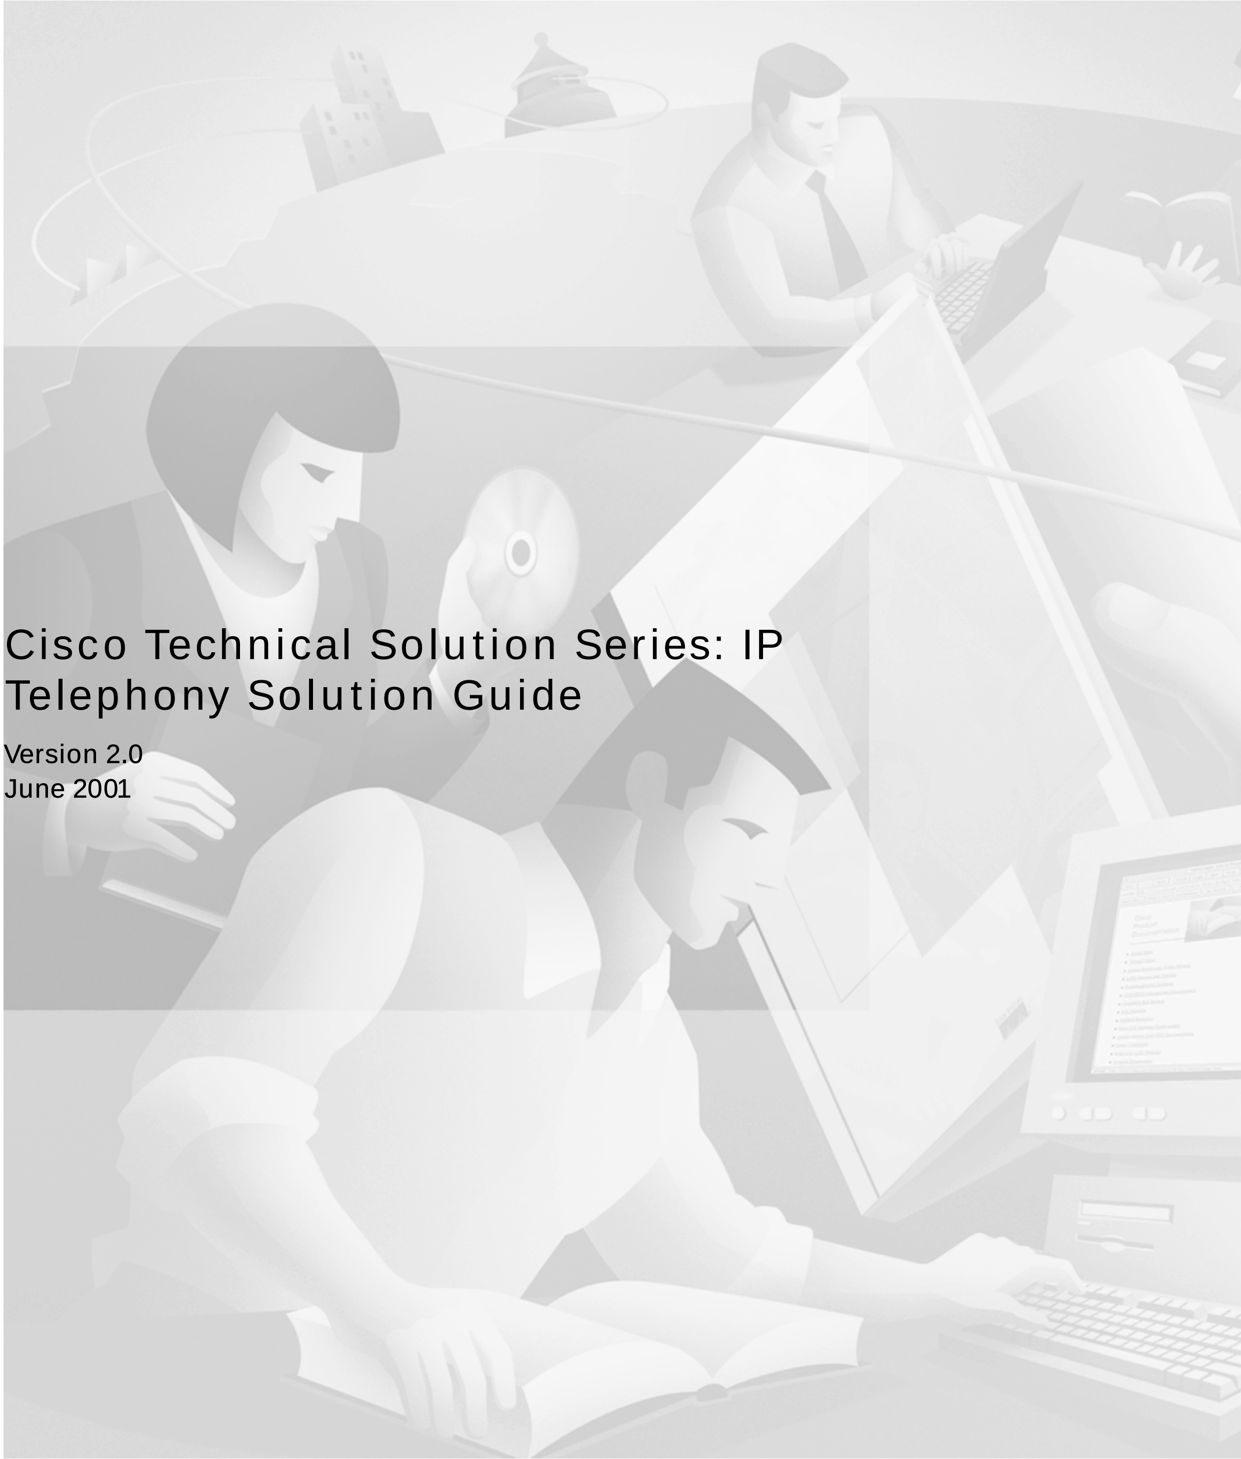

Cisco Technical Solution Series: IP Telephony Solution Guide

Version 2.0
June 2001

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

Cisco Technical Solution Series: IP Telephony Solution Guide

Introduction to IP Telephony

Overview

The Cisco IP Telephony Solution Guide is intended to help organizations implement and manage IP Telephony network solutions, which includes Planning, Design, Implementation, and Operations network phases. This method is called the PDIO model. Cisco experts in IP Telephony design, network design, customer support, high availability, network management, network implementation, and traditional telecom systems collaborated to create this document so that you can reduce guesswork, technical resources, and the time needed to ensure successful implementation of a Cisco IP Telephony network.

Organization

This solution guide consists of the following sections:

- Introduction to IP Telephony - provides a brief introduction to this manual.
- Chapter 2, "IP Telephony Architecture Overview" provides a general description of the IP Telephony architecture.
- Chapter 3, "Planning the IP Telephony Network" provides information necessary for planning IP Telephony solutions.
- Chapter 4, "Designing the IP Telephony Network" provides detailed design specifications for building IP Telephony networks.
- Chapter 5, "Implementing the IP Telephony Network" provides important information for successfully implementing IP Telephony.
- Chapter 6, "Operating the IP Telephony Network" provides information for successfully operating, networking, securing, and troubleshooting IP Telephony networks.

Audience

The Cisco IP Telephony Solution Guide is intended for the following audiences:

- Cisco customers involved with the planning, technical design, implementation, and operation of IP Telephony solutions
- Technical management or network planning personnel

- Cisco Sales Engineers, Technical Support Engineers, Cisco Professional Services, and Cisco Support Partners

This document also assumes some technical knowledge of Cisco switching, routing, Quality of Service, CallManager functionality, gateway functionality, and voice signaling principles.

Scope

The Cisco IP Telephony Solution Guide discusses the core components of the IP Telephony network:

- Current data network design for IP Telephony
- CallManager version 3.0
- Gateways supported under the current IP Telephony architecture
- Voice mail systems

The following applications are not discussed:

- uONE unified messaging
- TAPI or JTAPI

Contact your Cisco representative or visit the following Cisco website for available information on IP Telephony solution applications not covered in this solution guide:
www.cisco.com.

Revision History

Table 1-1 Cisco IP Telephony Solution Guide Revision History

Version	Date
Cisco Technical Solutions Series: IP Telephony Solution Guide Version 2.0	June 2001
Cisco Technical Solutions Series: IP Telephony Solution Guide Version 1.0	February 2001

Related Information

- IP Telephony Design Guide

  http://www.cisco.com/univercd/cc/td/doc/product/voice/ip_tele/index.htm
- IP Telephony Support Pages and Documentation

  http://www.cisco.com/warp/public/788/AVVID/avvid_index.shtml

CHAPTER 2

IP Telephony Architecture Overview

A previously published document, *The Architecture for Voice, Video, and Integrated Data*, can be found at the following Cisco.com location: http://www.cisco.com/warp/public/cc/so/neso/vvda/iptl/avvid_wp.htm.

Planning the IP Telephony Network

In this Chapter

This chapter consists of the following sections:

- Evaluating and Documenting the Existing Data Infrastructure
- Evaluating and Documenting the Existing Telecom Infrastructure
- Evaluating and Documenting the Existing Power/Cabling Infrastructure
- IP Telephony Availability Requirements
- Planning for WAN Deployments
- Operations and Implementation Planning

Related Information

- Data Sheet: Cisco VoIP Readiness Net Audit Planning for Migration to IP Telephony
  http://www.cisco.com/warp/public/cc/serv/mkt/sup/ent/avvid/nadit_ds.htm
- Cisco IP Telephony Network Design Guide
  http://www.cisco.com/univercd/cc/td/doc/product/voice/ip_tele/index.htm
- Westbay Engineers Limited Home Page
  http://www.erlang.com
- APC Home Page
  http://www.apcc.com
- Cisco IP Telephony Power Protection Page
  http://www.cisco.com/warp/public/779/largeent/avvid/solutions/powerpro.html

Evaluating and Documenting the Existing Data Infrastructure

Organizations need to evaluate their existing data infrastructure to help determine upgrade requirements for the IP Telephony solution. You may need to provide infrastructure for additional bandwidth, consistent performance, or higher availability required for the converging environment. This section describes both LAN and WAN requirements.

You should document and evaluate the existing data infrastructure in terms of:

- New voice performance requirements
- Availability requirements
- Feature requirements
- Potential network capacity or impact.

The required information for this includes network maps, device inventory information, and network baseline information. Analyzing these areas will help you understand the data network upgrade requirements needed to support IP Telephony and basic network availability, performance, and feature requirements.

To evaluate voice performance requirements, review the device inventory, network design, and baseline information. Links and devices should have sufficient capacity for the additional voice traffic. You may need to upgrade links with high peak or busy hour utilization. Target devices with high CPU utilization, high backplane utilization, high memory utilization, queuing drops, or buffer misses for additional inspection and potential upgrade. Peak utilization characteristics in the baseline will be valuable in determining potential voice quality issues.

To evaluate availability requirements for the IP Telephony network, review the network topology, feature capabilities, and protocol implementations. Review redundancy capabilities of the network to ensure you can meet availability goals with the current network design (or a new design) recommended for IP Telephony.

To evaluate current feature capabilities of the network, evaluate device characteristics including a chassis, module, and software inventory. This will be useful in determining IP Telephony feature capabilities in the existing environment.

You should also evaluate overall network capacity and impact to ensure that the network will meet overall capacity requirements and that there will be no impact on the existing network and application requirements. You should evaluate the network baseline in terms of the impact from IP Telephony requirements. You may need to add more CPU, memory, bandwidth, or features to ensure you meet both IP Telephony and existing network requirements.

Note Cisco can provide an IP Telephony readiness audit that provides the recommended baseline information.

LAN/Campus Environment

We recommend a LAN/Campus analysis for all LAN environments involving any of the four IP Telephony deployment models that include:

- Single site
- Networked with PSTN
- Multi-site with centralized call processing
- Multi-site with distributed call processing.

The LAN/Campus infrastructure analysis determines infrastructure and bandwidth issues that will affect IP Telephony voice quality and availability. You should collect the following types of information for the LAN/campus infrastructure analysis:

- LAN/campus topology
- IP addressing plan

- Location of TFTP servers, DNS servers, DHCP servers, firewalls, NAT (Network Address Translation) gateways, and PAT (Port Address Translation) gateways
- Potential location of gateways and CallManager clusters
- Protocol implementation including IP routing, Spanning Tree, VTP, IPX, and IBM protocols
- Device analysis including software versions, modules, ports, speeds, and interfaces
- Phone connection methodology (direct or daisy chain)
- Baseline showing network and resource control plane use

LAN/Campus Topology

You normally build LAN/campus infrastructures using a hierarchical access, distribution, and core model. One or two of these layers may be collapsed into smaller LAN/Campus environments. However, in general, you will have a standard deployment model with a standard distribution and core configuration. Read the *Campus Network Design* document to review Cisco's recommendations for a high availability campus design. This document can be found at the following location: http://cco/warp/public/779/largeent/design/campus_index.html.

You should create a simple map, such as Figure 3-1, that describes the layers, devices, media, and port speeds. The topology map should also show the location of TFTP servers, DNS servers, DHCP servers, firewalls, and gateways.

Review the following LAN/campus topology issues:

- Available average bandwidth
- Available peak or burst bandwidth
- Resource issues that may affect performance including buffers, memory, CPU, and queue depth
- Network availability
- IP phone port availability
- Desktop/phone QoS between user and switch
- CallManager availability
- Network scalability with increased traffic, IP subnets, and features
- Backup power capability
- LAN QoS functionality
- Convergence at Layers 2 and 3

Figure 3-1 LAN/Campus Topology

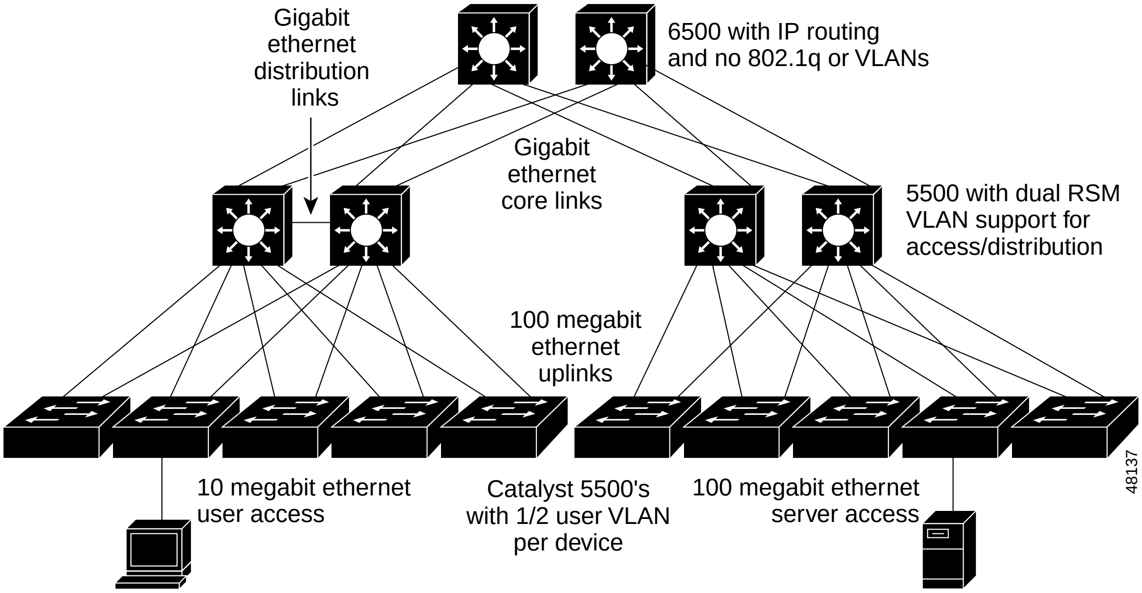
Gigabit ethernet distribution links
6500 with IP routing and no 802.1q or VLANs
Gigabit ethernet core links
5500 with dual RSM VLAN support for access/distribution
100 megabit ethernet uplinks
10 megabit ethernet user access
Catalyst 5500's with 1/2 user VLAN per device
100 megabit ethernet server access
48137

IP Addressing Plan

Review the following IP addressing plan and implementation characteristics:

- Phone IP addressing plan
- Average user IP subnet size use for the campus
- Number of core routes
- IP route summary plan
- DHCP server plan (fixed and variable addressing)
- DNS naming conventions

Potential considerations with IP addressing include:

- Route scalability with IP phones
- IP subnet space allocation for phones
- DHCP functionality with secondary addressing
- IP subnet overlap
- Duplicate IP addressing

Location of Servers and Gateways

Consider the location (or potential location) of servers and gateways prior to implementation and identify them in the LAN infrastructure planning phase as much as possible. Investigate other issues later to help ensure that service availability is consistent across the LAN infrastructure and for multiple sites. You should identify gateway and server network locations for the following:

- TFTP servers
- DNS servers
- DHCP servers

- Firewalls
- NAT or PAT gateways
- CallManager location
- Gateway location

Investigate these issues after you determine the location:

- Network service availability
- Gateway support (in conjunction with IP Telephony solution)
- Available bandwidth and scalability
- Service diversity

Protocol Implementation

Investigate overall protocol uses to determine IP Telephony scalability and any potential IP Telephony availability issues or additional protocol service issues. Review the following areas for the protocol implementation analysis:

- IP routing including protocols, summarization methods, NBMA (non-broadcast media access) configurations, and routing protocol safeguards
- Spanning Tree configuration including domain sizes, root designation, uplink fast, backbone fast, and priorities in relation to default gateways
- HSRP configuration
- VTP and VLAN configuration
- IPX, DLSW, or other required protocol services including configuration and resource usage

You should review the following issues in relation to protocol implementation:

- Protocol scalability
- Network availability
- Potential impact on IP Telephony performance or availability

Device Analysis

Analyze the existing network devices to help identify hardware and software issues associated with the IP Telephony deployment. Many devices may not have the desired control plane resources, interface bandwidth, QoS functionality, or power management capabilities. The following table displays device attributes that may be important:

- Device (type and product ID)
- Software version(s)
- Quantity deployed
- Modules and redundancy
- Services configured
- User media and bandwidth
- Uplink media and bandwidth
- Switched vs. shared media

- Users per uplink and uplink load sharing/redundancy
- Number of VLANS supported
- Subnet size, devices per subnet

Network Baseline

You can use a network baseline of the existing campus/LAN infrastructure for IP Telephony capacity planning. This will help determine potential voice quality issues and the impact to the existing environment. Measure the following characteristics as part of the baseline:

- Device average and peak CPU
- Device average and peak memory
- Peak backplane utilization
- Average link utilization (prefer peak hour average for capacity planning)
- Peak link utilization (prefer five minute average or smaller interval)
- Peak queue depth
- Buffer failures
- Average and peak voice call response times (before IP Telephony implementation)

Many different individuals and support organizations recommend different acceptable threshold values for these measured baseline issues. Remember that IP Telephony requires consistent performance and quality; therefore, all of the areas should be below safe recommended threshold values at all times. Use the following general guidelines on threshold issues:

- **CPU**—A requirement for all background processing in addition to some traffic processing requirements. Background processing includes route updates, keepalives, network management, and other critical processes for keeping the network up and stable. During stressful network times, such as route convergence or link flapping, significant CPU will be used to ensure the network remains stable and intact. Because significant CPU can be used during stress situations, a good rule of thumb is 50% peak CPU and 30% average CPU.
- **Memory**—Like CPU, main memory is used for background processing and traffic processing. And like CPU, significant amounts of memory can be used for a processing during link flap conditions, routing changes, and cache changes. Because significant changes can occur in memory requirements, a good rule of thumb is 50% peak and 30% average.
- **Backplane Utilization**—Can be a major issue in some devices if the port speed and density is higher than the available backplane capabilities. Backplane utilization over 50% may also indicate some port queuing or dropped traffic for trunks that have less bandwidth than the sum of all downstream bandwidth.
- **Link Utilization**—Critical to IP Telephony deployments because of VoIP performance and jitter requirements. First, remember that SNMP thresholds for peak utilization are still mainly done for five minute intervals. A good rule of thumb is to increase bandwidth utilization 40% above the five minute value to determine a true measure of peak utilization over the five minute average. Average link utilization may also be useless over time if peak-critical traffic occurs during a shorter interval of one hour. The telecom community thinks in terms of “busy hour” traffic. If you perform capacity planning using this busy hour utilization, then data network managers can consistently meet both voice and data requirements.

  To some extent, QoS capabilities at level II and level III will help minimize the need for significant bandwidth headroom. However, voice will add significant volume to the network and care should be taken to ensure that data traffic is not *starved*. Network designers also like to ensure that more bandwidth is available towards the core to help minimize or eliminate significant or critical congestion problems. Therefore, care should be taken for all core network links that have peak utilization in excess of 50% and average utilization above 30%. VoIP will likely work if it is higher, but there will be more opportunity for potential intermittent bandwidth problems that will first affect the critical voice traffic.

- **Queue Depth**—Indicates link congestion. Any transmit queues that experience any volume at all indicate that traffic is waiting. This directly impacts voice jitter and delay and indicates that link utilization is exceeding a peak recommended value.
- **Buffer Failures**—Indicates a temporary inability to perform control processing in the device and should be investigated in terms of overall network health. Some buffer failure issues could impact VoIP quality and should be investigated.

Note Cisco can provide a network baseline called the IP Telephony readiness Net Audit (http://www.cisco.com/warp/public/cc/serv/mkt/sup/ent/avvid/nadit_ds.htm).

WAN Environment

We recommend a WAN infrastructure analysis for multiple-site WANs with distributed call processing or multi-site IP WANs with centralized call processing. The WAN analysis determines infrastructure and bandwidth issues that will affect IP Telephony quality and reliability. You should collect the following information for the WAN environment analysis:

- WAN topology
- Location of gateways and servers
- WAN protocols
- Existing QoS requirements
- Device Analysis including software versions, modules, ports, speeds and interfaces
- WAN baseline

Note Review "LAN/Campus Environment" for information on location of gateways, IP addressing plan, and protocol implementation. We recommend some LAN analysis for all WAN sites supporting IP Telephony.

WAN Topology

You normally build WAN topology infrastructures using a hub and spoke model, meshed multi-site model, or a combination of both. You should create a WAN diagram showing potential IP Telephony sites, WAN devices, remote LAN devices, interface types, and bandwidth. The map should show the location of DNS servers, DHCP servers, firewalls, gateways, and potential CallManager locations. See Figure 3-2 for a sample WAN topology.

Review the following WAN topology issues:

- WAN availability, including bandwidth redundancy and resiliency
- WAN design or topology issues that may affect IP Telephony quality or performance

Note The *Cisco IP Telephony Network Design Guide* currently recommends a hub and spoke topology until call admission control using RSVP (Resource Reservation Protocol) is completely available. This document can be found at the following location: http://www.cisco.com/univercd/cc/td/doc/product/voice/ip_tele/index.htm.

- WAN scalability with increased traffic, IP subnets, and features
- Bandwidth and WAN service expectations
- Existing QoS requirements (see the "Designing for LAN/WAN QoS" section on page 4-2 for more information.)

Figure 3-2 WAN Infrastructure

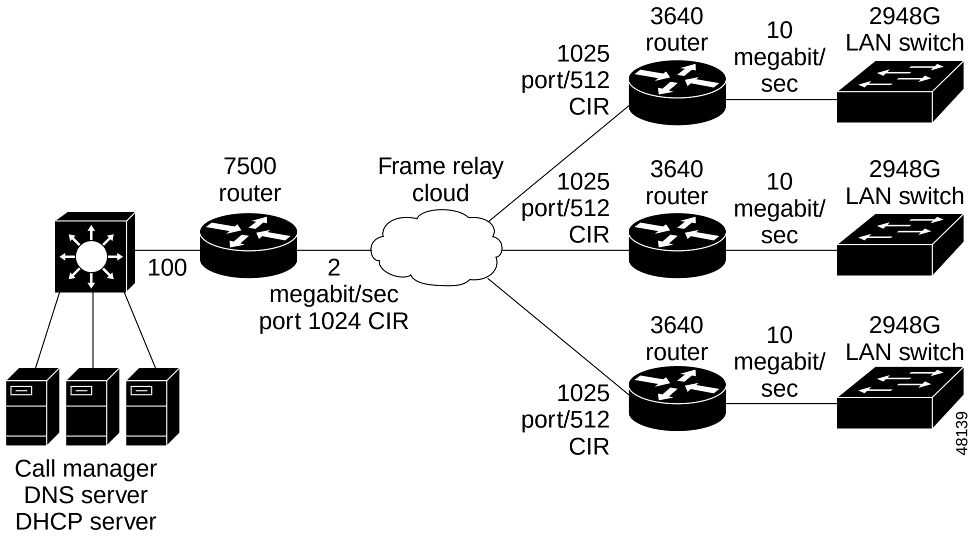

Location of Servers and Gateways

Consider the location (or potential location) of servers and gateways in the WAN prior to implementation and identify them in the WAN infrastructure planning phase as much as possible. Identify the following gateway and server network locations:

- TFTP servers
- DNS servers
- DHCP servers
- Firewalls
- NAT or PAT gateways
- CallManager location
- Gateway location

You should investigate the following issues after you determine the location:

- WAN outage impact and service diversity
- Gateway support (in conjunction with IP Telephony)
- Available bandwidth and scalability

WAN Protocols

You should investigate WAN protocols for issues that may impact IP Telephony quality or issues that may be affected by additional voice services. In many cases, the WAN may require further optimization to better support IP Telephony traffic. NBMA (non-broadcast multiaccess) environments may also be susceptible to protocol issues and overall reliability that can affect voice quality. Investigate the following specific issues:

- WAN IP protocol implementation and protocol overhead
- IP multicast implementation
- Carrier Service subscription rates including port speed, committed information rates, and expected performance
- NBMA protocol issues affecting voice quality and performance
- Other protocol overhead, including IPX and SNA

Analyze the following areas after investigating WAN protocol issues:

- Protocol optimization
- WAN scalability with increased traffic
- Expected network convergence with redundant topologies
- Carrier reliability and quality expectations with WAN protocols

Existing QoS Requirements

You should evaluate existing WAN QoS requirements to determine compatibility with Voice QoS requirements. You should identify applications and performance requirements, including application performance, burst requirements, and batch requirements. Investigate the following areas:

- Existing WAN QoS configurations
- Critical application requirements, including raw performance, burst bandwidth, and batch bandwidth
- WAN call admission control

Device Analysis

An analysis of existing network devices in the network helps identify hardware and software issues associated with the IP Telephony deployment. Software versions are important to determine QoS requirement compatibility. You can also use this information to create a network reliability path analysis to help determine potential network availability. The following table displays device attributes that may be important:

- WAN Devices
- Software Version(s)
- Remote LAN Devices
- Software Version(s)
- Quantity Deployed
- Modules and Redundancy
- Services Configured

- WAN Media/Bandwidth
- LAN Media/Bandwidth
- Switched vs. Shared Media
- User and IP Addressing per WAN Site

WAN Baseline

You can use a WAN baseline of the existing WAN and WAN site infrastructure for IP Telephony capacity planning. This will help determine potential voice quality issues and the impact to the existing environment. Measure the following characteristics as part of the baseline:

- Device average and peak CPU
- Device average and peak memory
- Average link utilization (prefer peak hour average for capacity planning)
- Peak link utilization (prefer five minute average or smaller interval)
- Peak queue depth
- Buffer failures
- Average and peak voice call response times (before IP Telephony implementation)

See the “Network Baseline” section on page 3-6 for specific guidelines for measuring these characteristics.

Evaluating and Documenting the Existing Telecom Infrastructure

You need to evaluate the existing Telecom infrastructure to help determine IP Telephony requirements. Perform this analysis for all sites implementing VoIP technology to determine the appropriate deployment model. IP Telephony supports the following deployment models:

- Single-site deployment
- Multiple single-site deployments interconnected via PSTN
- Distributed IP Telephony sites with centralized call processing
- Distributed IP Telephony sites with distributed call processing

The Telecom infrastructure analysis examines the products, services, and features used in the existing telecom environment including:

- PBX systems and locations
- Voice mail systems and locations
- Key systems
- PBX inter-connectivity
- Phone requirements
- PSTN trunking
- Voice mail trunking
- Site-to-site trunking

The analysis will then help determine the IP Telephony design criteria. You should examine the following issues:

- Existing PBX topology, including voice mail servers
- PBX and Key Systems
- Voice mail system
- Voice trunking
- Phones per site and phone features
- Existing dial plan
- Fax requirements

Examining the Existing Telecom Topology

The existing Telecom topology includes the location and internetworking connectivity for PBX systems, key systems, and voice mail servers. The topology should include the location of these devices and the trunks between systems used for connectivity. Trunking may include site-to-site trunks, PSTN trunks, and voice mail trunks. This section reviews the following existing Telecom topology issues:

- PBX system connectivity overview
- Trunking overview

See Figure 3-3 for an example telecom topology showing PBX systems, key systems, and voice mail systems:

Figure 3-3 Telecom Topology

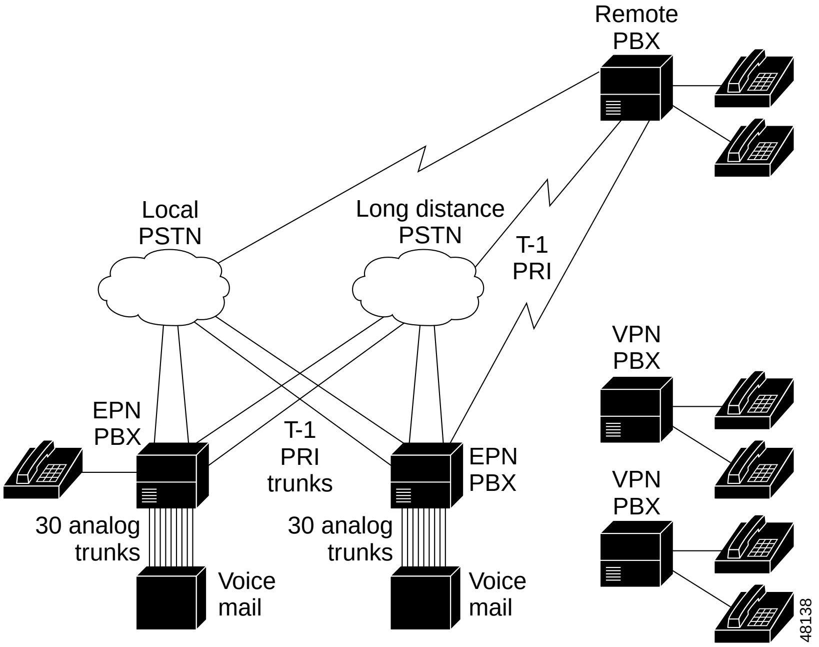

Examining PBX and Key Systems

You need PBX and key system information to help understand current voice features and functionality. The following information will help determine required features and PBX-to-IP Telephony connectivity requirements.

- PBX or KSU vendor and model
- Quantity and locations of PBX/KSU systems
- Release of software running on PBX or KSU
- Quantity and location of PBXs with which IP Telephony may interface
- Hardware models and revisions of installed cards
- Software features currently deployed, which may include call setup, conferencing, call transfer, call hold, call park, calling line identity, and calling party name
- Number of existing analog connections for each PBX or KSU and three expected to remain following deployment
- Number of existing digital connections for each PBX/KSU and those that will remain
- Number and capacity of ISDN trunks connected to each PBX

Examining Voice Mail Systems

You will need the following information to determine IP Telephony compatibility and feature capabilities:

- Voice mail system models and vendor
- Quantity and locations of voice mail systems
- Hardware model and revision cards of voice mail systems
- List of software features currently deployed with voice mail system
- Does the voice mail system have an SMDI interface?
- How is the voice mail system connected to the PBX?
- Is the message waiting indicator integrated into the voice mail solution?

Examining Voice Trunking

Use the existing voice trunking to determine the IP Telephony gateway requirements. In general, you should identify the trunks for voice mail, PSTN connectivity, and site-to-site trunking requirements. In addition, define the existing blocking factor for potential capacity issues. Cisco recommends a blocking factor of one percent for IP Telephony trunking. You may wish to complete a traffic analysis to understand busy hour trunking for the various trunking applications. You can then use an Erlang-B calculator (http://www.erlang.com) to determine new trunking requirements. PBX vendors can normally provide busy hour statistics as a support service. Use Table 3-1 to help identify overall trunking:

Table 3-1 Trunking Matrix

	Digital or Analog	Two-way Calling	DID Trunks	DOD Trunks
Voice Mail Trunks				
Local PSTN Trunks				
LD PSTN Trunks				
Trunks to Site X				
Trunks to Site Y				
Trunks to Site Z				

You may also use the following tables for planning and configuring the gateway trunks. In some cases, you may move these trunk Demarc locations to co-exist with IP Telephony equipment. In addition, you should document support responsibility for WAN carrier services for use in physical design documents.

Table 3-2 PBX WAN Trunk Information

Local Site A Name:								
Item No.	Local Location A	Remote Location B	Type (see below)	Speed (Kbps)	Framing	Coding	Local CSU/DSU A Vendor and Model	Remote CSU/DSU B Vendor and Model
1.								
2.								
3.								
4.								
5.								
6.								
7.								
8.								
9.								
10.								

Table 3-3 PBX WAN Trunk Cable Infrastructure Information

Local Site A Name:								
Item No.	PBX A Slot No./Port No.	PBX A Connector Type - Gender	PBX A/CSU Cable Length (ft.)	CSU A DTE Connector Type - Gender	CSU NET A Connector Type - Gender	CSU to Demarc Cable Length (ft.)	Demarc Connector Type - Gender	Remote PBX B Slot No./Port No.
1.								
2.								
3.								

Table 3-3 PBX WAN Trunk Cable Infrastructure Information

Local Site A Name:								
4.								
5.								
6.								
7.								
8.								
9.								
10.								

Note When ordering your DID, get a block of telephone (DID) numbers equal to or greater than the number of devices (phones, virtual phones, and H.323 devices such as NetMeeting) that will be connected to the network.

Table 3-4 PBX WAN Carrier and Circuit Information

Local Site A Name:						
Item No.	Local Carrier A Company Name	Local Carrier A Circuit ID	Long Haul Carrier Company Name	Long Haul Carrier Circuit ID	Remote Carrier B Company Name	Remote Carrier B Circuit ID
1.						
2.						
3.						
4.						
5.						
6.						
7.						

Table 3-4 PBX WAN Carrier and Circuit Information

Local Site A Name:						
8.						
9.						
10.						

ISDN PRI is a common PBX WAN trunk type. The following parameters are typically used when provisioning a T-1 or E-1 PRI span:

- Interface: ISDN Primary Rate Interface (PRI)
- Frame Format: Extended Super Frame (ESF)
- Line Encoding: B8ZS
- Number of B-Channels: 23 and 30 for Euro
- D-Channel: on channel 24th or Euro PRI it is the 16th
- Line Use: Voice

ISDN PRI provisioning also requires a switch type to be configured for Layer 3 protocols. There are four families of switch type protocols for PRI provisioning:

- AT&T, 4ESS, 5ESS, NII Called NI2 (National Protocols)
- DMS100 and DMS 250 (National Protocols)
- EUROPEAN PRI
- Custom 5ESS IntecomE

Common switch types and Layer 3 protocols include the following switch types for well-known PBX systems:

- Nortel (Meridian): 5ESS Custom NOTE: Gateway must be set to NETWORK
- Lucent (Definity): 4ESS or 5ESS
- MCI: DMS 250
- SPRINT: DMS 100 or DMS 250
- AT&T: 4ESS
- Madge (Teleos) BOX: 5ESS Custom
- Intecom: 5ESS Custom

Common switch types and Layer 3 protocols for IXCs and inter-exchange carriers include the following:

- AT&T: 4ESS
- MCI or SPRINT: DMS250
- When connecting to a local CO switch use the following:
- 5ESS (5E8 or 5E9)
- DMS 100
- NII

- Hunt Sequence: Float, Flex, or Fixed
- Out-pulse of digits: 4 is standard, but can be from 1 to 23 digits

Phones per Site and Phone Features

You will need the number of currently supported phones to properly size the CallManager platforms. You should identify phones that will convert to VoIP and potentially some analog phones for emergency and fax backup.You should also know the required phone features, which may include the following:

- Speaker capability
- Mute
- Call hold
- Call park
- Call transfer
- Calling line identity
- Calling party name
- Multi-party conferencing

Examining the Existing Dial Plan

Examine the existing dial plan architecture to understand the required call routing, abbreviated dialing, and route-group features for IP telephony migration. Call routing is used for PSTN or offnet access. Features associated with call routing include:

- Redundant or back-up paths (transparent to the user)
- Emergency dialing call patterns
- Automatic call distribution
- Call blocking where individual groups or numbers have limited offnet access.

Automatic call distribution allows many agents to answer calls from one published number. Call blocking is used to limit access to certain numbers such as 900 toll numbers or long distance PSTN access from building lobby phones. Abbreviated dialing is used to reduce the number of digits required for extension calls. In many cases, local extension dialing has been reduced to 4-digit numbers.

Questions for the IP Telephony deployment include:

- Will the organization use existing or distributed dial plans among multiple sites?
- Are there number ranges to be reserved for PBXs? If so, what are they?
- Are there number ranges to be reserved for analog phones? If so, what are they?

Use the following tables to document the existing dial plan:

Table 3-5, Part 1 Dial Plan Details

Site Name	IP Phone Dial Plan (indicate how many digits for on-net dialing)	PBX Dial Plan			Analog Dial Plan	
		PBX Dial Plan (indicate how many digits for dialing)	PBX	PBX Gateway	Analog Phone Dial Plan (indicate how many digits for dialing)	Analog Gateway
Site 1						
Site 2						
Site 3						
Site n						

You should also consider the following issues:

- What is the local PSTN access dial plan?
- Besides local PSTN access, is there a cellular network to be included in the deployment?
- What access code is used at each site for routing local off-net calls to PSTN?

To answer these questions, you can list all local PSTN calling patterns for each site and the route option for each calling pattern (route this pattern or block this pattern) in the following table format:

Note There could be multiple dial patterns for local PSTN access for non-North America dial plan areas.

Table 3-5, Part 2 Dial Plan Details

Site Name	Local PSTN Dial Pattern	Local PSTN Access Code	Route Option (route or block)	Outbound PSTN Gateway (list which gateway routes this dial pattern)
Site 1	Pattern 1			
	Pattern 2			
	Pattern 3 etc.			

Table 3-5, Part 2 Dial Plan Details

Site Name	Local PSTN Dial Pattern	Local PSTN Access Code	Route Option (route or block)	Outbound PSTN Gateway (list which gateway routes this dial pattern)
Site 2	Pattern 1			
	Pattern 2			
	Pattern 3 etc.			
Site 3 etc.	Pattern 1			
	Pattern 2			
	Pattern 3 etc.			

You should also consider the following issues:

- What is the Long Distance PSTN Access Dial Plan?
- Besides PSTN, is there a cellular network to be included in the deployment?
- Is it required to deploy long distance PSTN call toll bypass within the IP Telephony deployment sites (i.e., route long distance PSTN calls to the PSTN gateway in the destination city so that the call is via VOIP instead of PSTN)? Or is it required just to route all long distance off-net calls to PSTN?

To answer these questions, you can list all access codes, long distance calling pattern (including cellular net, if any, and both domestic and international calling patterns) for each city and the route option for each calling pattern (route this pattern or block this pattern) in the following table format:

Table 3-5, Part 3 Dial Plan Details

Site Name	Site Area Code	Long Distance PSTN Calling Pattern	Long Distance PSTN Access Code	Route Option (route of block)	Outbound PSTN Gateway (list which gateway routes this dial pattern)
Site 1	Pattern 1				
	Pattern 2				
	Pattern 3 etc.				

Table 3-5, Part 3 Dial Plan Details

Site Name	Site Area Code	Long Distance PSTN Calling Pattern	Long Distance PSTN Access Code	Route Option (route of block)	Outbound PSTN Gateway (list which gateway routes this dial pattern)
Site 2	Pattern 1				
	Pattern 2				
	Pattern 3 etc.				
Site 3 etc.	Pattern 1				
	Pattern 2				
	Pattern 3 etc.				

You should also list all special call routing and call distribution requirements here, such as:

- E-911 calling
- 900 call blocking
- ACD (Automatic Call Distribution)

Inbound Dial Plan

For each site, list the incoming called numbers for PSTN users to reach IP phone users.

Table 3-6 Inbound Dial Plan

Site Name	DID or Two-stage Dialing (specify IVR, auto-attendant)	Inbound PSTN Gateway (which PSTN gateway routes incoming calls)	DID Number Range or IVR/AA Main Incoming Called Number
Site 1			
Site 2			
Site 3			
Site 4 etc.			

Fax Requirements

CallManager versions 3.0.1 or later support fax-relay and modem pass-through on certain gateways. You must then define fax-relay requirements and modem pass-through requirements by identifying fax machines, locations, and fax numbers. You may potentially need to identify modems for a modem pass-through solution. You can use the following table to identify fax machines, modems, and information required to design the IP Telephony-based VoIP solution in this area:

Table 3-7 Fax Details

Device	Location	Current Number	PBX or Analog
Fax			
Fax			
Modem			
Modem etc.			

Evaluating and Documenting the Existing Power/Cabling Infrastructure

Another aspect of successful IP Telephony deployment and a high availability voice solution is power and cabling infrastructure. Traditional voice environments typically have well-planned power and cabling systems with UPS power backup and PBX-powered phones. This solution helps to create a more highly available voice implementation. To provide a similar high availability solution, the organization may need to better plan for a highly available cabling and power infrastructure for data equipment.

Without UPS power, organizations can expect to have approximately 1.66 hours of non-availability due to power alone. Refer to the following APC website for detailed information: http://www.apcc.com/go/machine/cisco/.

Cabling infrastructure issues can also cause availability problems due to poor installation practices, patch cord management, non-hierarchically structured installations, and non-standards-based installations. To ensure that you meet the availability requirement, you should understand your cabling infrastructures and plan for potential upgrades.

The first step is to examine the existing cabling and power infrastructure to ensure that power and cabling infrastructures are capable of handling IP Telephony requirements. Refer to the following locations for more information on power:

- Cisco IP Telephony Voice and Video Solutions

  http://www.cisco.com/warp/public/779/largeent/avvid/solutions/powerpro.html

- APC home page

  http://www.apcc.com

Basic questions that may help determine the current infrastructure readiness include:

- Does the building wiring conform to EIA/TIA-568A? Technical Services Bulletin 40 (TSB-40) defines the installation of category 5 wiring systems. TSB-67 defines the testing criteria to ensure compliance. If the wiring does not conform, contact your wiring contractor for testing and potential upgrade.
- Does your organization comply with National Electric code for powering and grounding sensitive equipment? If not or you are unsure, contact APCC for a power audit to determine the compliance and availability characteristics of your environment.
- Does your organization comply with the more rigorous IEEE 1100-1992 standard for recommended practices of grounding and powering sensitive equipment? If not or you are unsure, contact APCC for a power audit to determine the compliance and availability characteristics of your environment.
- Does the organization have standards for data center and wiring closet power that include circuit distribution, available power validation, redundant power supply circuit diversity, and circuit identification?
- Does the organization use UPS and/or generator power in the data center, wiring closet, phone systems, and internetworking devices?
- Does the organization have processes to SNMP manage or periodically validate and test backup power?
- Does your business experience frequent lightening strikes? Are there other potential natural disasters?
- Is the wiring to your building above ground?
- Is the wiring in your building above ground?
- Will the organization determine power draw, plug-type, and heat output for IP Telephony-sensitive equipment before installation to ensure adequate power, a smooth installation, and adequate power backup?

  Refer to the following power sizing guide for more information: http://www.cisco.com/warp/public/779/largeent/avvid/solutions/powerpro.html.

Data Center Power Requirements

CallManagers and gateway devices should be used in data center environments. Evaluate the existing data center in terms of available main power, UPS power, power plug compatibility, and heat dissipation for the potential IP Telephony equipment. You can use the following table to help determine overall IP Telephony power requirements and data center power and cooling requirements. You can find power draw in Watts, operating voltages, and plug type at the following APC website: http://www.apcc.com/template/size/apc/cisco_int/index.cfm.

You can locate heat dissipation, line voltage, and other environmental information in the Cisco IP Telephony data sheets. These documents can be found at the following location: http://www.cisco.com/warp/public/752/ds/english/iptel.html.

Table 3-8 Power Requirements

IP Telephony or Internetworking Device	Power Requirements in Watts	Dual Power Supply?	Operating Voltages	Plug Type	Heat Dissipation
MCS-7835	120	No	100, 120, 200, 208, 230	NEMA 5-15P	1475 BTU/hr.
Device 2					
Device 3					
Device 4					
Device 5					
Device 6					
Device 7 etc.					

Wiring Closet Power

Wiring closet power may require careful planning due to the use of Cisco Inline Power™ and the addition of wiring closet UPS systems. This will help ensure high availability to end phones. Inline Power requires space planning for powered patch panels. The organization should complete wiring closet worksheets, similar to the data center worksheet, to define power, UPS, and cooling requirements.

Table 3-9 Wiring Closet Power

IP Telephony or Internetworking Device	Power Requirements in Watts	Dual Power Supply?	Operating Voltages	Plug Type	Heat Dissipation
MCS-7835	120	No	100, 120, 200, 208, 230	NEMA 5-15P	1475 BTU/hr.
Inline Patch Panel	175	No	100, 120, 200, 208, 230	NEMA 5-15P	
Device 3					
Device 4					
Device 5					

Table 3-9 Wiring Closet Power

IP Telephony or Internetworking Device	Power Requirements in Watts	Dual Power Supply?	Operating Voltages	Plug Type	Heat Dissipation
Device 6					
Device 7 etc.					

IP Telephony Availability Requirements

You should design the IP Telephony network, infrastructure, and support services with targeted availability requirements. Availability planning is useful for several reasons:

- You can use availability as an overall SLA for the voice/data service.
- You can use availability modeling or measurement to determine the best availability level based on the cost of downtime, potential analysis, and a simple ROI (return on investment) calculation.
- You can use availability measurement in a quality availability improvement process to improve the level of service.

Cisco views availability as a combination of six major factors:

- Hardware availability
- Software reliability
- Link/Carrier availability
- Power/Environment availability
- Network Design reliability
- User error and network support processes

Each of these issues may impact different parts of the network in different ways. It is therefore useful to define availability requirements and models for different areas of the network. These may typically be the LAN, WAN, data center, or network core. Cisco currently has general availability classifications that correspond to the business requirements and cost of downtime experienced by the organization. These general classifications are:

- Reliable networks—Availability goal is approximately 99.5% over time (education and government).
- High availability networks—Availability goal is approximately 99.99% over time (high tech, manufacturing, and service).
- Non-stop networks—Availability goal is typically 99.999% and higher over time (financial or some medical environments).

Reliability block diagrams help an organization model availability requirements. Since each availability factor can occur independently of others, the factors are multiplied together to achieve a final result. The result is that if one area is weak, overall availability will be affected more severely. See the following example:

Table 3-10 Reliability Matrix

Network Area	Hardware Reliability	Software Reliability	Link/Carrier Reliability	Power Environment	Network Design	User-error Process	Overall Estimated Ability
Core	HA	HA	HA	HA	non-stop	HA	HA
LAN	HA	HA	HA	HA	non-stop	HA	HA
WAN	Reliable	HA	HA	Reliable	non-stop	HA	Reliable

For instance, hardware reliability typically uses MTBF (mean time between failure) analysis combined with MTTR (mean time to repair) to better model theoretical hardware reliability. In other cases, precise modeling is not possible, but general overall characteristics of the “best practices” have been identified. The following sections provide more detail regarding the best practices within each of the three availability types for each availability factor.

Hardware Reliability

You should measure the reliability of hardware in the network given the network topology, amount of redundancy, and expect time required to repair broken hardware. You can perform hardware reliability calculations using the MTBF of each device (and device modules) and the MTTR for hardware replacement. Hardware reliability for the IP Telephony solution should also include the path between the IP phone and the CallManager, as well as the path between the IP phone and the called party.

Your Cisco account team can provide precise hardware modeling using the expected MTTR. The account team will then contact manufacturing quality to determine the MTBF for each device and module so that they can perform a path analysis on a case-by-case basis. MTBF information is based on the BellCore standard for component quality. BellCore has identified the expected lifetime of more than 500,000 components that are used to manufacture Cisco modules and chassis.

The MTTR is critical in overall hardware reliability. Organizations relying on standard Cisco SmartNET hardware replacement can expect an average of 24-48 hours for hardware replacement, which will lower the availability of the overall solution.

The following table will help characterize the availability of the IP Telephony solution given different areas of the network, replacement times, and the amount of hardware redundancy:

Table 3-11 Hardware Reliability Matrix

	Reliable Networks	HA Networks	Non-stop Networks
Network Core	• Non-redundant simple network core with four hour MTTR • Redundant complex network core with 24 hour MTTR	• Redundant core required • 4 hour MTTR recommended	• Redundant core required • 1-2 hour MTTR required
User/IP Phone/IP Telephony Components	Non-redundant 8 hour MTTR for phone, IP Telephony, and local access switch components	• Redundant CallManager required • Redundant gateways required • Non-redundant IP phone • 4 hour MTTR recommended	• Redundant IP phones • Redundant network hardware infrastructure • Redundant CallManagers • Redundant gateways • 1-2 hour MTTR required
WAN	Non-redundant 8 hour MTTR for phone, IP Telephony, and local access switch components	• Redundant CallManager and gateway required • Complete redundant hardware path • 4 hour MTTR recommended	• Distributed redundant call processing required • Redundant hardware paths required • 1-2 hour MTTR required

Software Reliability

Organizations cannot easily control software reliability since software quality is primarily the responsibility of the vendor. Cisco strives to release only high quality software at or above 99.999% reliability. However, in many cases, early deployment software and early release software falls short of this goal due to unexpected and untested traffic patterns or load related issues.

Cisco IOS also has several classifications of software that correspond to the expected reliability. These include GD (general deployment), LD (limited deployment), and ED (early deployment). In addition, the IOS may have untested interim releases and experimental releases. GD code is considered highly reliable and generally has a proven track record of 99.999% availability.

Some organizational processes contribute to higher software reliability and availability within the organization. The first is software version control. This practice involves maintaining only a few versions of software on the network that have proven track records and that have been tested or piloted within the network to prove reliability.

Another best practice is software testing. Software testing includes feature testing and "what-if" testing to determine the software impact to the existing environment. Cisco offers a testing service called Network Verification Services (NVS) that can help organizations better test their software, hardware, and network designs. Refer to the *Network Verification Service Option* document for more information on the NVS service. This document can be found at the following location: http://www.cisco.com/warp/public/cc/serv/mkt/sup/ent/nsa/welcome/nsan_ov.htm.

NVS is currently only available to ANS or BES customers. Other testing tools and modeling tools are also available. These tools can capture packets, generate traffic, and even insert delay. It is especially important to create a lab where you can adequately test products and features prior to deployment. One popular test device is the SmarBits traffic generator. Refer to the following Spirent Communications SmartBits website for more information:
http://www.netcomsystems.com/.

You will also need processes to manage overall network consistency including topology, software versions, configuration, and features. Consistency within the software configuration can always help contribute to availability as less code is exercised and fewer opportunities for problems can occur.

Table 3-12 helps to identify factors involved in overall software reliability at the various levels.

Table 3-12 Software Reliability Matrix

Reliable Networks	High Availability Networks	Non-stop Networks
• General Release software only • Software version control recommended • Lab testing or solution pilot recommended • Configuration consistency required	• General Release software only • Software version control required • Testing including what-if analysis, feature testing and load testing required • Configuration consistency required	• General Deployment or older proven software required • Software version control required • Testing including what-if analysis, feature testing, and load testing required • Configuration consistency required

Link/Carrier Availability

The following factors can contribute to link/carrier availability:

- Campus fiber installation quality, redundancy, and geophysical diversity
- Building riser installation quality, redundancy, and diversity
- Cable testing and validation
- Patch cord management and cable labeling
- Cabling installation age
- Building entrance facilities and diversity
- Local loop carrier resiliency, redundancy, and diversity
- Long distance carrier resiliency, redundancy, and diversity
- Bandwidth redundancy

Many of these factors are controllable within an organization with the possible exception of local loop and long distance resiliency, redundancy, and diversity. Many carrier infrastructures, especially in third-world countries and parts of Asia, have limited resiliency and almost no redundancy or diversity. Infrastructure repairs can also take days instead of hours or minutes. In developed countries, organizations often have the opportunity for path diversity, multiple carriers, and higher overall availability and yet failures still can occur that affect redundant configurations. This is a because the entire circuit path is not diverse and redundant. In many cases, complete redundancy and diversity is still not an option. The organization should strive to understand cable paths and single points of failure within the local loop system and carrier infrastructure to help understand all the potential availability issues.

The following table will help characterize link/carrier infrastructure requirements at the three availability levels:

Table 3-13 Link/Carrier Matrix

Reliable Networks	High Availability Networks	Non-stop Networks
• Patch cord management and cable labeling recommended • Cabling infrastructure younger than 10 years recommended • Local loop resilient infrastructure required • Long distance resilient infrastructure required	• Copper Cabling infrastructure follows EIA-TIA 568 standard and is tested in conformance with TSB-67 • Fiber cabling tested to ensure DB loss within spec for required media technology • Redundant riser and campus cabling infrastructure recommended • Patch cord management and cable labeling required to ensure cable traceability and to ease hardware replacement (if necessary) • Campus core link diversity recommended • WAN long distance carrier redundancy and diversity required • Significant Local Loop resiliency and redundancy and bandwidth redundancy recommended	• Copper Cabling infrastructure follows EIA-TIA 568 standard and is tested in conformance with TSB-67 • Fiber cabling tested to ensure DB loss within spec for required media technology • Redundant riser and campus cabling infrastructure required • Multiple building entrance facilities required • Patch cord management and cable labeling required to ensure cable traceability and to ease hardware replacement (if necessary) • Campus core link diversity required • WAN long distance carrier redundancy and diversity required • Local loop redundancy and geophysical diversity required • Bandwidth redundancy required

Power/Environment

The following factors *affect* power and environment-related availability:

- Power backup systems
- Network management systems used to monitor UPS and environmental conditions
- IP Telephony equipment
- Environmental conditions
- The processes used to provision equipment and manage power

The following factors *contribute* to power/environment availability:

- Environmental cooling and temperature control
- Equipment BTU and determination
- Environmental conditioning provisioning processes
- Power provisioning process to ensure circuit wattage availability and circuit redundancy for redundant equipment and power supplies
- Equipment surge protection
- UPS battery backup systems
- Generator systems
- SNMP or other remote management processes for UPS systems
- Geographic location of equipment where lightening strikes, floods, earthquakes, severe weather, tornados, or snow/ice/hail storms can affect power reliability
- Power cabling above ground
- Construction near or within equipment facility
- Power cabling infrastructure conformance to NEC and IEEE wiring standards for safety and ground control

The following information relates to availability at the three availability levels for different equipment types and corresponds to studies done by the APC corporation (http://www.apcc.com/).

Table 3-14 Power Matrix

	Reliable Networks	High Availability Networks	Non-stop Networks
IP Phones	Inline Power™ with surge protection and 30 minute UPS battery backup recommended	Inline Power with surge protection and 1 hour UPS battery backup recommended	Inline Power with surge protection and 8 hour UPS battery backup recommended
CallManager and gateways	• 30 minute UPS battery backup recommended • Equipment BTU determination and environmental provisioning process recommended • Power provisioning process recommended	• 1 hour UPS battery backup recommended • Equipment BTU determination and environmental provisioning process required • Power provisioning process required • UPS SNMP management process required	• 8 hour UPS battery and generator backup recommended • Equipment BTU determination and environmental provisioning process required • Power provisioning process required • UPS SNMP management process required

Table 3-14 Power Matrix

	Reliable Networks	High Availability Networks	Non-stop Networks
Data Center	• 30 minute UPS battery backup recommended • Equipment BTU determination and environmental provisioning process recommended • Power provisioning process recommended	• 4 hour UPS battery backup recommended • Equipment BTU determination and environmental provisioning process required • Power provisioning process required • UPS SNMP management process required	• 8 hour UPS battery and generator backup recommended • Equipment BTU determination and environmental provisioning process required • Power provisioning process required • UPS SNMP management process required
Internetworking infrastructure	• 30 minute UPS battery backup recommended • Equipment BTU determination and environmental provisioning process recommended • Power provisioning process recommended	• 1 hour UPS battery backup recommended • Equipment BTU determination and environmental provisioning process required • Power provisioning process required • UPS SNMP management process required	• 4 hour UPS battery and generator backup recommended • Equipment BTU determination and environmental provisioning process required • Power provisioning process required • UPS SNMP management process required

Network Design

Factors that contribute to network design availability and reliability include:

- Modular and hierarchical logical design
- Hierarchical IP routing infrastructure supporting IP route summarization
- Consistent hardware, software, and device configuration
- High availability, high performance media and devices
- High availability convergence capabilities at level II and level III
- QoS functionality in the LAN/WAN supporting low voice delay and jitter
- Design testing
- Vendor approval
- Capacity and Performance management processes

- Network change management that promotes network consistency and change validation for higher risk change
- Minimize or eliminate spanning tree with IP routing

The following table provides basic network design characteristics that you need at the three defined availability levels:

Table 3-15 Design Characteristics for Availability Levels

Reliable Networks	High Availability Networks	Non-stop Networks
• Modular and hierarchical logical network design required • Consistent hardware, software, and configuration recommended • Hierarchical IP routing infrastructure supporting IP summarization recommended • High performance media and devices recommended • QoS in LAN at level II/III recommended • QoS in WAN with Call admission control recommended • Fast converging spanning tree parameters recommended • HSRP gateway redundancy features recommended • Fast converging IP routing protocol EIGRP or OSPF recommended	• Modular and hierarchical logical network design required • Consistent hardware, software, and configuration required • Hierarchical IP routing infrastructure supporting IP summarization required • High performance media and devices required • QoS in LAN at level II/III required • QoS in WAN with call admission control required • Fast converging Spanning Tree parameters required • HSRP gateway redundancy features required • Fast converging IP routing protocol EIGRP or OSPF required • Change management validation and vendor approval recommended • Capacity and performance management processes for baselining, exception management, and what-if analysis recommended	• Modular and hierarchical logical network design required • Consistent hardware, software, and configuration required • Hierarchical IP routing infrastructure supporting IP summarization required • High performance media and devices required • QoS in LAN at level II/III required • QoS in WAN with Call admission control required • Fast converging Spanning Tree parameters required • HSRP gateway redundancy features required • Fast converging IP routing protocol EIGRP or OSPF required • Change management validation and vendor approval required • Capacity and performance management processes for baselining, exception management, and what-if analysis required

User Error and Process

User error and process are major contributors to non-availability. The Gartner Group maintains that approximately 40% of all availability issues are user error and process related. To achieve higher levels of availability, higher levels of expertise and processes are needed to manage the network. There are four goals of network management processes:

- Fault or problem detection and speedy resolution
- Proactive problem identification and resolution before fault condition occurs
- Consistent modular, hierarchical architecture with consistent versions and configuration
- Successful migration to new solutions and technologies

To achieve these goals, a large organization needs well-defined processes associated with the planning, design implementation, and operation of the network. People, process, and expertise must also be applied to processes that may impact availability in relation to the fault, configuration, accounting, performance, and security of the network infrastructure and the IP Telephony solution.

The required expertise needed to manage IP Infrastructures and the IP Telephony solution includes:

- VoIP Technology—H.323, VoIP concepts and protocols, RTP
- Network Architecture—TCP/IP, IP subnetting, routing protocols including EIGRP, OSPF, RIP, BGP
- Services/Peripherals—DNS, DHCP, TFTP, WEB server, QoS
- Access—Signaling (ISDN-PRI, EIR2, CCS and CAS), FXS, FXO, ground/loop start
- IP Telephony—NT 4.0 server administration, IOS, Codecs including G.711, G.729
- General skills—project management, troubleshooting, security, network management, server administration
- Recommended experience—Cisco VoIP configuration, Catalyst Ethernet switches, Cisco IOS, QoS techniques, knowledge and understanding of voice dial plans, exposure to traditional PBXs

We recommend the following processes in relation to network infrastructures and the IP Telephony solutions:

- Well-defined architecture and design process that includes testing and validation
- Solution implementation templates that define standard connectivity, hardware devices/modules, configuration, software versions, out-of-band management practices, and network management tool update requirements
- Operational support plan for new technology that defines training requirements, problem types, problem priorities, resolution goals, escalation paths, and support personnel responsible
- Consistent implementation processes that ensure consistent deployment. Process may include staging, physical design, documentation and approval, as-built documentation, and step-by-step implementation guide
- Fault management processes detecting link-down/device-down conditions
- Fault management processes detecting critical Syslog message or exception conditions and providing speedy resolution
- Hardware sparing plan that ensures consistent, quick hardware replacement
- Performance and capacity planning processes that provide baselines, trends, exception condition identification, and resolution; what-if analysis to test or validate capacity or performance impact of network changes
- Help Desk systems and ticket generation/accountability for all problems

- Security policies and operational processes
- Configuration management practices for hardware/software version control and device configuration management
- TACACS device access

Planning for WAN Deployments

IP telephony WAN deployments require significant planning. This section explains a detailed process that includes collateral and checklists needed for IP telephony WAN deployments using centralized call processing. It only covers the requirements for planning WAN connectivity and does not include general centralized call manager design steps, gateway design steps, or operational planning requirements. WAN analysts should refer to this section to analyze voice requirements and build a WAN environment suitable for IP telephony.

The process covers WAN capacity planning and current network analysis, voice capacity planning, upgrade planning, deployment, and validation for WAN environments. In each of these areas, a descriptive checklist of important critical success factors is provided, along with required planning collateral and references.

The six major steps for WAN deployment are:

Step 1 Collecting Information on the Current WAN Environment

Step 2 Determining Voice Bandwidth Requirements

Step 3 Analyzing Upgrade Requirements

Step 4 Performing Upgrades and Implementing Tuning

Step 5 Assessing Results

Step 6 Operational Turnover and Production

Collecting Information on the Current WAN Environment

The first step in planning WAN deployments is to collect information from the WAN and WAN devices. Use this information to analyze gaps after determining bandwidth and device requirements. Collect several categories of information, including:

- the existing WAN topology, which includes includes logical design information and bandwidth subscription rates.
- device information, which includes includes router models, memory, CPU, interface card modules/versions and software versions.
- resource utilization, which includes peak utilization and, to a lesser extent, average utilization across WAN links. It also includes delay, packet loss, and jitter information for each WAN link to calculate the delay and jitter budget.

WAN Topology

Topology information should be required for the central site and remote sites that implement IP telephony. You may also need topology information for other sites when shared circuits are implemented on one physical interface, which is common with frame relay central locations. Create an accurate map showing the central location and all WAN sites that will be implementing IP phone connectivity. The topology map should include all WAN routers supporting primary and backup WAN connectivity, service providers, interfaces, and media. Also include information such as media, endpoints, devices, device interfaces, WAN providers, WAN subscription rates, and redundant connectivity options. WAN subscription rates are important and include port speeds, committed information rates, class of service, and how ports are shared with other virtual circuits. Use the following table to help collect the required WAN topology information.

Table 3-16 WAN Topology Information

WAN Edge Location	Primary WAN Media	Backup WAN Media	WAN Subscription Rates	Device/Interface Endpoints
San Jose to Denver	Frame Relay	None	SJ: T-1 shared port Denver: T-1 port SJ to Denver: 64kCIR Denver to SJ: 64k CIR	Sj-gw1: S3/1 Dvr-gw1: S0
San Jose to Los Angeles	ATM w/ABR class of service	Frame-relay	ATM: 3 megabit ABR SJ: T-3 shared FR Denver: T-1 port SJ to Denver: 768 CIR Denver to SJ: 512 CIR	Sj-gw2:ATM3/0 La-gw1: ATM1 Sj-gw1:S4/1 La-gw1:S2
San Jose to Seattle	Pt-to-Pt T-1	None	N/A	Sj-gw1: S3/0 Stl-gw1: S0

Device Information

You must evaluate device hardware and software upgrade requirements. Device issues involved with IP telephony include software features, hardware support issues, CPU capabilities, and memory. The following table illustrates the required information for potential upgrades for WAN router equipment.

Note LAN equipment is not evaluated in this chapter and basic switched 10/100 connectivity is assumed for all remote IP phones and central IP telephony resources.

Table 3-17 Router Upgrade Information

Device Name	Hardware Platform/CPU/Memory/Flash	Software Version	Hardware Modules
Sj-gw1	7500/RSP4/128MB/16MB	12.0(13)	VIP2-40, PA-8, PA-FE-TX
Sj-gw2	7206/NPE-200/64MB/20MB	12.0(13)	PA-A1, PA-2FE-TX
Dvr-gw1	2620/50Mhz/64MB/16MB	12.0(13)	NM-1FE-TX, NM4T
La-gw1	3640/100Mhz/64MB/20MB	12.0(14)	NM-1FE-TX, NM-1A-T3, NM4T

Stl-gw1	2620/50Mhz/64MB/16MB	12.0(14)	NM-1FE-TX, NM4T

Resource Utilization

Resource utilization includes issues like bandwidth utilization, CPU utilization ,and memory utilization with bandwidth being the primary concern in IP telephony deployments. You should also investigate CPU and memory. Use this information to determine total resource requirements after estimating the IP telephony requirements.

Determine bandwidth utilization by performing a baseline on potential IP telephony WAN links. In the previous example, this includes frame relay links between San Jose to Denver and San Jose to Los Angeles, the ATM link between San Jose and Los Angeles, and the T-1 link between San Jose and Seattle. The baseline should include peak bandwidth utilization and busy hour utilization. Peak utilization is the peak five minute average utilization over the baseline period. Busy hour utilization is the peak hour utilization during the baseline period. The baseline period should be two to seven days. We normally recommend at least one week with three weeks being ideal. Busy hour utilization is important because it mimics telecom capacity planning time periods and helps to understand peak hourly utilization requirements during busy periods. Average utilization is less useful because it generally includes weekend or nighttime utilization that misrepresents business periods.

Pay close attention to detail when collecting and analyzing the peak link utilization value since several factors can misrepresent the baseline results. First, link utilization must be represented by both inbound and outbound utilization. WAN links are generally full-duplex transmissions where an equal amount of capacity is available inbound and outbound. For example, a point-to-point T-1 serial link that uses SDLC link layer protocol provides roughly 1.54 megabits/sec in each direction. Ethernet, on the other hand, is a shared media. This can be a problem with some network management tools because they may average the inbound and outbound delta traffic volume or simply compute utilization as the greater of the two. Optimally, WAN link utilization should be provided in each direction. If your tool does not provide this, investigate the method for computing utilization. On WAN links, it is not uncommon to have most utilization occurring outbound to the remote location since this is the primary location of servers and information available to the remote office. Remember that voice traffic can be affected in both directions. The important point is that you should ideally collect peak or hourly average utilization for both inbound and outbound traffic, especially if different subscription rates are applied to the circuit, which may be the case with frame relay. This is especially important with ATM or frame relay point-to-multipoint network topologies since most of the traffic is outbound.

Another issue is determining the total bandwidth available for the link. For example, how does SNMP or the performance management tool know the available bandwidth? In some cases, the SNMP tool simply polls the MIB II **IfSpeed** value to determine the bandwidth. This can be a problem on WAN links because the **IfSpeed** value represents the bandwidth of the physical interface. With ATM or frame relay, the actual bandwidth may be quite different because of rate limiting or multiple virtual circuits configured on one physical interface. The result is that if the tool uses the **IfSpeed** MIB II variable to compute available bandwidth, you may need to re-configure using the **bandwidth** command to override the physical interface bandwidth contained in the default **IfSpeed** variable.With ATM, configure to the subscribed bandwidth. With Frame Relay, configure to the CIR value because frame relay traffic shaping to CIR is recommended for frame relay links with IP Telephony. This may be a problem when different CIR values exist for inbound and outbound traffic. Also, remember that the **bandwidth** command is also used for routing metrics. With many tools, such as CiscoWorks 2000 TrafficDirector, the bandwidth is input directly into the tool and the **IfSpeed** MIB II variable is not used in the utilization calculation. In these cases, enter the subscribed bandwidth for ATM or rate-limited links and enter the CIR bandwidth for frame relay links.

Another important consideration is that peak utilization, as measured by SNMP for a five-minute polling interval, is not actually a peak. Data traffic is generally very bursty. The graph below shows how the five-minute polling value returned (shown in green) is an average for the five-minute period. The red line shows a potential threshold for problems. The blue line is a theoretical utilization curve that shows the threshold was exceeded during two polling intervals. However, the five-minute average was below the threshold. Therefore, be careful when interpreting the five-minute peak value. Tests with protocol analyzers show that data generally bursts 40% above the five-minute returned value some time during the polling period.

Figure 3-4 Resource Utilization Graph

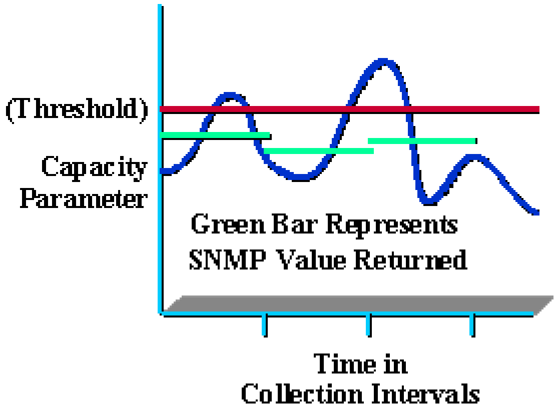

Organizations wishing to collect baseline network utilization have a variety of options. A performance management tool is usually available within the organization. If not, the organization may consider tools such as CiscoWorks 2000 TrafficDirector or the Cisco WAN SwitchProbe. You can also outsource network baselines to a variety of network service or consulting firms. Cisco also has a VoIP audit that provides peak and hourly average utilization for WAN links.

Organizations should also consider a CPU and memory baseline for the network. This is important because you may need additional CPU and memory for configurations such as QOS, rate-limiting, and link-fragmentation interleaving. You can also test CPU and memory requirements in a lab environment to help ensure that adequate resource are available for any given platform and memory allocation.

The baseline utilization information is used to determine bandwidth requirements for data and IP telephony. The device information is used to determine upgrade requirements once overall requirements have been identified.

Performance Information

A baseline of current performance information, including delay, jitter, and packet loss, is valuable for budgeting current jitter and delay for the WAN link. This helps identify potential delay or jitter problems and estimate expected delay and jitter following network upgrades:

- Delay is the amount of time in milliseconds it takes for a packet to traverse the WAN link.
- Jitter is the range of delay seen in packets traversing the WAN link.
- Packet loss is the number or percentage of lost packets within the network due to queuing overflow or errors.

You can measure delay, jitter, and packet loss in a few ways. You can use ping scripts to determine minimum and maximum delay, jitter, and packet loss. In this case, packets sent should approximate voice packets in size (about 60 bytes).

You may also have a tool such as Internet Performance Manager (IPM), part of the WAN management suite of CiscoWorks 2000, which can send voice type packets to determine delay, jitter, and packet loss. The table that follows is an example performance baseline for WAN links.

Table 3-18 Sample Baseline Performance

WAN Link and Direction	Minimum Delay	Maximum Delay	Potential Jitter	Packet Loss Percentage
San Jose to Denver	20 ms.	100 ms.	80 ms.	.01%
Denver to San Jose	20 ms.	40 ms.	20 ms.	.003%
San Jose to Los Angeles	8 ms.	29 ms.	21 ms.	none
Los Angeles to San Jose	8 ms.	18 ms.	10 ms.	none
San Jose to Seattle	26 ms.	120 ms.	94 ms.	none
Seattle to San Jose	26 ms.	102 ms.	76 ms.	none

Determining Voice Bandwidth Requirements

Voice bandwidth requirements depend on a number of parameters, including the sampling rate, codec, link type, header compression techniques, and the number of simultaneous voice calls. This section identifies voice bandwidth requirements for WAN links.

The most important issue is the number of simultaneous calls that are permitted across the WAN link. The organization can limit the number of calls via call admission control; however, with a centralized call processing model, attempts that are made at calling capacity will simply busy out. This is the same behavior one would expect with a centrex solution when an outside line is not available. In addition, no exact rules exist for the percentage of lines to phones since it depends on business requirements. The organization can start by determining how many outside lines and phones the remote location currently has implemented. The organization may also be able to collect statistics on how frequently a busy-out occurred with the current solution to determine additional needs. Since you are implementing a new solution, this is a good time to investigate whether the number of lines available at the remote location is sufficient for the remote location.

Once you determine the maximum number of calls allowed across the WAN link, you can better understand the additional bandwidth requirements for that link. The next step is to understand the amount of bandwidth required for one voice call. Unfortunately, if you ask for the bandwidth requirements for a voice call, you are likely to get a variety of different answers, all of which may be correct for any given situation. It is therefore important to understand the components of a VoIP packet and the different variables that affect overall utilization. The following diagram identifies the components in a VoIP packet. In addition to packet size, the sampling rate will affect bandwidth requirements. The configuration of VAD (voice activity detection), will also impact bandwidth requirements. VAD reduces bandwidth requirements using the theory that in any given voice call, only one party is talking at a time. Using this theory, periods of non-activity are then not transmitted across the link. VAD is expected to save overall bandwidth by as much as 50%. Planners must be careful however because at any one time 100% of the expected bandwidth may be required for one voice stream.

Figure 3-5 VoIP Packet Components

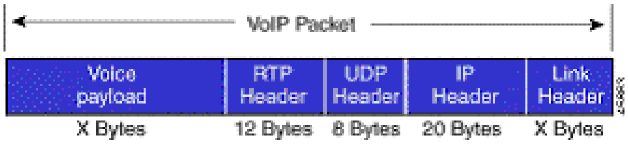

The organization can start by understanding payload requirements. Two different encoding methods are currently available: G.711 and G.729A. In general, Cisco recommends G.711 encoding for LAN environments and G.729A across the WAN. When both are tested for voice quality, G.711 slightly edges out G.729 because of the additional payload information. The more important factor is the sampling rate, which is the period of time allocated to encoding information before the packet is transmitted. 20 millisecond sampling rates have higher quality because less voice information is needed for a 20 ms. time slot. 30 millisecond sampling rates have lower voice quality than 20 ms. It is possible to configure sampling rates above 30 milliseconds; however, this is not recommended because it usually results in very poor voice quality. The following chart shows payload requirements for G.711 and G.729A for 20ms and 30ms sampling rates. Note that the sampling rate does not significantly impact bandwidth for the payload. However, when overhead is added, there are significant increases with 50 packets per second rather than 33 packets per second. The bandwidth consumption is also required for each VoIP stream. In any conversation, two such streams are required: one in each direction.

Table 3-19, Part 1 Payload Requirements

Codec	Sampling Rate	Voice Payload in Bytes	Packets per Second	Bandwidth per Conversation
G.711	20 ms.	160	50	64 kbps
G.711	30 ms.	240	33	63.4 kbps
G.729A	20 ms.	20	50	8 kbps
G.729A	30 ms.	30	33	7.9 kbps

The next table takes into account the additional overhead of headers including RTP headers, UDP headers, IP headers, and link headers. All of the media types below include RTP, UDP, and IP headers at 40 bytes per packet.

Table 3-19, Part 2 Payload Requirements

Codec	Ethernet 14 Bytes of Header	PPP 6 Bytes of Header	ATM 53-Byte cells with 48-Byte payload	Frame-Relay 4 Bytes of Header
G.711 at 50 pps	85.6 kbps	82.4 kbps	106 kbps	81.6 kbps
G.711 at 30 pps	56.5 kbps	54.4 kbps	70 kbps	54 kbps
G.729A at 50 pps	29.6 kbps	26.4kbps	42.4 kbps	25.6 kbps
G.729A at 33 pps	19.5 kbps	17.4 kbps	28 kbps	17 kbps

You can improve bandwidth allocations using RTP header compression and VAD. RTP header compression reduces the size of the RTP header from 12 bytes to 2 bytes. VAD reduces the bandwidth requirement by approximately 50 percent since bandwidth is only allocated to the talking party.

The values below for VAD can be misleading since it is unclear where the speaking party is at any one point in time. For this reason, use caution in simply reducing the bandwidth requirement by a full 50 percent. The following table shows bandwidth requirements for all major media with and without RTP header compression and VAD on a per-stream basis. Remember that any conversation has two streams, one in each direction.

Table 3-19, Part 3 Payload Requirements

Codec	Ethernet 14 Bytes of Header	PPP 6 Bytes of Header	ATM 53-Byte cells with 48-Byte payload	Frame-Relay 4 Bytes of Header
G.711 at 50 pps	85.6 kbps	82.4 kbps	106 kbps	81.6 kbps
With cRTP	81.6 kbps	78.4 kbps	102 kbps	77.6 kbps
With VAD	42.8 kbps	41.2 kbps	58 kbps	40.8 kbps
With cRTP & VAD	40.8 kbps	39.2 kbps	51 kbps	38.8 kbps
G.711 at 33 pps	56.5 kbps	54.4 kbps	70 kbps	54 kbps
With cRTP	54.1 kbps	52.0 kbps	67.6 kbps	51.6 kbps
With VAD	28.3 kbps	27.2 kbps	35 kbps	27 kbps
With cRTP & VAD	27.1 kbps	26 kbps	33.8 kbps	25.8 kbps
G.729A at 50 pps	29.6 kbps	26.4kbps	42.4 kbps	25.6 kbps
With cRTP	25.6 kbps	22.4 kbps	38.4 kbps	21.6 kbps
With VAD	14.8 kbps	13.4 kbps	21.2 kbps	12.8 kbps
With cRTP & VAD	12.8 kbps	11.4 kbps	19.1 kbps	10.8 kbps
G.729A at 33 pps	19.5 kbps	17.4 kbps	28 kbps	17 kbps
With cRTP	16.2 kbps	14.1 kbps	24.8 kbps	13.8 kbps
With VAD	9.8 kbps	8.8 kbps	14 kbps	8.5 kbps
With cRTP & VAD	8.6 kbps	7.6 kbps	12.8 kbps	7.3 kbps

Using the information above, network planners can estimate the bandwidth required for each WAN site. Remember that WAN links are generally full-duplex so an equal amount of bandwidth should be allocated in each direction for one voice call. Be careful when estimating bandwidth using VAD because the values above do not represent the actual required bandwidth in any one direction at a particular point in time.

For example, the Acme Corporation has a remote field site that has 20 permanent employees. The site currently has three Centrex lines and users sometimes busy-out because an outside line is not available. Network planners talked to the provider and found out that the busy-out condition was occurring roughly ten times a day. This was considered unacceptable to the office so the network planner agreed to provide bandwidth for four simultaneous voice calls. The site is currently connected via frame relay. The network planner also tested different compression techniques and decided to use G.729 encoding with cRTP and VAD over frame relay at 50 pps. Using the available information, the network planners found that each call would use approximately 10.8 kbps per stream. However, the planner was a bit uncomfortable with only 43.2 kbps in each direction because VAD does not guarantee that bandwidth requirements will be this low in each direction at one time. The planner decided that to better guarantee voice quality, 64kbps should be allocated across frame relay. The site currently has 64 kbps CIR over frame relay so the planner intends to double CIR to 128k and configure the appropriate QOS, traffic shaping, and link fragmentation interleaving to provide acceptable voice quality.

Analyzing Upgrade Requirements

You now need to consider upgrade requirements for hardware, software, and WAN connectivity to ensure success with IP telephony over the WAN. Hardware upgrades may be recommended to ensure adequate processing power for IP telephony features. Software upgrades may be needed to support IP telephony features. WAN upgrades may be needed to support the additional traffic load. A good place to start is to identify the required software features and a potential software image. This may drive hardware requirements, including memory flash and CPU. Bandwidth requirements may also drive specific hardware features to support higher bandwidth WAN media such as ATM or HSSI.

Investigate IOS software requirements to ensure that a targeted release contains the required feature sets for WAN-based IP telephony. The design portion of this guide includes the recommended IOS features for WAN-based IP telephony. Class based weighted fair queuing (CBWFQ), link fragmentation interleaving (LFI), traffic shaping, compressed RTP (cRTP), and IOS gateway support for location based call admission control are some of the features that may be needed in the targeted IOS version to support centralized call processing. Once you identify these features, research potential versions on Cisco.com at http://cco/warp/public/732/releases/ and http://cco/warp/public/732/Tech/voice/index.html. A tool that may be helpful is the Cisco IOS software feature navigator at http://cco/cgi-bin/Support/FeatureNav/FN.pl. Before targeting a specific release, also investigate potential bugs at http://cco/support/bugtools/. You may also validate the software feature analysis and image selection with Cisco sales and support representatives. The results of the investigation should then be a targeted software release for WAN routers supporting IP telephony.

Before researching hardware platforms, the organization should determine WAN upgrade requirements to ensure that media support requirements will be met. In most cases, simply add current data requirements to anticipated peak voice requirements. If you chooses to starve a portion of data traffic during peak periods, conduct careful research to ensure that data requirements will be continually met. In addition to bandwidth requirements, examine other WAN media considerations:

- Does the carrier or WAN connectivity support the required availability for IP telephony?
- Will redundancy be required across the WAN to support high availability voice?
- Does the WAN connectivity have the required hardware, local loop, and long distance redundancy and diversity to support the required availability level?

With the centralized call processing model, the loss of WAN connectivity means that IP phones in the remote location are down. This problem will be addressed some time in the future in the form of IOS CallManager/gateway support, but until then, maintain redundancy and diversity with WAN connectivity to support high availability IP telephony service in the remote location. Carefully consider if current WAN data availability will be acceptable for IP phone service.

- Can congestion occur within the carrier network? At what subscription levels? Will this affect jitter? Is traffic shaping required to achieve consistent performance?
- Will the WAN connection consistently meet maximum delay requirements?

Analyze WAN carrier connectivity to determine what performance can be expected from the carrier network. With pt-to-pt T1, E1, or T3 links, simply determine if you meet minimum delay requirements. QOS features help ensure that requirements are consistently met for IP telephony. In the case of ATM or frame relay, different carriers provide different levels of service. In some networks, only the committed information rate can be expected to meet consistent performance levels. In others, bursting up to port speed still results in adequate delay and jitter. The organization should talk to the carrier to understand what can be expected in the way of performance and at what subscription levels. This drives traffic shaping, and possibly bandwidth requirements for the carrier network.

When you have analyzed WAN bandwidth, redundancy, diversity, carrier, and media requirements, you may need to investigate carrier services for IP telephony or upgrade current services. Potential services include WAN pt-to-pt (T1, E1, T3 & E3), ATM, frame relay, and fractional T3 services that support HSSI, POS, or frame relay link level protocols.

Next, investigate hardware router platforms and sizing. The 2500 router in the WAN may not support the new image size requirement, new media, or CPU needed to support QOS and other IP telephony services. In general, Cisco recommends the 2600 and 3600 platform for the centralized call processing model. Other platforms may meet requirements, but should be thoroughly tested to ensure that resource requirements are met. The main issues is hardware sizing are:

- WAN routers may need additional CPU for QOS, LFI, and traffic shaping processes under load.
- WAN routers may need additional CPU and memory to support H.323 location-based call admission control.
- WAN routers may need additional CPU to support newly identified routing protocol or redundancy requirements needed for high availability.
- WAN routers may need additional DRAM or flash memory to support the identified image size.
- WAN routers may need additional memory to support routing protocols or features implemented to support redundancy and convergence.
- WAN routers must support newly identified port modules needed for data and voice bandwidth requirements.

When the organization has determined software, hardware, and WAN requirements, the next step is to implement the solution without voice traffic and test the solution prior to production.

Performing Upgrades and Implementing Tuning

At this point, the organization should order, schedule, and implement the required upgrades for IP telephony services. We recommend that all required centralized call processing, WAN, and LAN upgrades occur prior to voice deployment so the solution can be verified prior to live voice traffic. All upgrade changes should be implemented under guidelines for moderate to high risk. If possible, set up and/or stage the upgrades in a lab environment to ensure success in the production network.

Refer to the configuration guidelines for VoIP WAN router configuration at http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/12cgcr/voice_c/vcprt1/vcvoip.htm. There is additional configuration documentation for frame relay configuration at http://www.cisco.com/warp/public/788/voip/voip_fr_frts_frf12.htm.

You may also wish to investigate the QOS policy manager under CiscoWorks 2000, which aids in the configuration and management of policy-based QOS on Cisco routers.

Assessing Results

Before voice traffic is added to the WAN network, validate performance of voice traffic across the WAN links. This helps catch any configuration or performance issues that may still exist, validate the intended design, and ensure production success. There are several assessment methods:

- Live testing
- Internet Performance Manager (bundled with CiscoWorks 2000)
- Cisco VoIP Readiness Net Audit (version II)

Physically test and evaluate voice quality and performance between the WAN site and the central site. This method is the simplest; however, the danger exists that users will be dissatisfied with quality and less willing to identify problems or continue working with the solution if quality is less than adequate. For this reason, it is a good idea to test the solution with a more technical team that has allocated time to objectively identify and track down potential problems.

Another method is the Internet Performance Manager for CiscoWorks 2000. IPM measures network performance based on the “synthetic traffic generation” technology within Cisco IOS software, known as the Service Assurance Agent (SAA). You can configure these agents on performance agents called “collectors” in the router and other computer platforms. Synthetic traffic is generated from SAA devices and performance delay and jitter information collected via SNMP from the SAA collectors. The advantage of testing using this tool is that the tool is helpful following production to troubleshoot IP telephony performance problems. Refer to http://cco/warp/public/cc/pd/wr2k/nemo/ for additional information.

The Cisco VoIP Readiness Net Audit is a one-time audit of the network that shows potential performance problems. The audit is recommended at this stage to quickly pinpoint issues that can impact VoIP across the LAN or WAN. The audit requires a one-time information collection from the network using a Net Audit collector device. Refer to http://www.cisco.com/warp/partner/synchronicd/cc/serv/mkt/sup/ent/avvid/nadit_ds.htm for more information.

Operational Turnover and Production

Define operational turnover and IP telephony production/support of the WAN site prior to going live:

- Detailed network documentation including maps and configurations to operations
- Operations support training
- Support contacts and escalation guidelines
- On-site support and hardware replacement service definition
- Operations troubleshooting guides and support tools
- Service level goals or service level agreements for IP telephony support
- Move, Add, Change procedures for remote site

Note These items are discussed in more detail in the Operations and Implementation sections of this guide.

Operations and Implementation Planning

In addition to design planning, the organization needs to plan implementation and operation processes that will help ensure a smooth implementation. You can accomplish this by identifying the design, implementation, and operations processes that are required to achieve ongoing success of the solution. This includes processes such as capacity planning, network management planning, staffing, and operations planning.

Since IP Telephony is a major technology change for most organizations, you should initiate significant operations planning during the design phase. This helps ensure that you meet operations requirements for the solution and that the operations group for the organization has the necessary resources and levels of expertise to manage the solution. You may require the following types operations and management planning:

- IP Telephony Capacity Planning
- Solution Manageability Requirements
- Staffing and Expertise Requirements
- Operations Support Plan

IP Telephony Capacity Planning

Capacity planning is a critical process for enterprise VoIP migration and overall success. In traditional data environments, capacity planning focuses on link utilization and control-plane resources for network devices, including CPU and memory. IP Telephony VoIP networks require three distinct capacity planning processes:

- CallManager processing requirements

  CallManager processing requirements help to ensure that the CallManager servers have sufficient resources for normal call processing, voice conferencing, and other IP Telephony services. This planning and design process typically leads to an improved network design that can better support the organization’s requirements.

- Network capacity/performance

  Network capacity/performance planning helps to ensure that the network can support the additional VoIP and traffic and, more importantly, that the traffic will meet delay and jitter requirements critical to VoIP networks. This process may also lead to an improved overall design that will better support the VoIP requirements.

- Trunking capacity

  Trunking capacity investigates PSTN trunking requirements that include normal PSTN trunking, voice mail server trunking, and site-to-site trunking where needed to support off-net inter-domain trunking and/or inter-domain overflow.

Organizations can assure sufficient capacity and performance only when organizations investigate and follow through with each of these processes. The purpose of this section is to provide best practices for capacity planning. You should perform these processes during the network planning and design phases of IP Telephony implementation.

This section also discusses capacity planning tools and Cisco IP Telephony support services that can help with these processes, as well as design and configuration guidelines.

CallManager Processing

CallManager processing investigates CPU and memory requirements and supported configurations. In larger applications with CallManager 3.0 and later, servers are configured in clusters. Clustering guidelines include the use of a dedicated server for database replication or publishing and backup CallManager capabilities for the desired number of IP phones. CallManager services are also assigned weights for their impact to CPU and memory requirements within each CallManager configuration.

You must investigate the number of desired IP phones in the configuration and the desired service weight to properly design a configuration that includes the server model and number of servers configured in the cluster. We recommend database publishing for configuring 2500 IP phones or more for a dedicated server for TFTP services.

Refer to the tables in the *Cisco IP Telephony Network Design Guide* to learn more about CallManager resource requirements. These tables can be found at the following location: http://www.cisco.com/univercd/cc/td/doc/product/voice/ip_tele/network/dgclustr.htm.

Planning for Network Capacity and Performance

Network performance and capacity planning helps to ensure that the network will consistently have available bandwidth for data and VoIP traffic and that the VoIP packets will consistently meet delay and jitter requirements. We recommend the following six-step process for network capacity and performance planning:

1. Baseline existing network utilization and peak load requirements.
2. Determine VoIP traffic overhead in required sections of the network based on busy hour estimates, gateway capacities, and/or CallManager capacities.
3. Determine minimum bandwidth requirements.
4. Determine the required design changes and QoS requirements, based on IP Telephony design recommendations and bandwidth requirements (over-provision where possible).
5. Validate baseline performance.
6. Determine trunking capacity.

Each of these steps are described in detail below:

Step 1 Baseline Existing Network Utilization

Baselining network utilization helps to determine the current traffic load and data traffic requirements for a combined VoIP and data architecture. You should perform this step for major distribution, backbone, and required WAN links. If you have a relatively homogenous environment, you may perform a sample baseline from representative links rather than collect large amounts of data for an entire network. You normally use two types of utilization statistics to describe link utilization: peak utilization and average utilization. Both of these measurements have limitations and you should carefully evaluate them.

- Peak Utilization

  You normally collect peak utilization data by SNMP polling of the “average” utilization over a short time interval. In most cases, the time interval is five minutes but may be as high as fifteen minutes. The problem with this information is that it is still an average of utilization over the specified time period. Studies have been done showing that over a five minute period, the actual peak utilization is 40% higher than the reported result. This means that a reported peak utilization of 70% may indicate points in time during the five minute period where utilization is 98%. In VoIP environments, this may already impact delay and jitter of voice packets. Another issue is that the information may not represent duplex links where utilization may be quite different in each direction.

- Average Utilization

  Average utilization may not provide complete information for other reasons. For example, data traffic may be much higher during peak periods, such as mornings and afternoons. If the average utilization is reported for a 24-hour period, then the reported result may not show heavily utilized periods during the day.

To be more useful, peak and average utilization should be more closely tied to actual peak periods. Average utilization would be much more valuable if it depicted "busy hour" utilization for the busiest periods during the day. Peak utilization would be more valuable if it could determine the actual peak periods, perhaps seconds or even milliseconds in duration. Unfortunately, peak utilization is not easily collected over many areas of the network and doesn't scale well. For this reason, peak utilization values should use the 40% rule to estimate actual peak requirements for five minute SNMP collection intervals. If you are using QoS in both the LAN and WAN, then peak utilization may not be necessary. However, busy hour utilization is valuable information that you can use to determine if the combined data and voice traffic will have sufficient bandwidth.

To collect busy hour utilization, you should simply poll all distribution, backbone, and WAN links that will carry voice traffic on an hourly basis and report the busy hour utilization over each link after about a week of data collection. Tools such as Concorde NetHealth can provide this information for a large number of links or you can use any SNMP collector to gather data over a representation of LAN links. In general, you should investigate all WAN links that will be used for VoIP. You will then use this information in conjunction with the expected busy hour voice traffic to determine bandwidth requirements. This process is sufficient for any network that plans of implementing VoIP in a QoS controlled environment. If QoS is not planned for the LAN environment, you may need to inspect peak utilization more closely. In many cases, an organization may also simply choose to over-subscribe the LAN bandwidth to avoid link utilization problems either now or in the future.

Once you collect the data and compare it with voice requirements, you should investigate the desired QoS requirements, expected network growth, and overall requirements to define upgrade requirements. Closely investigate expected combined busy hour LAN link utilization over 50% and WAN link utilization over 75%. This will ensure you meet data and voice requirements.

Step 2 Determine VoIP Overhead

You should determine VoIP overhead by building or by WAN site. However, network planners will not typically be concerned with the low amount of VoIP overhead in the LAN environment. No specific rules exist to determine the traffic requirements; however, your Telecom department may be able to provide busy hour traffic statistics for different areas of the current voice network. You should take special care to differentiate buildings or sites where more or less voice traffic is required. For example, there may be an organization where one building houses engineers that don't typically spend much time on the phone and another building with Support Center personnel who are always on the phone.

We recommend you estimate busy hour call volume and use an Erlang calculator to determine busy hour call requirements. This can then be multiplied by the VoIP encoding method to determine busy hour bandwidth requirements. For example, in a building with 100 people, the Telecom department believes that busy hour traffic for this building is 16.66 hours based on a busy hour call volume of 100 calls and average call duration of 10 minutes. You can then perform an Erlang B calculation based on 16.66 hours and a blocking factor of one percent. The blocking factor is then the confidence level that the estimated bandwidth will be sufficient for the voice requirements. The Erlang B calculator (available on their web site—http://www.erlang.com) computes that 26 lines will be needed to support the 99% confidence level. If we multiply 26 times the encapsulation method, we can determine the busy hour traffic volume.

Assuming the G.711 encoding is used with 80 Kilobyte packets and constant bi-directional voice traffic, we estimate that two megabits/second will be needed to provide adequate bandwidth. Since voice traffic is generally not constant, this is considered an acceptable estimate of the voice traffic requirement.

If the Telecom department cannot provide busy hour traffic volume, you should investigate the voice usage within the building, perhaps even with a visual inspection. This may not be too critical in LAN environments since even constant phone use by 500 users only consumes 40 megabits/second of data bandwidth.

Refer to the *Cisco IP Telephony Network Design Guide* for WAN bandwidth requirements. This document can be found at the following location: http://www.cisco.com/univercd/cc/td/doc/product/voice/ip_tele/network/dgclustr.htm.

Step 3 Determine Bandwidth Requirements

Over WAN links, you should then combine the busy hour data traffic and busy hour voice traffic to determine link bandwidth requirements. Keep in mind that this value does not offer any room for growth. It is up to the organization to determine the growth requirements over time and to perform link trending to determine overall bandwidth requirements once data and voice are implemented. We recommend that you baseline and trend link utilization over time to help ensure that the network consistently meets both data and voice requirements.

Step 4 Determine Design Changes and Required QoS

Design changes should first follow IP Telephony deployment guidelines for LAN and WAN topologies.

Note Refer to the Cisco IP Telephony Network Design Guide for detailed information: http://www.cisco.com/univercd/cc/td/doc/product/voice/ip_tele/network/index.htm.

Once these requirements are met, you should investigate data growth and voice traffic requirements, as well as overall scalability, to define an overall network infrastructure solution for IP Telephony. In addition to overall network design, QoS should be configured for both LAN and WAN environments to validate the network design with a performance baseline.

Step 5 Validate Baseline IP Telephony Performance

You should then validate VoIP performance with a performance baseline prior to implementation to determine voice readiness. Cisco Systems can perform a voice readiness audit that measures delay and jitter across major identified paths. The audit also investigates potential network issues across major network paths. Issues may include queuing delay, CPU utilization, link utilization, buffer utilization, and error rates. You should also consider this step after full IP Telephony implementation in order to baseline the working solution and help determine potential issues with the additional traffic load and performance requirements.

Step 6 Determine Trunking Capacity

In addition to CallManager capacity issues and network capacity issues, you need to determine trunking requirements for voice mail connectivity, PSTN connectivity, and possibly site-to-site connectivity for off-net trunking or off-net overflow. Capacity planning for trunk for voice trunking is already an established capacity planning process in the voice world and the same process applies to trunking capacity for IP Telephony.

The process is normally performed by determining the busy hour traffic statistics for the required trunking application. Some typical trunking applications are:

- Local PSTN trunking
- Long distance PSTN trunking
- Voice mail trunks connected to voice mail systems

- Site-to-site trunks over leased lines or data/voice TDM systems
- Site-to-site trunks for off-net overflow used with Call admission control

You can normally determine busy hour traffic for each application type from the existing telecom switch vendor. This may be an additional service to the organization. Busy hour traffic can be defined as the number of call hours during the busiest hour period. Busy hours may even be the busiest hour of a month. For example, if there are 200 calls with an average duration of ten minutes, then busy hour traffic is 2000 minutes or 33.33 hours. If busy hour traffic is not available, you may simply choose to implement a similar (or slightly higher) number of trunks for IP Telephony than the existing Telecom solution.

Once you determine busy hour traffic for the desired application, you can use the Erlang B calculator (http://www.erlang.com) to arrive at a trunking quantity that satisfies the organization based on a desired blocking factor. Most Telecom organizations currently use a blocking factor of one percent, which means that one percent of the time during the busy hour, users are expected to wait or busy out because of trunk non-availability. This calculation shows that you need 45 trunks for the 2000 minutes. This equates very closely to two T-1s with 23 voice channels each.

You must simply define each application, determine or estimate the busy hour traffic, use the Erlang calculator, and then determine the gateway product or combination of products that will best suit your specific requirements.

Solution Manageability Requirements

Today’s data and voice networks definitely have unique manageability requirements to help ensure consistent service. A management plan for the IP Telephony solution should consider existing requirements for data or voice and unique management requirements needed for the VoIP solution. This solution manageability plan is important in the planning phase because manageability may impact product choice, network design, and/or network management architecture.

When investigating solution manageability requirements, it will help to reference a standard network management model. The most common reference model is the ISO network management model. The model outlines five functional areas in dealing with various aspects of managing a network infrastructure. This model, known as *FCAPS*, includes:

- Fault
- Configuration
- Accounting
- Performance
- Security

The model and its functional areas allow you to clearly define the scope and objectives in the evaluation of network management requirements. Refer to Chapter 6, “Operating the IP Telephony Network” for more information on the FCAPS model and requirements for IP Telephony management.

Functional Areas of Network Management

There are five functional areas of network management:

- Fault Management

  The objective of fault management is to detect problems on network elements within the IP Telephony infrastructure. Problems in hardware, software, or link connectivity can lead to disruption or degradation of network services. You can detect these problems via element functionality, proper element management configuration, forwarding to a management system, and notification of the problem through screen alerts or paging/e-mail. Based on the severity of faults reported by the elements, you can take the proper steps to minimize impact on network availability and performance. You should carefully implement fault management to ensure effectiveness of fault detection, consistent monitoring of all elements, and timely problem resolution.

- Configuration Management

  Configuration management involves managing configuration files, software, addresses, dial plans, network maps, and a detailed inventory of network elements. An up-to-date configuration management system can help reduce the time required in troubleshooting and help to identify potential problems before they cause faults.

- Accounting Management

  In larger organizations and for service providers, it may sometimes be necessary to track the use of application and network resources by user or functional group for budgeting or billing purposes. Accounting information may also be critical to audit or security processes such as fraud detection on voice systems.

- Performance Management

  Performance management involves measuring and allocating resource use in the network to help ensure consistent availability. You can measure performance by exception reporting, baseline reporting, or trending. Measurement may include device resource utilization, link utilization, or end-to-end performance measurements.

- Security Management

  Security management involves various aspects of controlling access to resources in the infrastructure to protect against data theft, data loss, and attacks.

IP Telephony Infrastructure Management

You should determine manageability requirements for a combined data voice infrastructure by investigating existing data requirements, existing Telecom requirements, and evaluating any new requirements. The following table can help you understand major manageability requirements. Once you identify exact requirements for the solution, you can define product and tool requirements from the manageability perspective.

Table 3-20 IP Telephony Infrastructure Management

Management Category	FCAPS Category	Manageability Requirements

Table 3-20 IP Telephony Infrastructure Management

Element Management	Inventory	Maintain memory quantity, chassis/cards installed and versions, IOS versions, IP address allocation
	Configuration	TFTP location for IOS and Catalyst software, backup configuration files, show date/user/change of differences
	Fault	SNMP link/device polling, Syslog reporting, viewing and Syslog summary reporting
	Performance	Interface utilization peaks and 1 hour busy utilization, CPU/memory peak utilization and trends, device backplane utilization
	Security	Device access authentication, authorization, and audit trail
Network Layer Management	Performance	End-to-end performance/delay validation capability
IP Telephony Management		
CallManager Management	Inventory	Maintain chassis type, memory and CPU, software versions
	Configuration	Bulk configuration changes, backup of CallManager configuration and changes
	Fault	SNMP device polling of CallManager status, call status or debug logging capability
	Performance	CPU and memory utilization of CallManager
	Accounting	Call Detail Record (CDR) output for call detail reporting
	Security	Device access authentication, authorization and audit trail
Gateway Management	Fault	SNMP Status of D and B channels on Gateway PRI, SNMP device polling of gateway, Syslog capability for device status messages
	Configuration	Identify hardware modules, hardware versions, software versions
IP Phone Management	Configuration	Quantity, location and firmware versions of IP phones
	Configuration	IP addressing and (DHCP) registration status
	Fault	Obtaining call statistics including termination reason and quality

Staffing and Expertise Requirements

You should begin thinking about staffing and expertise requirements during the planning phase so you can plan for additional resources and/or training requirements. Larger organizations also need to have well-defined provisioning and management processes in place prior to implementation.

Many organizations will simply consider how many people are needed to manage an IP Telephony network. This depends on the size of the organization, the availability requirements of the organization, the automation capabilities, and efficiency of the organization. Instead of defining how many people are needed, this guide recommends that you:

1. Identify your needs.
2. Identify the expertise and staffing level required to fill individual requirements. You can best identify the expertise and staffing by dividing the requirements into the PDIO (plan, design, implement, operate) model.

The following tables can help you plan, design, implement, operate, and help identify the resource requirements for a successful IP Telephony implementation. Additional processes may be required in some environments.

Note The tables contain some examples, but are not intended to be an exact fit for each organization.

Table 3-21 Staffing and Expertise Requirements

IP Telephony Planning				
Process	*Expertise Required*	*Resources Identified*	*Training Required*	*Time Estimated for Process*
Document Existing Environment	LAN/WAN technology, IOS, switches, routers	Data Engineer II	No	2 days
Document Existing Telecom Environment	PBX, trunking, dial plan experience	Voice Engineer II	No	2 days
Document Existing Power/Cabling Infrastructure				
Baseline existing Data Network		Network Management Group		
Determine IP Telephony management requirements	Network management concepts, tools and operations, LAN/WAN technology, IOS, switches, routers	Network Management Group		
Operations support requirements	Management business goals	Organization management		
Determine Availability Requirements	Management business goals	Organization management		
IP Telephony Design				

Table 3-21 Staffing and Expertise Requirements

Process	*Expertise Required*	*Resources Identified*	*Training Required*	*Time Estimated for Process*
VoIP proof of concept or trial	Cisco VoIP configuration, VoIP concepts, CallManager config, signaling concepts			
Infrastructure design	Catalyst Ethernet switches and routers, IOS, QoS techniques, H.323, VoIP concepts and protocols, RTP, TCP/IP, IP subnetting, routing protocols including EIGRP, OSPF, RIP, BGP			
IP Telephony Design	H.323, VoIP concepts and protocols, RTP, NT 4.0 server administration and management, IOS, Codecs including G.711, G.729, Signaling (ISDN-PRI, EIR2, CCS and CAS), FXS, FXO, ground/loop start			
Solution Validation	All technical expertise from infrastructure and IP Telephony design	Lab group		
Network Management Design	H.323, Voice over IP concepts and protocols, RTP, NT 4.0 server administration and management, IOS, tools experience, system administration, database administration.			
IP Telephony Implementation				
Process	*Expertise Required*	*Resources Identified*	*Training Required*	*Time Estimated for Process*

Table 3-21 Staffing and Expertise Requirements

IP Telephony Testing IP Telephony Pilot	Project management, TCP/IP, IP subnetting, DNS, DHCP, TFTP, WEB server, QoS, NT 4.0 server administration and management, IOS, Codecs including G.711, G.729, Cisco VoIP configuration, Catalyst Ethernet switches, IOS, QoS techniques, knowledge and understanding of voice dial plans, exposure to traditional PBXs			
Phone Provisioning Process	Process engineering, DHCP, DNS, IP subnetting, phone training, phone provisioning process documentation			
Migration Plan	TCP/IP, IP subnetting, routing protocols including EIGRP, OSPF, RIP, BGP, project management, troubleshooting, security, network management, server administration			
Facilities Management	Power, environmental, space planning , project management			
Site Survey	Power, environmental, space planning, project management, product physical characteristics			

Table 3-21 Staffing and Expertise Requirements

Implementation Configuration Templates and Implementation Guidelines	Process engineering, project management, TCP/IP, IP subnetting, DNS, DHCP, TFTP, WEB server, QoS, NT 4.0 server administration and management, IOS, Codecs including G.711, G.729, Cisco VoIP configuration, Catalyst Ethernet switches, Cisco IOS, QoS techniques, Knowledge and understanding of voice dial plans, Exposure to traditional PBXs			
Implementation	Cisco VoIP configuration, Catalyst Ethernet switches, IOS			
As-built Documentation	TCP/IP, IP subnetting, project management, troubleshooting experience, documentation tool expertise, VoIP concepts and protocols			
IP Telephony Operations				
Process	*Expertise Required*	*Resources Identified*	*Training Required*	*Time Estimated for Process*
Operations Hand off Requirements	Operations management, project management, process management			
IP Telephony Fault Management Tools and Reporting	SNMP tools, NOC tools, Syslog server, Trap Server, RMON, UNIX or NT administration, database administration, tool training			

Table 3-21 Staffing and Expertise Requirements

IP Telephony Performance Management Tools and Reporting	SNMP tools, NOC tools, RMON, UNIX or NT administration, database administration, tool training,			
IP Telephony Configuration Management Tools and Reporting	SNMP tools, NOC tools, RMON, UNIX or NT administration, database administration, tool training, TACACS+,			
NOC problem identification process	NOC Training			
NOC escalation process	NOC Training			
Security—Fraud Detection	Telecom CDR reporting			
Accounting Management—Billing	Telecom CDR reporting			
Disaster Recovery Plan	Operations management, project management, process management			

Operations Support Plan

Understanding how the network must be supported is critical to the design and operations phase of the network lifecycle. Part of the planning for IP Telephony must then be an understanding of the required service levels within the organization and a plan to support those service levels. This starts with defining the availability and performance requirements for the network. The design group can then build a network that meets these requirements and the operations group can define the processes needed to achieve these availability and performance goals.

Refer to Chapter 6, "Operating the IP Telephony Network" for more information on operations support planning and the process of defining availability and performance standards, defining network service levels, and creating reactive or proactive service definitions.

CHAPTER 4

Designing the IP Telephony Network

In this Chapter

This chapter consists of the following sections:

- Overview
- Introduction to IP Telephony Design
- Designing the Campus Infrastructure
- Designing for LAN/WAN QoS
- Designing Cisco CallManager Clusters
- Selecting Gateways
- Dial Plan Architecture and Configuration
- Designing a Multi-site WAN with Distributed Call Processing
- Designing a Multi-site WAN with Centralized Call Processing
- Catalyst DSP Provisioning
- Cisco Packet Fax and Modem Support Guidelines
- E911 and 911 Emergency Services
- Security Considerations for IP Telephony Networks
- Integrating Voice Mail
- Migrating to an IP Telephony Network

Related Information

- Cisco IP Telephony Network Design Guide

  http://www.cisco.com/univercd/cc/td/doc/product/voice/ip_tele/network/index.htm
- Cisco Auto-negotiation Troubleshooting Document

  http://www.cisco.com/warp/public/473/3.html

- Cisco Policing and Shaping Document

  http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/12cgcr/qos_c/qcpart4/
- Voice Messaging with Cisco uOne 4.1E

  http://www.cisco.com/univercd/cc/td/doc/product/voice/uone/index.htm

Overview

This chapter provides guidance for designing the IP Telephony network and is organized based on the PDIO design model. While most of the sections in this chapter are part of another document (the *Cisco IP Telephony Network Design Guide* which is located at http://www.cisco.com/univercd/cc/td/doc/product/voice/ip_tele/network/index.htm), the information is arranged to follow the PDIO design philosophy. Read these sections in order to gain the most complete understanding of IP Telephony design in relation to PDIO.

Introduction to IP Telephony Design

http://www.cisco.com/univercd/cc/td/doc/product/voice/ip_tele/network/dgintro.htm

Designing the Campus Infrastructure

http://www.cisco.com/univercd/cc/td/doc/product/voice/ip_tele/network/dgcampus.htm

Designing for LAN/WAN QoS

This section provides a blueprint for implementing the end-to-end Quality of Service (QoS) that is required for successful IP Telephony deployment. This document will *not* examine all possible QoS configurations available on all Cisco products. You should also have a base knowledge of Cisco IOS, CatOS, Cisco products, and QoS theories. This section addresses the following deployment solutions:

- Enabling the High Speed Campus
- Building a Branch Office
- Enabling the Wide Area Network

The Importance of QoS

Two factors affect voice quality: lost packets and delayed packets. Packet loss causes voice *clipping* and *skips*. Current Cisco DSP algorithms can correct for up to 30 Msec of lost voice. Cisco VoIP technology uses 20 Msec samples of voice payload per VoIP packet. Therefore, only a single packet can be lost during any given time period for the Codec correction algorithms to be effective. Packet delay can cause either voice quality degradation, due to the end-to-end voice latency, or packet loss, if the delay is variable. If the end-to-end voice latency becomes too long (250 Msec, for example), the conversation begins to sound like two parties talking on a CB radio. If the delay is variable, there is a risk of jitter buffer overruns at the receiving end. Eliminating drops and delay is even more imperative when including fax and modem traffic over IP. By examining the causes of packet loss and delay, you can understand why QoS is needed in all areas of enterprise networks.

Network Quality

Voice packets can be dropped if:

- the network quality is poor.
- the network is congested.
- there is too much variable delay in the network.

Poor network quality can lead to sessions frequently going out of service because of loss of physical or logical connections.

Note Because VoIP design and implementation assumes that the physical and logical network follow sound design methodologies and are extremely stable, this section does not address network quality.

Network Congestion

Network congestion can lead to packet drops and variable packet delays. Voice packet drops from network congestion are usually caused by full transmit buffers on the egress interfaces somewhere in the network. As links or connections approach 100 percent utilization, the queues servicing that connection will fill. When a queue fills, packets attempting to enter the full queue are discarded. This can occur on a campus Ethernet switch as easily as in a service provider’s Frame Relay network.

Because network congestion is typically sporadic, delays from congestion tend to vary. These variable delays are caused by long wait times from the egress interface queue or large serialization delays. See the next section for a discussion of both of these issues.

Delay and Jitter

Delay is the amount of time it takes a packet to reach the receiving endpoint after being transmitted from the sending endpoint. This time period is called the *end-to-end delay* and can be divided into two areas: 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.

You should examine fixed network delay during the initial design of the VoIP network. The International Telecommunication Union (ITU) standard G.114 states that a 150 Msec one-way delay budget is acceptable for high voice quality. The Cisco Technical Marketing team has shown that there is a negligible difference in voice quality scores using networks built with 200 Msec delay budgets.

Examples of fixed network delay include the propagation delay of signals between the sending and receiving endpoints, voice encoding delay, and the voice *packetization* time for various VoIP Codecs. Propagation delay calculations work out to almost 6.3 Usec/km. The G.729A Codec, for example, has a 25 Msec encoding delay value (2 10 Msec frames + 5 Msec look-ahead) and an additional 20 Msec of packetization delay.

Congested egress queues and serialization delays on network interfaces can cause variable packet delays. Without Priority or Low-Latency Queuing (LLQ), queuing delay times equal serialization delay times as link utilization approaches 100 percent.

Serialization delay is a constant function of link speed and packet size. The larger the packet and the slower the link clocking speed, the greater the serialization delay. While this is a known ratio, it can be considered variable because a larger data packet can enter the egress queue before a voice packet at any time, or not at all. If the voice packet must wait for the data packet to serialize, the delay incurred by the voice packet is its own serialization delay plus the serialization delay of the data packet in front of it.

Using Cisco Link Fragmentation and Interleave (LFI) techniques, discussed in the "Enabling the Wide Area Network" section on page 4-47, you can configure serialization delay to be a constant delay value.

Table 4-1 Packet Size and Link Speed

	Packet Size					
Link Speed	**64 Bytes**	**128 Bytes**	**256 Bytes**	**512 Bytes**	**1024 Bytes**	**1500 Bytes**
56 Kbps	9 Msec	18 Msec	36 Msec	72 Msec	144 Msec	214 Msec
64 Kbps	8 Msec	16 Msec	32 Msec	64 Msec	128 Msec	187 Msec
128 Kbps	4 Msec	8 Msec	16 Msec	32 Msec	64 Msec	93 Msec
256 Kbps	2 Msec	4 Msec	8 Msec	16 Msec	32 Msec	46 Msec
512 Kbps	1 Msec	2 Msec	4 Msec	8 Msec	16 Msec	23 Msec
768 Kbps	.640 Msec	1.28 Msec	2.56 Msec	5.12 Msec	10.24 Msec	15 Msec

Because network congestion can occur at any point in time within a network, buffers can fill instantaneously. This can lead to a difference in delay times between packets in the same voice stream. This difference (jitter) is the variation between when a packet is expected to arrive and when it is actually received. To compensate for these delay deltas between voice packets in a conversation, VoIP endpoints use jitter buffers to turn these delay variations into a constant value so voice can be played out smoothly. A jitter buffer is used to temporarily hold the packets prior to voice playout in order to smooth out the variations in packet delay values.

Cisco VoIP endpoints use domain specific part (DSP) algorithms, which have an adaptive jitter buffer between 20-50 Msecs. The actual size of the buffer varies between 20-50 Msecs based on the expected voice packet network delay. These algorithms examine the time stamps in the Real-time Transport Protocol (RTP) header of the voice packets, calculate the expected delay, and adjust the jitter buffer size accordingly. When this adaptive jitter buffer is configured, a 10 Msec portion of *extra* buffer is configured for variable packet delays. For example, if a stream of packets enters the jitter buffer with RTP time stamps indicating 23 Msec of encountered network jitter, the jitter buffer for the receiving VoIP endpoint will be sized at a maximum of 33 Msecs. If a packet's jitter is greater than 10 Msec above the expected 23 Msec delay variation (23+10 = 33msec of dynamically allocated adaptive jitter buffer space), the packet is dropped.

Figure 4-1 Delay and Jitter

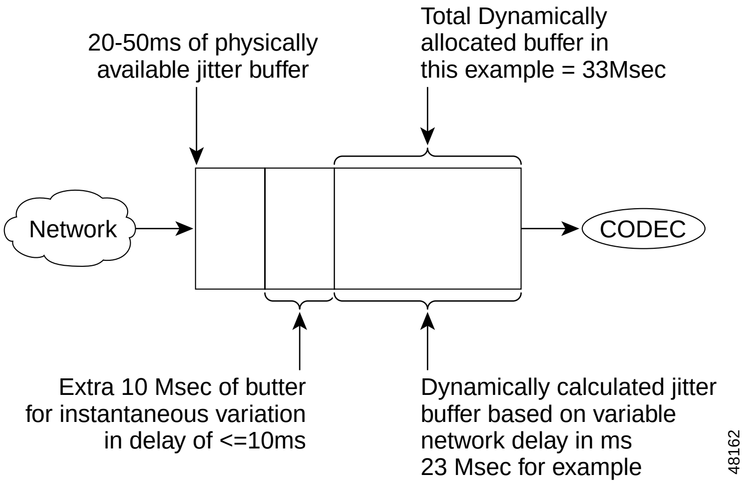
20-50ms of physically available jitter buffer
Total Dynamically allocated buffer in this example = 33Msec
Network
CODEC
Extra 10 Msec of butter for instantaneous variation in delay of <=10ms
Dynamically calculated jitter buffer based on variable network delay in ms 23 Msec for example
48162

Voice quality is only as good as the quality of your weakest network link. Packet loss, delay, and delay variation all contribute to degraded voice quality. Also, because instantaneous buffer congestion can occur at any time in any portion of the network, network quality is an end-to-end design issue. The QoS tools discussed throughout this section are a set of mechanisms to increase voice quality on data networks by decreasing dropped voice packets during times of network congestion and minimizing both the fixed and variable delays encountered in a given voice connection.

These QoS tools can be separated into three categories and are described below:

- Classification
- Queuing
- Network Provisioning

The Cisco QoS tools configured in the examples illustrated in this section are based on Figure 4-2.

Figure 4-2 QoS Network Diagram

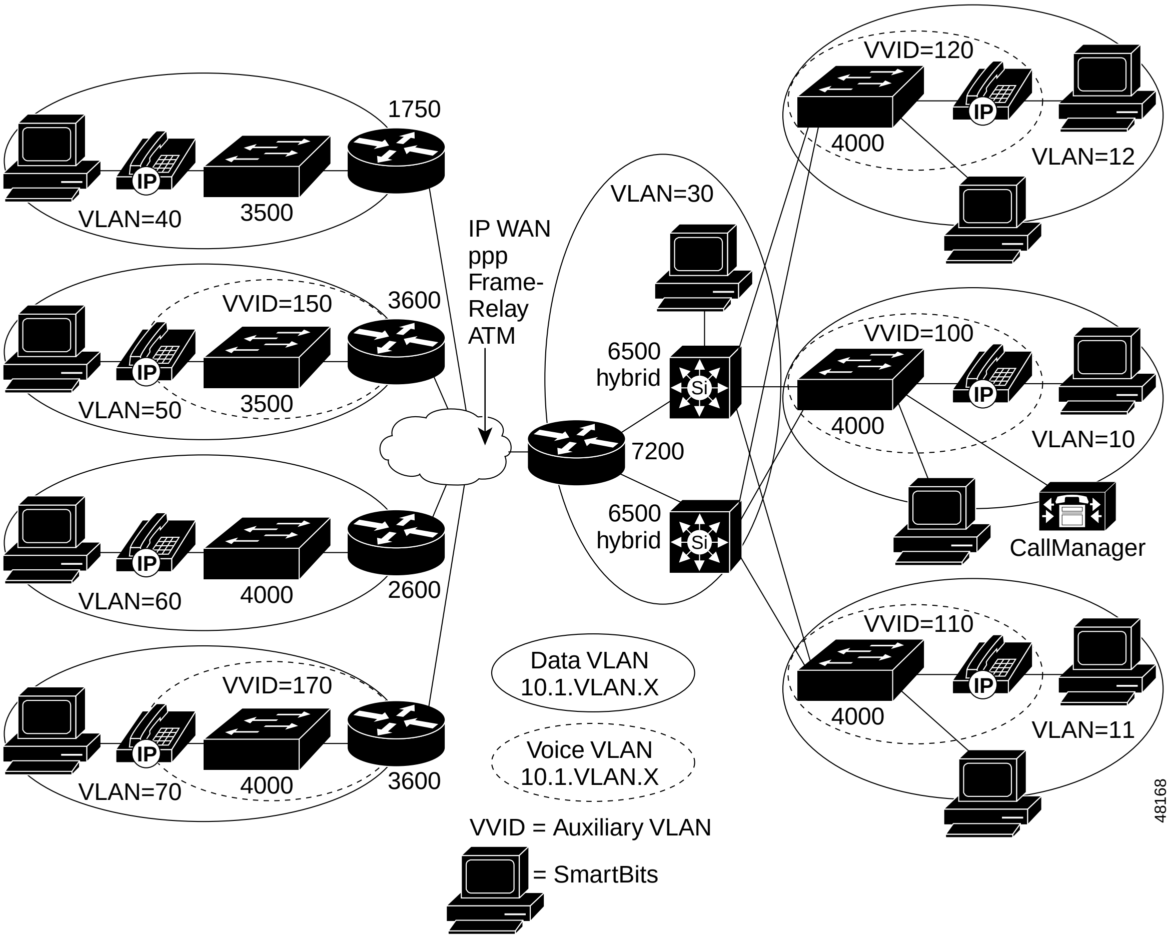

Classification

Classification means marking a packet or flow with a specific priority. This marking establishes a trust boundary that must be enforced.

Classification should take place at the network edge, typically in the wiring closet or within the IP phones or voice endpoints. Packets can be marked as important by using Layer 2 Class of Service (CoS) settings in the User Priority bits of the 802.1p portion of the 802.1Q header (see Figure 4-3) or the IP Precedence/Differentiated Services Code Point (DSCP) bits in the Type of Service (ToS) byte in the IPv4 header (see Table 4-2).

All IP phone RTP packets should be tagged with a values of CoS=5 for the Layer 2 802.1p settings and IP Precedence=5 for Layer 3 settings. Additionally, all Control packets should be tagged with a Layer 2 CoS value of 3 and a Layer 3 ToS of 3.

The previous example uses IP Precedence to mark traffic as a transitional step until all IP devices support the DiffServ Code Point. Ideally, all Cisco VoIP endpoints will use a DSCP value of Explicit Forwarding (EF) for the RTP voice bearer flows and DSCP = Assured Forwarding 31 (AF31) for VoIP Control traffic. See the "Connecting the IP Phone" section on page 4-8 for a detailed discussion of classification.

Figure 4-3 Classification

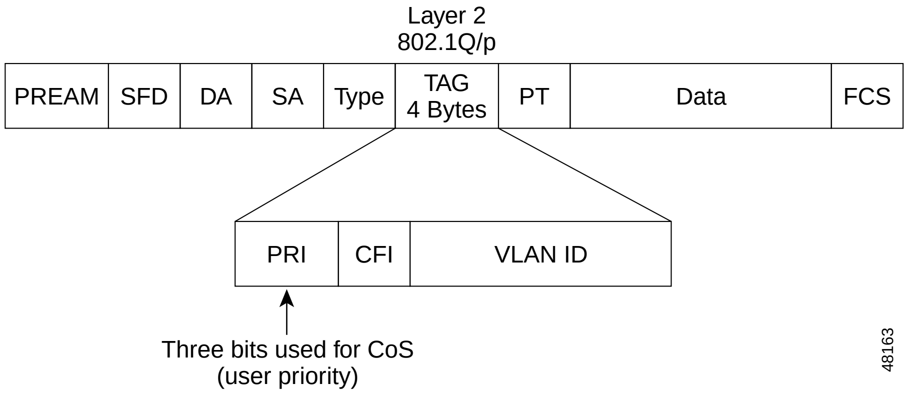

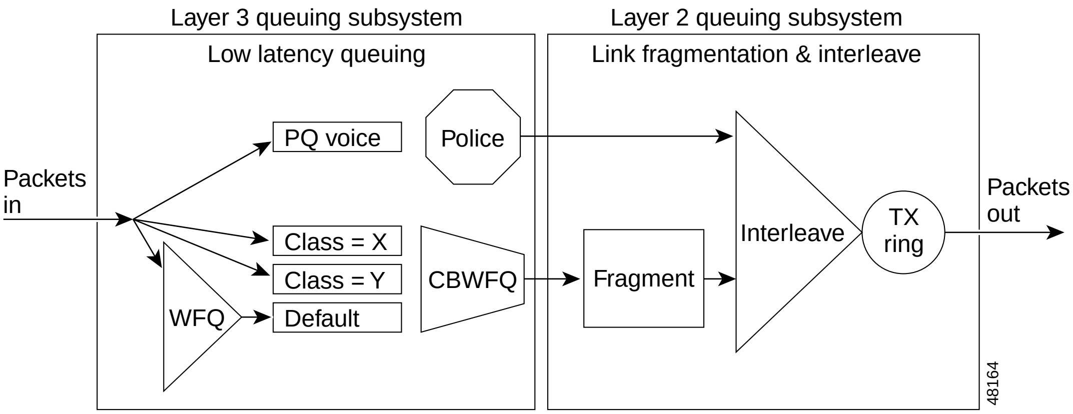

Table 4-2 Classification

Layer 2 Class of Service	IP Precedence	DSCP
CoS 0	Routine (IP precedence 0)	0-7
CoS 1	Priority (IP precedence 1)	8-15
CoS 2	Immediate (IP precedence 2)	16-23
CoS 3	Flash (IP precedence 3)	24-31
CoS 4	Flash-override (IP precedence 4)	32-39
CoS 5	Critical (IP precedence 5)	40-47
CoS 6	Internet (IP precedence 6)	48-55
CoS 7	Network (IP precedence 7)	56-63

Queuing

Queuing means assigning a packet or flow to one of multiple queues, based on classification, for appropriate treatment in the network.

When data, voice, and video are placed in the same queue, packet loss and variable delay are much more likely to occur. By using multiple queues on the egress interfaces and separating voice into a different queue than data, network behavior becomes much more predictable. Queuing is addressed in all sections of this document because buffers can reach capacity in any portion of the network.

Addressing serialization delay is considered part of an overall queuing solution. Because serialization delay is only a factor on slow speed links (links of 768 Kbps or below), it is discussed in the "Enabling the Wide Area Network" section on page 4-47.

Network Provisioning

Network provisioning entails accurately calculating the required bandwidth needed for voice conversations, all data traffic, any video applications, and necessary link management overhead, such as routing protocols.

When calculating the required amount of bandwidth for running voice over the WAN, it is important to remember that all the combined application traffic (voice, video, and data) should only equal 75 percent of the provisioned bandwidth. The remaining 25 percent is used for overflow and administrative overhead, such as routing protocols. VoIP bandwidth calculations, ATM cell overhead, and other details involved in network provisioning are discussed in "Enabling the Wide Area Network" section on page 4-47.

Connecting the IP Phone

Section Highlights and Recommendations

- Use auto-negotiation for port settings on the wiring closet switch.
- Separate IP phones onto a voice-specific subnet.
- Use PortFast to decrease IP phone boot time.
- Extend the classification trust boundary to the phone using **trust-ext** commands.
- Never allow PC applications to send traffic at a CoS or ToS value of 5-7.
- Use only Layer 3/4 intelligent wiring closet switches with SoftPhone.

There are four ways to connect the IP phone to a campus network (see Figure 4-4):

- Use a single cable.
- Use multiple cables.
- Use the SoftPhone application running on a PC.
- Use separate switches.

Each of these connectivity methods has challenges for providing guaranteed voice quality, such as:

- What speed and duplex settings should be used to connect the phone?
- What VLAN and IP addressing scheme should be used?
- How is classification and queuing handled for VoIP flows on the IP phone?

Figure 4-4 IP Phone Connectivity Models

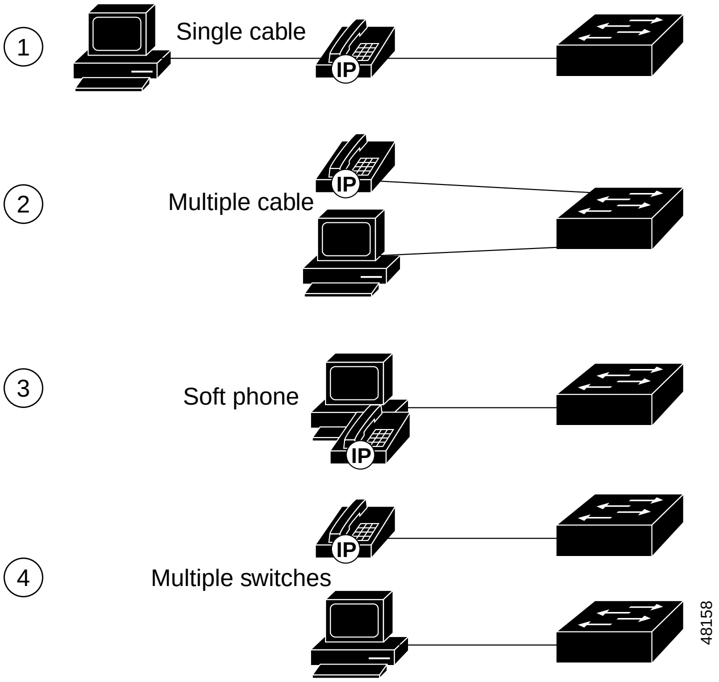

Note You cannot attach IP phones to any shared media devices, such as an Ethernet hub. You also cannot cascade IP phones at this time. Correctly connecting the IP phone is the first step in enabling QoS in the enterprise network.

Single Cable IP Phone Installation

Most enterprises deploy IP Telephony networks using the single-cable IP phone installation model because of the following reasons:

- Ease of installation
- Savings on cabling infrastructure
- Cost savings on wiring closet switch ports

With these cost savings comes requirements for additional switch features, particularly where QoS is concerned. For example, you must correctly configure the Ethernet link speed and duplex, Layer 2 CoS, and queuing on both the IP phone and the wiring closet Ethernet switch. Figure 4-5 illustrates possible QoS problem areas.

Figure 4-5 Possible QoS Problem Areas

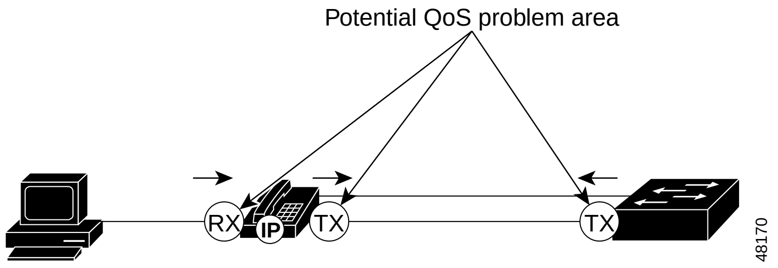

Speed and Duplex Settings

The IP phone's 10/100 Ethernet ports support auto-negotiation for configuring speed and duplex (this is not user-configurable). If the PC's network interface card also uses auto-negotiation, but if the Ethernet switch port is configured for 10BaseT half-duplex, then a link speed mis-match could potentially lead to interface buffer overflow problems. While this half-duplex connection between the IP phone and the switch should not normally be problematic, buffer congestion can arise through the aggregation of 100BaseT full-duplex-to-10BaseT half-duplex aggregation. During periods of intense traffic (such as an extremely high-speed video stream), the half-duplex connection can lead to packet loss from deferred packets. The deferred packets are caused by excessive collisions on the segment. Both the switch and the IP phone, which uses a priority queue for voice, will always send voice traffic out first. However, the high-speed video stream will also be sending as many packets as possible. When either the switch or the phone attempts to send the voice traffic (depending on which direction the video flows are going), it can encounter collisions when attempting to transmit. This results in deferred voice packets.

Note This extreme traffic loss example is rare within a normal enterprise environment, but is reproducible in the lab using SmartBits to simulate MPEG video streams.

This problem is best addressed by setting the switch port to auto-negotiation for all phone connection options. If the port is statically set to 100BaseT full-duplex, the IP phone will automatically set its port to 100BaseT half-duplex, resulting in a duplex mis-match. For an explanation of this problem, please refer to the *Configuring and Troubleshooting Ethernet 10/100Mb Half/Full Duplex Auto-negotiation* technical tip, which can be found at the following location:
http://www.cisco.com/warp/public/473/3.html.

We recommend this IP Telephony configuration because any user can modify the NIC configuration and enable 100BaseT full-duplex. Also, because auto-negotiation enables 100BaseT full-duplex port speeds, implementing it in the IP Telephony rollout ensures the infrastructure is ready to support high-speed video applications.

You can also use the CatOS PortFast mechanism to configure the phone access port to immediately move into a forwarding state. This decreases IP phone boot time. Set the **port host** command on the Catalyst 4000 and 6000 and the **spanning-tree portfast** command on the 2900XL and 3500XL to turn off DTP and PAgP and enable portfast.

Note Phone boot times should normally not be a problem because the phone stays powered and connected at all times.

- Catalyst 4000 and 6000—on the Catalyst 4000, 2948G, 2980G, and 6000 line of Ethernet switches, you can configure this as:

```
cat6k-access> (enable) set port inlinepower 5/1-48 auto -> Default (only avail for
Powered Ethernet linecards)
cat6k-access> (enable) set port speed 5/1-48 auto
cat6k-access> (enable) set port host 5/1-48
```

- Catalyst 3500/2900 XL—on the Catalyst 3500 and 2900 XL switches, you can configure the same functionality as:

```
interface FastEthernet0/1
 power inline auto
 speed auto
    spanning-tree portfast
```

IP Addressing

Once you configure the speed and duplex settings for the IP phone, you need to consider IP addressing. There are three phone IP addressing options:

- Create a new subnet and use that for IP phones in a different IP address space (registered or RFC 1918 address space).
- Provide an IP address in the same subnet as the existing data device (PC or workstation).
- Start a new subnet in the in the existing IP address space (may require you to re-create the entire IP addressing plan for the organization).

You can implement all of these options using either DHCP or static configuration. Adding IP phones can potentially double the organization's need for IP address space. While this may not be an issue in some enterprises, others may not have the available address space in particular subnets or even throughout the enterprise. Because of the IP address space concerns, as well as the requirement of separation between the voice and data networks for administrative and QoS reasons, we recommend you create a new subnet for the IP phones.

Note Using a separate subnet, and potentially a separate IP address space, may not be an option for some small branch offices. See the "Building a Branch Office" section on page 4-42 for information on single address space configurations for connecting both IP phones and data devices.

Figure 4-6 IP Addressing

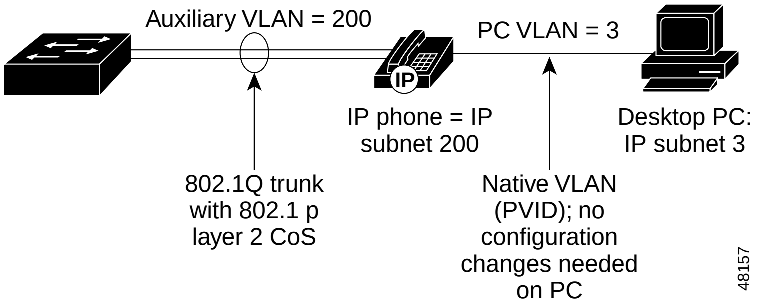
Auxiliary VLAN = 200
PC VLAN = 3
IP
IP phone = IP subnet 200
Desktop PC: IP subnet 3
802.1Q trunk with 802.1 p layer 2 CoS
Native VLAN (PVID); no configuration changes needed on PC
48157

- Catalyst 4000 and 6000—use the **set port auxiliaryvlan** CatOS command to create these IP phone 802.1Q access trunks in the Catalyst 2948Gs, 2980Gs, 4000s, and 6000s:

```
cat6k-access> (enable) set vlan 10 name 10.1.10.0_data
cat6k-access> (enable) set vlan 110 name 10.1.110.0_voice
cat6k-access> (enable) set vlan 10 5/1-48
cat6k-access> (enable) set port auxiliaryvlan 5/1-48 110
```

```
cat4k> (enable) set vlan 11 name 10.1.11.0_data
cat4k> (enable) set vlan 111 name 10.1.111.0_voice
cat4k> (enable) set vlan 11 2/1-48
cat4k> (enable) set port auxiliaryvlan 2/1-48 111
```

- Catalyst 3500/2900 XL—in the Catalyst 3500 and 2900 XL series, you can configure this same functionality using a different set of commands:

```
interface FastEthernet0/1
switchport trunk encapsulation dot1q
 switchport trunk native vlan 12
 switchport mode trunk
 switchport voice vlan 112
 spanning-tree portfast
vlan database
vlan 112
```

Note You can configure the VLAN to match the subnet address for easier troubleshooting.

Classification and Queuing on the Cisco IP Phone

Classifying (marking) traffic as close to the edge of the network as possible has always been an integral part of the Cisco network design architecture. When connecting the IP phone using a single cable model, the phone is now the edge of the managed network. Therefore, the IP phone can and should classify traffic flows.

Three User priority bits in the 802.1p portion of the 802.1Q header are used for signaling Layer 2 CoS information. By default, all VoIP RTP bearer flows from the IP phone are set to a Layer 2 CoS value of 5 and a Layer 3 IP Precedence value of 5. Using IP Precedence is a transitional step as all Cisco VoIP devices will eventually migrate to the DSCP for Layer 3 classification. At that time, Cisco VoIP endpoints using DSCP, instead of IP Precedence, will use a DSCP value of 46, or Expedited Forwarding (EF). These CoS and ToS values are significant when examining how classification and queuing works both within an IP phone and in an enterprise network.

At the heart of a Cisco IP Phone is a 3-port 10/100 switch. One port, P0, is an internal port used for connecting the actual voice electronics in the phone. Port P1 is used to connect a daisy-chained PC and Port P2 is used to uplink to the wiring closet Ethernet switch. Each port has four queues with a single threshold (4Q1T) configuration. One of these queues, Queue 0, is a high priority queue for all bridge protocol data unit (BPDU) and CoS=5 traffic. These queues are all serviced in a round-robin fashion with a timer used on the high priority queue. If this timer expires while the queue scheduler is servicing the other queues, the scheduler will automatically move back to the high priority queue and empty it’s buffer, ensuring voice quality. See Figure 4-7 for an illustration of IP phone queueing.

Figure 4-7 IP Phone Queueing

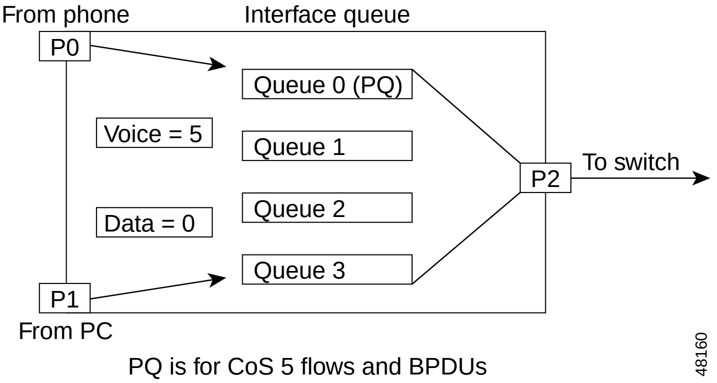

Because the IP phone's high priority queue is accessible to any Layer 2 CoS=5 traffic, it is critical to make sure the PC connected to the IP phone's access port is not classifying traffic. We recommend you extend the Ethernet switch's trust boundary to the IP phone, but not beyond.

Figure 4-8 IP Phone Trust Boundaries

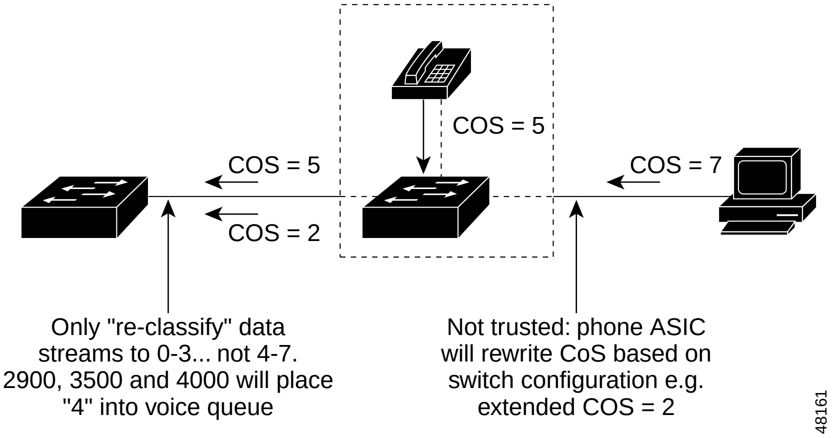
COS = 5
COS = 5
COS = 7
COS = 2
Only "re-classify" data streams to 0-3... not 4-7. 2900, 3500 and 4000 will place "4" into voice queue
Not trusted: phone ASIC will rewrite CoS based on switch configuration e.g. extended COS = 2
48161

- Catalyst 6000—use the **set port trust-ext** command in CatOS 5.5. This command instructs the IP phone to mark all VoIP frames as CoS=5 and all data traffic from the attached PC as CoS=0. Once you configure the phone to manipulate the CoS value, you also need to configure the linecard to accept the IP phone's CoS. The best way to accomplish this is to configure an ACL to trust all CoS classification on Ethernet ports in the Auxiliary VLAN. See the following CLI commands for these switches:

```
cat6k-access> (enable) set port qos 5/1-48 trust-ext untrusted ->Default
```

- Catalyst 2948G, 2980G, and 4000—it's important to note that the Catalyst 2948G, 2980G, and 4000 do not currently offer the **set port qos <mod/port> trust trust-ext** commands. Therefore, these switches must rely on the default configuration of the IP phone, which uses CoS=5 for all VoIP streams and re-classifies CoS on all PC traffic to "0".
- Catalyst 3500/2900 XL—when connecting IP phones to Catalyst 3500s and 2900 XLs using the single cable model, the same functionality is needed. To configure the IP phone not to trust the CoS settings from the PC, use the following commands:

```
interface FastEthernet0/1
 switchport priority extend cos 0
```

Multiple Cable IP Phone Installations

Some enterprises will want to deploy IP Telephony networks using the “multiple cables for IP phone” installation model for the following reasons:

- To deploy IP phones that do not have a second Ethernet port for attaching a PC.
- To create a physical separation between the voice and data networks.
- To easily provide in-line power to IP phones without having to upgrade the data infrastructure.
- To limit the number of switches which need UPS power.
- To limit the amount of CatOS upgrades in the network.
- To limit the Spanning Tree configuration in the wiring closet switches.

Speed and Duplex

Because there is no PC behind the IP phone, port speed and duplex settings aren’t as critical. Although it is a safe configuration to use the identical configuration model as a single cable IP phone installation, this is not required.

IP Addressing

We recommend you use a separate IP subnet and separate VLANs for IP telephony.

Note Using a separate subnet, and possibly a separate IP address space, may not be an option for some small branch offices due to the IP routing configuration. See the “Building a Branch Office” section on page 4-42 for more information on single address space configurations for connecting both IP phones and data devices. If the IP routing can handle an additional subnet at the remote branch, you can use the Cisco Network Registrar and secondary addressing.

Classification and Queuing for the IP Phone

Since the IP phone and any data PCs will be on separate physical cables, queuing on the IP phone is not required. However, since the IP phone is still a managed device, classification should still take place on the phone or ingress access switch port. This classification for VoIP packets can be handled in a variety of methods depending on which hardware is used in the wiring closet switch. See the following scenarios.

- Catalyst 6000—In Figure 4-9, a Catalyst 6000 is used as a wiring closet switch. Ports 3/1-24 connect to IP phones and ports 3/25-48 connect to data-only PCs. Because this is a tightly managed environment, all Layer 2 CoS settings are enforced on the Catalyst 6000.

Figure 4-9 Catalyst 6000 Example with Multiple Cables

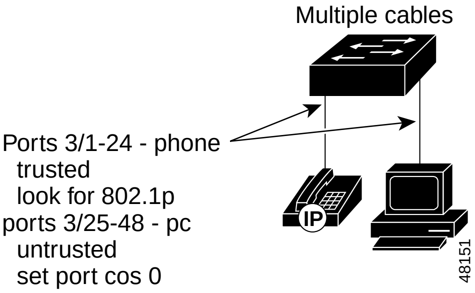

```
cat6k-access> (enable) set port inlinepower 6/1-24 auto
cat6k-access> (enable) set port inlinepower 6/25-48 off
cat6k-access> (enable) set vlan 110 6/1-24
cat6k-access> (enable) set vlan 10 6/25-48
cat6k-access> (enable) set port auxialaryvlan 6/1-24 dot1p
cat6k-access> (enable) set port host 6/1-24
cat6k-access> (enable) set port qos 6/1-24 trust-ext untrusted
```

- Catalyst 4000—currently there is not a “dot1p” extension to the **auxialaryvlan** command on the Catalyst 2948G, 2980G, and 4000 switches. In order to use the IP phone’s 802.1p classification for switch QoS, the Auxiliary VLAN is configured with the same value as the port VLAN ID. This enables the IP phone to mark packets with the proper CoS settings.

```
cat4k> (enable) set vlan 11 2/25-48
cat4k> (enable) set vlan 111 2/1-24
cat4k> (enable) set port host 2/1-48
cat4k> (enable) set port auxiliaryvlan 2/1-24 111
```

- Catalyst 3500/2900 XL—another option when configuring trust is to set this at the port level. On the Catalyst 3500 and 2900 XL series switches, you can use either 802.1p or port-based CoS settings for classifying traffic. A port-based configuration would look similar to the one shown in Figure 4-10:

Figure 4-10 Catalyst 3500/2900 Example with Multiple Cables

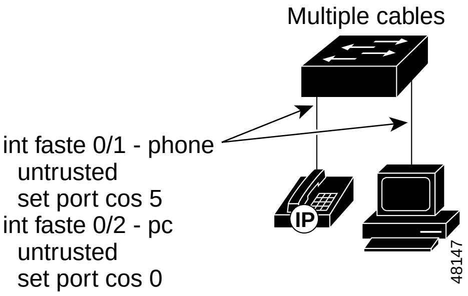

```
interface FastEthernet0/1
    description IP Phone port
 spanning-tree portfast
 switchport mode access
 switchport access vlan 112
interface FastEthernet0/2
    description Data-only PC port
 switchport mode access
 switchport access vlan 12
```

SoftPhone

Some enterprises consider deploying IP telephony using the IP Telephony SoftPhone application. Additionally, there are many more enterprises that want to evaluate the viability of using PC-based VoIP applications. Because the SoftPhone is a Layer 4 application, all VoIP bearer traffic classification, and therefore priority queuing, can only take place at the first Layer 4-capable network device. The Layer 4-enabled wiring closet Ethernet switch models, which also support multiple queues, is limited to the Catalyst 6000 with a PFC installed.

Speed and Duplex

You should set all wiring closet switch access ports to full-duplex 100BaseT. Because PC CPUs are so fast, data has the potential to starve voice on a half-duplex 10BaseT connection. See the "Single Cable IP Phone Installation" section on page 4-9 for details on setting speed and duplex on Catalyst switches.

IP Addressing

IP addressing is not an issue because the SoftPhone application operates on a PC.

Classification and Queuing on the IP Phone

Because the SoftPhone application operates on the PC (sharing its MAC and IP addresses), it is very difficult to establish a classification mechanism to label VoIP packets as priority traffic. The access Ethernet switch cannot trust any Layer 2 CoS markings from the PC because the PC could be marking data packets. Setting port-based CoS will not work for the same reason. In fact, after examining all the options, the only viable way to classify SoftPhone VoIP streams is to filter on Layer 4 UDP port numbers. This requires the access switch to be Layer 3/4 aware because of the need to prioritize voice traffic before the first uplink to the distribution Layer. This limits the choice of wiring closet switches with multiple queues to the Catalyst 6000 with a PFC installed.

- Catalyst 6000 with PFC—in this example, a Catalyst 6000 is being used as a wiring closet switch. The 6000's supervisor engine has a PFC daughter card installed providing Layer 3/4 QoS intelligence. The access port connected to the PC is not trusted, so any CoS or ToS settings on traffic from the PC will be ignored. An ACL (ACL_SOFTPHONE) is configured to filter on VoIP bearer streams and re-classify them to a DSCP value of **EF**.

```
cat6k-access> (enable) set qos enable
cat6k-access> (enable) set port qos 7/1-48 port-based
cat6k-access> (enable) set port qos 7/1-48 trust untrusted
cat6k-access> (enable) set qos acl ip ACL_SOFTPHONE dscp 46 udp any any range 16384
32767
cat6k-access> (enable) commit qos acl ACL_SOFTPHONE
cat6k-access> (enable) set qos acl map ACL_SOFTPHONE 7/1-48
```

Separate Access Layer Switches

Some enterprises are considering deploying IP telephony using completely separate switches in the wiring closet. These customers either choose not to upgrade their current data switches or believe in keeping the voice and data networks completely separate. These installations are very similar to the scenario of using separate ports on the wiring closet switch.

Speed and Duplex

Because there is no PC behind the IP phone, port speed and duplex settings aren’t as critical. Although a safe configuration is to use the identical configuration model as a single cable IP phone installation, this is not required.

IP Addressing

We recommend you use a separate IP address space and separate VLANs for IP telephony. In this case, the entire second switch, the newly installed VoIP-only Ethernet switch, will run a single VLAN. No trunking is necessary between the IP phone and the Ethernet switch. However, we recommend you use 802.1p for tagging VoIP packets from the IP phone as *important*.

```
cat4k> (enable) set port inlinepower 2/1-48 auto
cat4k> (enable) set port inlinepower 2/25-48 off
cat4k> (enable) set vlan 111 2/1-48
cat4k> (enable) set port host 2/1-48
cat4k> (enable) set port auxiliaryvlan 2/1-24 111
```

Classification and Queuing for the IP Phone

Since the IP phone and any data PCs will be on separate physical cables, queuing on the IP phone is not required. However, since the IP phone is still a managed device, classification should still take place on the IP phone. You can handle this classification for VoIP packets in a variety of methods, depending on which hardware is used in the wiring closet switch. If the wiring closet switch is a Layer 2-only device, then the IP phone’s CoS setting will be used for classification at the Access Layer and into the Distribution Layer. See the following example:

- Catalyst 3500/2900 XL

```
interface FastEthernet0/1
switchport trunk encapsulation dot1q
   switchport trunk native vlan 112
   switchport mode trunk
      switchport voice vlan dot1p
 spanning-tree portfast
```

Enabling the High Speed Campus

Section Highlights and Recommendations:

- You must have multiple queues on all interfaces to guarantee voice quality.
- Use UplinkFast in wiring closet switches to enable fast convergence on the Catalyst 2900 XL, 3500, 2948G, 2980G, 4000, and 6000 switches. These switches have multiple egress queues.
- Set all IP Telephony Control and Management traffic to maximum CoS and ToS values of 3.
- Never allow PC applications to send traffic at a CoS or ToS value of 4-7.
- Distribution Layer switches must have the ability to map Layer 3 ToS to Layer 2 CoS values.

Campus Switching Designs for IP Telephony

Until recently, the conventional wisdom was that QoS would never be an issue in the enterprise campus due to the bursty nature of network traffic and the capability of buffer overflow. However, we now know that buffering, not bandwidth, is the primary issue in the campus.Therefore, QoS tools are required to manage these buffers to minimize loss, delay, and delay variation.

Figure 4-11 QoS Problems Areas (1 of 3)

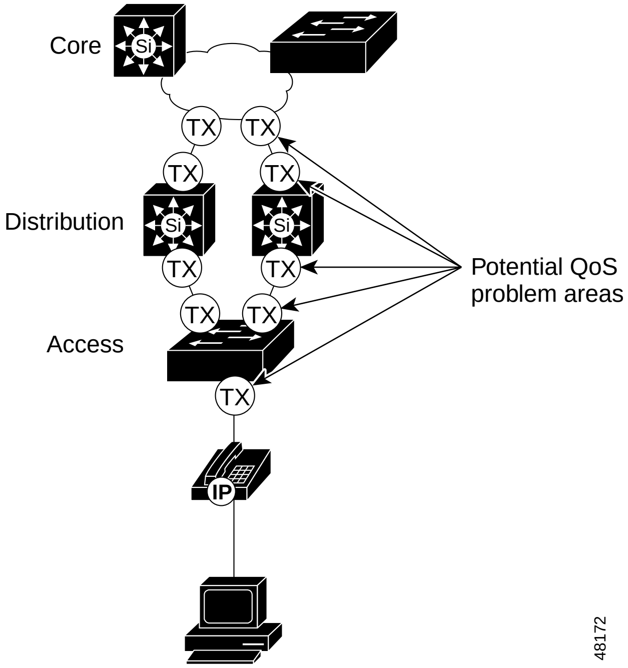

Transmit buffers have a tendency to fill to capacity in high-speed campus networks due to the bursty nature of data networks combining with the high volume of smaller TCP packets. If an output buffer fills, ingress interfaces are not able to place new flow traffic into the output buffer. Once the ingress buffer fills (which can happen very quickly), packet drops will occur. These drops will typically be more than a single packet in any given flow. Current Cisco DSP algorithms can correct for 30 Msec of lost voice. Cisco VoIP technology uses 20 Msec samples of voice payload per VoIP packet, which means that only a single voice RTP packet can be lost during any given time period. If two successive voice packets are lost, voice quality will degrade.

Figure 4-12 QoS Problem Areas (2 of 3)

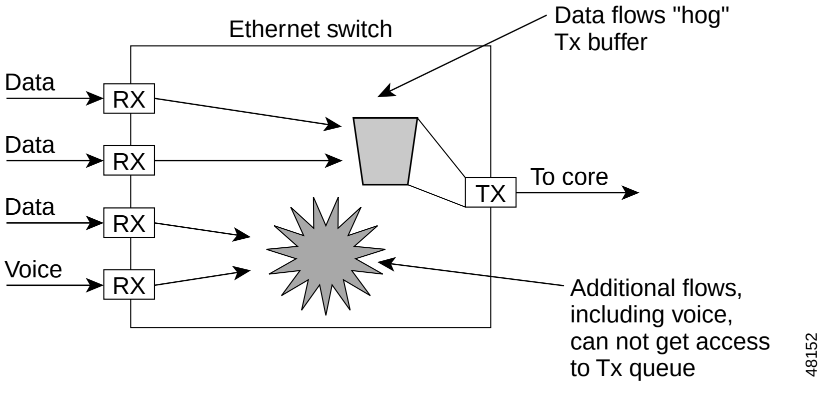

VoIP traffic is sensitive to delay and drop. As long as a campus uses Gigabit Ethernet trunks, which have extremely fast serialization times, delay should never be a factor, regardless of the size of the queue buffer. However, drops will adversely affect voice quality in the campus. Using multiple queues on transmit interfaces is the only way to eliminate the potential of drops caused by buffers operating at 100 percent capacity. By separating voice and video into their own queues, flows are never dropped at the ingress interface if data flows fill up the data transmit buffer.

Figure 4-13 QoS Problem Areas (3 of 3)

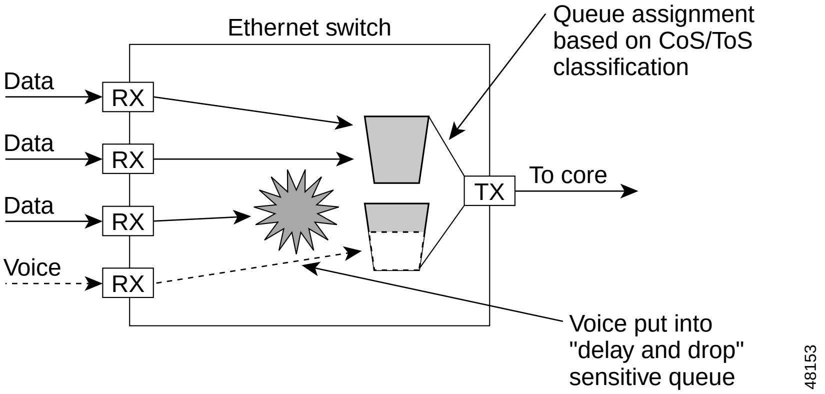
Ethernet switch
Queue assignment based on CoS/ToS classification
Data
RX
Data
RX
Data
RX
Voice
RX
TX
To core
Voice put into "delay and drop" sensitive queue
48153

Note It is critical to remember to verify Flow Control is disabled when enabling QoS (multiple queues) on Catalyst switches. Flow Control will interfere with the configured queuing behavior by acting on the ports before queuing is activated. Flow Control is disabled by default.

The scheduler process can use a variety of methods to service each of these transmit queues. The easiest method is a Round-Robin (RR) algorithm, which services queue 1 through queue N in a sequential manner. While not robust, this is an extremely simple and efficient method that can be used for branch office and wiring closet switches. Distribution Layer switches use a Weighted Round-Robin (WRR) algorithm so higher priority traffic is given a scheduling *weight*. Another option is to combine RR or WRR scheduling with priority scheduling for delay and drop sensitive applications. This uses a priority queue (PQ), which, when there are packets in the queue, is always serviced first. If there are no frames in the PQ, then the additional queues are scheduled using RR or WRR.

An important consideration is how many queues are actually needed on transmit interfaces in the campus:

- Should you add a queue to wiring closet switches for each CoS value?
- Should you add eight queues to the Distribution Layer switches?
- Should you add a queue for each of the 64 DSCP values?

It is important to remember that each port has a finite amount of buffer memory. A single queue will have access to all the memory addresses in the buffer. As soon as you add a second queue, the finite buffer amount is split into two portions, one for each queue. All frames entering the switch not classified for entry into the newly created second queue are now contending for a much smaller portion of buffered memory registers. Therefore, during periods of high traffic, the buffer will fill and frames will be dropped at the ingress interface. Considering that the vast majority of network traffic is TCP-based (comprising TCP ACKs (40 Bytes), TCP SYN/ACKs (44 Bytes) and 512-1024 Byte TCP application traffic (SMTP, HTTP, FTP)), dropping a packet results in a re-send. In other words, dropped packets within TCP-oriented networks increase network congestion. Queuing should be used cautiously and only when particular drop and delay sensitive priority traffic is traversing the network.

Two queues are adequate for wiring closet switches, where buffer management is less critical. Whether these queues are serviced in a RR, WRR, or Priority Queuing manner is less critical because the scheduler process is extremely fast when compared to the aggregate amount of traffic.

Distribution switches require much more complex buffer management because of the flow aggregation occurring at this layer. Not only do you need priority queues, but you also need thresholds within the standard queues.

Cisco chose to use multiple thresholds within queues instead of continually increasing the number of interface queues. Each time a queue is configured and allocated, all of the memory buffers associated with that queue can only be used by frames meeting the queue entrance criteria. For example, we will assume that a Catalyst 4000 10/100 Ethernet port has two queues configured: one for VoIP (VoIP bearer and control traffic) and the default queue which is used for HTTP, e-mail, FTP, logins, NT Shares, and NFS. The 128KB queue is split into a 7:1 transmit and receive ratio. The TX buffer memory is then further separated into high and low priority partitions in a 4:1 ratio. If the default traffic (the web, E-mail, and file shares) begins to congest the default queue, which is only 24KB, then packets begin dropping at the ingress interfaces regardless of whether or not the VoIP control traffic is using any of its queue buffers. The dropped packets of the TCP-oriented applications will cause each of these applications to re-send the data again, aggravating the congested condition of the network. If this same scenario were configured with a single queue, but multiple thresholds used for congestion avoidance, then the default traffic would share the entire buffer space with the VoIP control traffic. Only during periods of congestion, when the entire buffer memory approaches saturation, would the lower priority traffic (HTTP and e-mail) be dropped.

This is not to say that multiple queues are not a vital component in IP Telephony networks. As discussed earlier, the VoIP bearer streams *must* use a separate queue to eliminate the adverse affects of drops and delays to voice quality. However, every single CoS or DSCP value should not get its own queue because the small size of the resulting default queue will cause many TCP re-sends and actually increase congestion.

The VoIP bearer channel is also a poor candidate for queue congestion avoidance algorithms like Weighted Random Early Detection (WRED). Queue thresholding uses the WRED algorithm to manage queue congestion when a preset threshold value is set. WRED works by monitoring buffer congestion and discarding TCP packets if the congestion begins to increase. The result of the drop is that the sending endpoint detects the dropped traffic and slows the TCP sending rate by adjusting the window size. A WRED drop threshold is the percentage of buffer utilization at which traffic with a specified CoS value is dropped, leaving the buffer available for traffic with higher-priority CoS values. The key is the word “Random” in the algorithm name. Even with *Weighting* configured, WRED can still discard packets in any flow; it is just statistically more likely to drop from the lower CoS thresholds.

Marking Control and Management Traffic

In networks with high traffic loads, such as Cisco’s internal network, managing the delivery of control traffic is critical in ensuring that the user experience is positive. For example, with the Delay to Dial-Tone (DTT) time periods, the IP phones use the Skinny Station Protocol to communicate with the CallManager. When an IP phone goes off-hook, it *asks* the CallManager what to do. The CallManager instructs the IP phone to play Dial-Tone. If this Skinny Client Protocol management and control traffic is dropped or delayed within the network, quality is lower. This same logic applies to all signaling traffic for gateways and phones.

To ensure that this control and management traffic is marked as important (but not as important as the actual RTP stream), ACLs are used to classify these streams on Layer 3/4-aware Catalyst 5000 and 6000 switches. Examples of these configurations are detailed below. For designs where a Cisco IOS router is the first Layer 3/4 access point, access lists are used. See the “Enabling the Wide Area Network” section on page 4-47 for examples.

Figure 4-14 Marking Control and Management Traffic

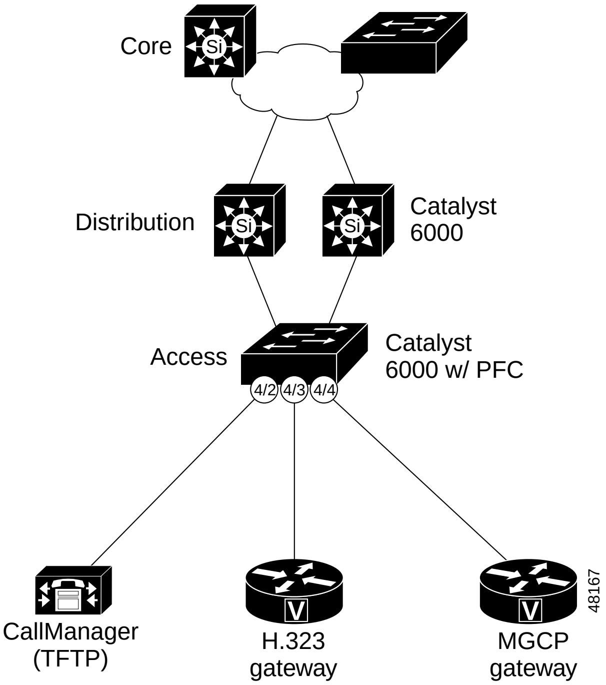

Skinny Protocol

The CallManager communicates with IP phones and gateways using TCP ports 2000-2002. The commands below will classify all skinny traffic from IP phones and gateways (VLAN 110) and the CallManager (4/2) as DSCP 26 (AF31, which is backward-compatible with IP Precedence 3).

Note With the release of *Encore*, the Cisco CallManager will include the ability to configure the CoS and ToS values for all VoIP control and management traffic from the CallManager, IP phones, and Skinny Gateways. When this classification is supported, network element access lists will no longer be required for marking Skinny VoIP control traffic. H.323 and MGCP traffic will still require external, network element marking for the foreseeable future.

```
cat6k-access> (enable) set qos enable
cat6k-access> (enable) set qos acl ip ACL_IP-PHONES dscp 26 tcp any any range 2000 2002
cat6k-access> (enable) set qos acl ip ACL_IP-PHONES trust-cos ip any any
cat6k-access> (enable) set qos acl ip ACL_VOIP_CONTROL dscp 26 tcp any any range 2000 2002
cat6k-access> (enable) set port qos 5/1-48 trust trust-cos
cat6k-access> (enable) set port qos 5/1-48 vlan-based
cat6k-access> (enable) set port qos 5/1-48 trust-ext untrusted
cat6k-access> (enable) set port qos 4/2 port-based
cat6k-access> (enable) commit qos acl all
cat6k-access> (enable) set qos acl map ACL_IP-PHONES 110
cat6k-access> (enable) set qos acl map ACL_VOIP_CONTROL 4/2
```

Each of the commands above performs the following functions:

- Enables switch-wide QoS.
- Creates an access list (ACL_IP-PHONES) marking all Skinny Client/Gateway Protocol traffic from the IP phones and Skinny gateways with a DSCP=26 (AF31) value.

- Adds to the ACL_IP-PHONE access-list trusting all DSCP markings from the IP phone so the ToS=5 RTP traffic is not re-written.
- Creates an access list (ACL_VOIP_CONTROL) marking all Skinny Client/Gateway Protocol traffic from the CallManager with a DSCP=26 (AF31) value.
- Accepts incoming Layer 2 CoS classification (Current 10/100 “1” linecards must have trust-cos enabled even though the parser returns an error).
- Informs the port that all QoS associated with the port will be done on a VLAN basis.
- Instructs the IP phone to re-write CoS from the PC to CoS=0 within the IP phone Ethernet ASIC.
- Informs the CallManager port (4/2) that all QoS associated with the port will be done a port basis.
- Writes the access list to hardware.
- Maps the ACL_IP-PHONE access list to the Auxiliary VLAN.
- Maps the ACL_VOIP_CONTROL access list to the CallManager port.

H.323

The CallManager communicates with H.323 gateways using TCP ports 1720 (H.225) and 11xxx (H.245). The commands below classify H.323 control traffic from the CallManager (4/2) and H.323 gateways (4/3) as DSCP 26 (AF31, which is backwards compatible with IP Precedence 3).

```
cat6k-access> (enable) set qos acl ip ACL_VOIP_CONTROL dscp 26 tcp any any eq 1720
cat6k-access> (enable) set qos acl ip ACL_VOIP_CONTOL dscp 26 tcp any any range 11000
11999
cat6k-access> (enable) set port qos 4/2 port-based
cat6k-access> (enable) set port qos 4/3 port-based
cat6k-access> (enable) commit qos acl ACL_VOIP_CONTROL
cat6k-access> (enable) set qos acl map ACL_VOIP_CONTROL 4/2
cat6k-access> (enable) set qos acl map ACL_VOIP_CONTROL 4/3
```

MGCP

The CallManager communicates with MGCP gateways using UDP port 2427. The commands below will classify MGCP control traffic from the CallManager (4/2) and the MGCP gateway (4/4) as DSCP 26 (AF3,1 which is backwards compatible with IP Precedence 3).

```
cat6k-access> (enable) set qos acl ip ACL_VOIP_CONTROL dscp 26 udp any any eq 2427
cat6k-access> (enable) set port qos 4/2 port-based
cat6k-access> (enable) set port qos 4/4 port-based
cat6k-access> (enable) commit qos acl ACL_VOIP_CONTROL
cat6k-access> (enable) set qos acl map ACL_VOIP_CONTROL 4/2
cat6k-access> (enable) set qos acl map ACL_VOIP_CONTROL 4/4
```

Verify the ACLs are attached to the correct VLANs and ports using the **show** command in the following example:

```
cat6k-access> (enable) sh qos acl map run all

ACL name                         Type Vlans
-------------------------------- ---- ---------------------------------
ACL_IP-PHONES                      IP 110,111,112

ACL name                         Type Ports
-------------------------------- ---- ---------------------------------
ACL_IP-PHONES                      IP
```

```
ACL name                         Type Vlans
-------------------------------- ---- ---------------------------------
ACL_VOIP_CONTROL                   IP

ACL name                         Type Ports
-------------------------------- ---- ---------------------------------
ACL_VOIP_CONTROL                   IP 4/2,4/3,4/4
```

Catalyst 6000 Access Layer

One of the most popular campus configurations for IP Telephony is to use Catalyst 6000 switches in both the wiring closet and distribution/core layers. There are several reasons for this:

- The Catalyst 6000 can provide in-line power to the IP phones.
- The Catalyst 6000 offers the highest growth potential.
- The Catalyst 6000 supports the most advanced Layer 2/3 Campus QoS tools in the Cisco product line.

The Catalyst 6000 example QoS configurations in this section are modeled after the the example in Figure 4-15.

Figure 4-15 Catalyst 6000 QoS Configuration

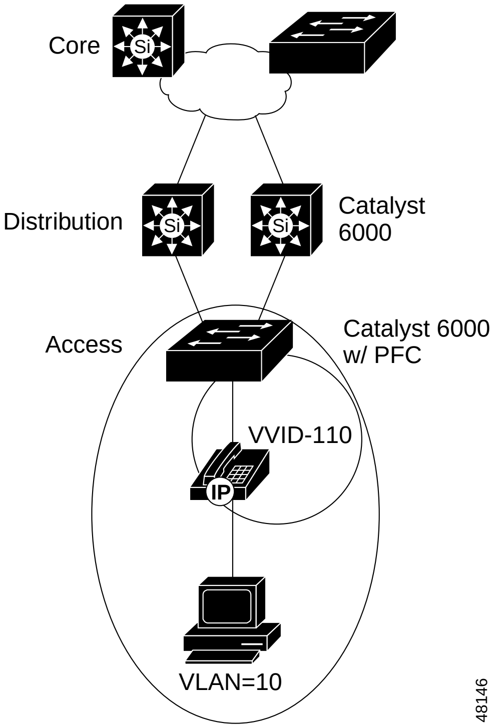

With the addition of the PFC daughter card, the Catalyst 6000 is inherently a Layer 2, 3, and 4 QoS-aware platform. The PFC can be used to enable advanced QoS tools such as packet classification and marking, scheduling, and congestion avoidance based on both either Layer 2 or Layer 3/4 header information. Multiple receive and transmit queues, with thresholds, can be configured and utilized based on the QoS policy rules configured in the switch.

The Catalyst 6000 has two versions of Supervisor Engines: the Sup1 and Sup1A. There are also two versions of 6000 linecards, also denoted by the "A" product numbers. All Catalyst 6000 Ethernet modules support one single receive queue with four thresholds and two transmit queues, each with two thresholds. The "A" cards include enhanced QoS features, specifically an additional priority queue for both ingress and egress interfaces. These queues are serviced in a WRR method, except the priority queue, which is always serviced as soon as frames have entered the queue. To see how a port is configured, issue the **show port capabilities <*mod/port*>** CatOS command. The default QoS capabilities of the port can be changed using the **set qos map <*port_type*> rx | tx <*queue#*> <*threshold#*>** and **set qos wred-threshold** commands. When modifying the queue thresholds, it is important to remember that the higher priority queue has a higher numerical value.

Scheduling for the Catalyst 6000 transmit interfaces is managed by the WRR algorithm. Each queue is given a weight, which is user-configurable. By default, the *high* queue is given 98 percent of the scheduler time and the *low* queue is given just two percent. This ratio ensures that packets with a low delay tolerance are not delayed in a queue. This is also the reason for giving the *low* queue a much higher percentage of the overall interface buffer. If the priority Queue is configured, it will always be serviced first. If no frames reside in the PQ, WRR begins to schedule the other two queues.

Catalyst 6000 Port Scheduling and Queuing Schemes

Receive Interface

1Q4T:

- One standard queue with 4 drop thresholds
- 8KB receive buffer for 10/100 Mbps
- 64KB receive buffer for 1000 Mbps
- All 10/100/1000 Mbps modules

Table 4-3 Default Values for Drop Thresholds

Percentage of Buffer Capacity	Drop CoS Value
50	0-1
60	2-3
80	4-5
100	6-7

1P1Q4T:

- One Priority queue, one standard queue with four drop thresholds
- Only certain versions of 10/100/1000 Mbps modules; line card-dependant
- By default, all CoS 5 frames are placed in the Priority Queue, which uses a strict priority scheduling algorithm

Table 4-4 Default Values for Drop Thresholds in Standard Queue

Queue Number	Percentage of Buffer Capacity	Drop CoS Value
1	50	0-1
1	60	2-3

Table 4-4 Default Values for Drop Thresholds in Standard Queue (continued)

Queue Number	Percentage of Buffer Capacity	Drop CoS Value
1	80	4
1	100	6-7
2	100	5

Transmit Interface

2Q2T:

- Two standard queues with two drop thresholds.
- The high priority queue is allocated 20 percent of the total queue size. The low priority queue is allocated 80 percent of the total queue size.
- All 10/100/1000 Mbps modules

Table 4-5 Default Values for Drop Thresholds

Queue Number	Percentage of Buffer Capacity	Drop CoS Value
1 - Low priority - 80% of total queue size	40	0-1
	100	2-3
2 - High priority - 20% of total queue size	40	4-5
	100	6-7

1P2Q2T:

- One PQ, two standard queues with two drop thresholds. By default, all CoS 5 frames are placed in the PQ, which uses a strict priority scheduling algorithm.The PQ is always serviced first and then, once the PQ is empty, WRR is used on the remaining queues. The PQ is allocated 15 percent of the total queue size, as well as the high priority queue. The low priority queue is allocated 70 percent of the total queue size.
- Only certain versions 10/100/1000 Mbps modules; linecard dependant.

Table 4-6 Default Values for Drop Thresholds

Queue Number	Percentage of Buffer Capacity	Drop CoS Value
1 - Low priority - 80% of total queue size	40	0-1
	100	2-3
2 - High priority - 20% of total queue size	40	4-5
	100	6-7
3 - Priority queue - 15% of total queue size	100	5

Configuring QoS Parameters

After the IP phone has been connected to the wiring closet switch (See the "Connecting the IP Phone" section on page 4-8), you must configure the QoS parameters on the switch. This includes the following steps:

1. Set up multiple queues on all ports.

2. Configure access to the queues.
3. Set thresholds for traffic drops.
4. Configure the Uplink Interface to the Distribution Switch.
5. Configure MLS and Catalyst QoS.
6. Configure the Catalyst 6000 Transmit Queue.
7. Configure the Catalyst 6000 CoS/ToS to DSCP Mappings.
8. Verify CoS/ToS-to-DSCP Mapping.

The following sections describe these in detail.

Step 1 Set up IP Phone Port Queues

Following the single cable IP phone installation scenario, this access port is configured to trust the IP phone and not the attached PC. It also uses multiple transmit queues, one being a priority queue for voice traffic.

Figure 4-16 IP Phone Port Queueing

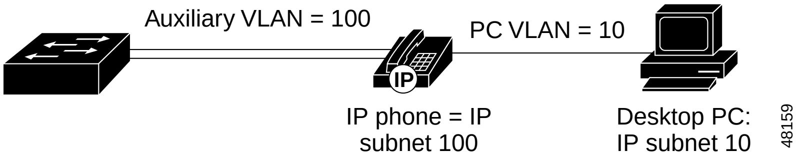

The commands listed below enable QoS on the Access Layer Catalyst 6000:

```
cat6k-access> (enable) set qos enable
cat6k-access> (enable) set port qos 5/1-48 vlan-based
cat6k-access> (enable) set port qos 5/1-48 trust-ext untrusted
cat6k-access> (enable) set port qos 5/1-48 trust trust-cos
cat6k-access> (enable) set qos acl ip ACL_IP-PHONES trust-cos any
cat6k-access> (enable) commit qos acl ACL_IP-PHONES
cat6k-access> (enable) set qos acl map ACL_IP-PHONES 110
```

Each command performs the following functions:

- Enable switch-wide QoS.
- Inform the port that all QoS associated with the port will be done a VLAN basis.
- Instruct the IP phone to re-write CoS from the PC to CoS=0 within the IP phone Ethernet ASIC.
- Accept incoming Layer 2 CoS classification (Current 10/100 “1” linecards must still have trust-cos enabled, even though the parser returns an error).
- Create an access list which accepts incoming Layer 3 ToS classification (only necessary on 10/100 ports).
- Write the access list to hardware.
- Map the access list to the Auxiliary VLAN.

Step 2 Configure Access to the Queues

Once you enable QoS on the Access Layer Catalyst 6000, use the command below to place all CoS=3 (VoIP control traffic) into the second transmit queue with a low drop threshold to ensure successful call control during periods of heavy congestion. All CoS=5 (VoIP RTP Bearer traffic) is automatically placed into the second queue.

```
cat6k-access> (enable) set qos map 2q2t tx 2 1 cos 3
```

Step 3 Set the Thresholds for Traffic Drops

One of the fundamental pieces of implementing QoS is verifying that the configurations are actually performing as expected. On the Catalyst 6000 Access Layer switch, verify its performance under periods of high congestion by examining the output of the following commands:

```
show port qos <mod/port>
show qos info runtime <mod/port>
show mac <mod/port>
show qos statistics l3
show qos stat <mod/port>
```

See the following sample out from the commands shown below:

```
cat6k-access> (enable) sh port qos 5/1
QoS is enabled for the switch
QoS policy source for the switch set to local.

Port  Interface Type Interface Type Policy Source Policy Source
      config         runtime        config        runtime
----- -------------- -------------- ------------- -------------
 5/1      vlan-based     vlan-based          COPS         local

Port  TxPort Type  RxPort Type  Trust Type   Trust Type    Def CoS Def CoS
                                config       runtime       config  runtime
----- ------------ ------------ ------------ ------------- ------- -------
 5/1          2q2t         1q4t    trust-cos    trust-cos*       0       0

Port  Ext-Trust Ext-Cos
----- --------- -------
 5/1  untrusted       0

(*)Runtime trust type set to untrusted.

Config:
Port  ACL name                         Type
----- -------------------------------- ----
No ACL is mapped to port 5/1.
ACL is mapped to VLAN

Runtime:
Port  ACL name                         Type
----- -------------------------------- ----
No ACL is mapped to port 5/1.


cat6k-access>(enable) sh qos info run 5/1
Run time setting of QoS:
QoS is enabled
Policy Source of port 5/1: Local
Current 10/100 "1" linecards support 2q2t/1q4t only
Tx port type of port 5/1 : 2q2t
Rx port type of port 5/1 : 1q4t
Interface type: vlan-based
ACL is mapped to VLAN
ACL attached:
The qos trust type is set to trust-cos.
Warning: Runtime trust type set to untrusted.
```

```
Default CoS = 0
Queue and Threshold Mapping for 2q2t (tx):
Queue Threshold CoS
----- --------- ---------------
1     1         0 1
1     2         2
2     1         3 4 5
2     2         6 7
Queue and Threshold Mapping for 1q4t (rx):
Queue Threshold CoS
----- --------- ---------------
1     1         0 1
1     2         2
1     3         3 4 5
1     4         6 7
<snip…>


cat6k-access> (enable) sh mac 5/1

Port     Rcv-Unicast          Rcv-Multicast        Rcv-Broadcast
-------- -------------------- -------------------- --------------------
 5/1                    267223                   37                    4

Port     Xmit-Unicast         Xmit-Multicast       Xmit-Broadcast
-------- -------------------- -------------------- --------------------
 5/1                  28748894                 5206                   72

Port     Rcv-Octet            Xmit-Octet
-------- -------------------- --------------------
 5/1                  17178128           1840430081

"Out-Discards" are packets drooped due to congestion in the tx interface buffers
MAC      Dely-Exced MTU-Exced  In-Discard Out-Discard
-------- ---------- ---------- ---------- -----------
 5/1              0          0          0      262140


cat6k-access> (enable) sh qos stat l3
VoIP Control packets that have been re-written with CoS=3/DSCP=26 (AF31)
Packets dropped due to policing:                0
IP packets with ToS changed:                 1885
IP packets with CoS changed:                  781
Non-IP packets with CoS changed:                0


cat6k-access> (enable) sh qos stat 5/1
All packets dropped are in the 1st drop threshold of queue #1
Tx port type of port 5/1 : 2q2t
Q #  Threshold #:Packets dropped
---  -----------------------------------------------
1    1:393210 pkts, 2:0 pkts
2    1:0 pkts, 2:0 pkts
Rx port type of port 5/1 : 1q4t
Q #  Threshold #:Packets dropped
---  -----------------------------------------------
1    1:0 pkts, 2:0 pkts, 3:0 pkts, 4:0 pkts
```

Step 4 Configure the Uplink Interface to the Distribution Switch

Once you configure all of the access port queuing, you need to configure the uplink interfaces to the distribution/core switch.

Step 5 Configure MLS and Catalyst QoS

If the IP phones are in a different VLAN than the CallManager, then you need additional configurations. Any time a packet of flow is sent to the MSFC for Layer 3 switching, the CoS will be set to “0”. Because most configurations will have the MSFC located in the Distribution Layer switch, the Access Layer switch must trust all DSCP tagging on the uplink trunk from the Distribution Layer. This enables the DSCP marking to be retained and used for DSCP-to-CoS Layer 3 classification in the wiring closet switch. Use **trust-cos** for Layer 2 uplinks and **trust-dscp** for Layer 3 uplinks. See the following sample command:

```
cat6k-access> (enable) set port qos 1/1 trust trust-dscp
```

Step 6 Configure the Catalyst 6000 Transmit Queue

All VoIP (CoS=5) traffic will be placed into the egress interface Priority Queue on 1p2q2t interfaces and Queue #2 on 2q2t interfaces as soon as QoS is enabled. You must also configure the Catalyst 6000 CoS queue admission rules to ensure CoS=3 traffic flows (VoIP control traffic) are placed into the second queue.

```
cat6k-access> (enable) set qos map 1p2q2t tx 2 1 cos 3
cat6k-access> (enable) set qos map 2q2t tx 2 1 cos 3
```

Step 7 Configure the Catalyst 6000 CoS/ToS to DSCP Mappings

Cisco follows the IETF recommendations for setting the DSCP classification values for both the VoIP Control plane traffic and VoIP Bearer or Media plane traffic. The recommended settings are:

- DSCP=AF31 for VoIP Control plane.
- DSCP=EF for VoIP Bearer plane.

To correctly map the Layer 2 CoS and Layer 3 IP Precedence settings to these DSCP values, you must modify the default CoS/ToS to DSCP mappings as in the following example:

```
cat6k-distrib> (enable) set qos cos-dscp-map 0 8 16 26 32 46 48 56
cat6k-distrib> (enable) set qos ipprec-dscp-map 0 8 16 26 32 46 48 56
```

Step 8 Verify CoS/ToS-to-DSCP Mapping

Use the cards in the example below to verify CoS/ToS-to-DSCP mapping:

```
cat6k-distrib> (enable) sh qos map run cos-dscp-map
CoS - DSCP map:
CoS   DSCP
---   ----
  0   0
  1   8
  2   16
  3   26 -> 26 = AF31
  4   32
  5   46 -> 46 = EF
  6   48
  7   56


cat6k-distrib> (enable) sh qos map run ipprec-dscp-map
IP-Precedence - DSCP map:
IP-Prec   DSCP
-------   ----
      0   0
      1   8
      2   16
      3   26 -> 26 = AF31
      4   32
      5   46 -> 46 = EF
      6   48
```

7 56

Catalyst 4000 Access Layer

Another popular campus configuration for IP Telephony is to use Catalyst 2948G, 2980G, and 4000 series switches in the wiring closets. There are several reasons for this:

- The Catalyst 4006 can provide in-line power to the IP phones.
- The Catalyst 4000 offers a very low price per port.
- Extremely scalable, high-speed switching.

With the release of CatOS 5.2, the Catalyst 4000 lines support dual-transmit queues on every interface. Admission to the queues are based on Layer 2 CoS markings and is configurable in 802.1p User Priority pairs.

Catalyst 4000 Port Scheduling and Queuing Schemes

Receive Interface

FIFO—One standard FIFO (First-In, First-out) queue

Transmit Interface

2Q1T—Two standard queues with a single threshold. Scheduling is done on a RR basis. Admission to the queues are based on 802.1p CoS value and are user-configurable in pairs. If you enable QoS but do not modify the CoS-to-transmit queue mappings, switch performance could be affected because all traffic is assigned to queue 1. Once QoS is enabled on the Catalyst, CoS mappings must be changed to use the newly created queue.

Table 4-7 Default 4000 Queue Admission Criteria

Queue Number	Queue Admission CoS Value
1	0-7
2	Broadcast, Multicast, and Unknown Traffic

The Catalyst 4000 QoS configuration samples in this document are modeled after the example below:

Figure 4-17 Catalyst 4000 QoS Configuration Example

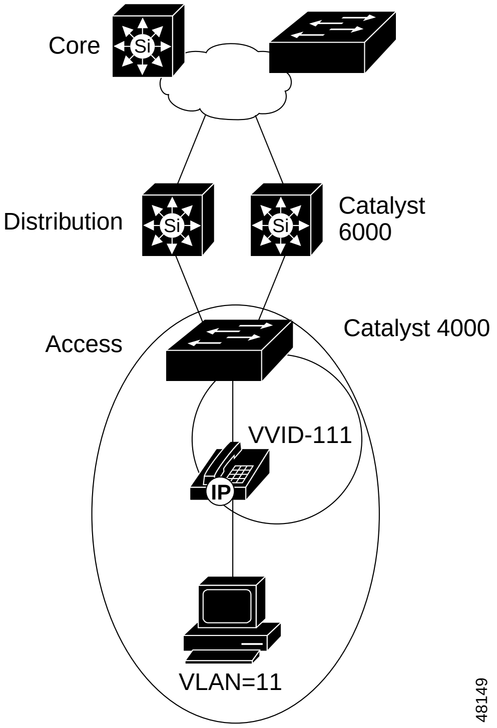

After the IP phone has been connected to the wiring closet switch (See "Connecting the IP Phone"), you must configure the QoS parameters on the switch. This includes the following steps:

1. Establish Switch-wide QoS.
2. Verify Catalyst 4000 Queue Admission Configuration.
3. Set up Phone Port Queuing.
4. Configure the Uplink Interface to the Distribution Switch.

The following sections describe these steps in detail.

Step 1 Establish Catalyst 4000 Switch-wide QoS

By default, only one queue is enabled on the Catalyst 4000 line of switches. Use the **set qos map** commands to enable the use of the second queue in CatOS 5.5.1. VoIP Control (CoS=3) frames should be placed into the second queue in the Catalyst 4000. These maps must be configured in pairs of CoS values as the Catalyst 4000 only examines the first two CoS bits. See the following example:

```
cat4k> (enable) set qos enable
cat4k> (enable) set qos map 2q1t 1 1 cos 0-1
cat4k> (enable) set qos map 2q1t 2 1 cos 2-3
cat4k> (enable) set qos map 2q1t 2 1 cos 4-5
cat4k> (enable) set qos map 2q1t 2 1 cos 6-7
```

Step 2 Verify Catalyst 4000 Queue Admission Configuration

Use the following command to verify the Catalyst 4000 queue administration configuration:

```
cat4k> (enable) show qos info runtime
Run time setting of QoS:
QoS is enabled
All ports have 2 transmit queues with 1 drop thresholds (2q1t).
```

```
Default CoS = 0
Queue and Threshold Mapping:
Queue Threshold CoS
----- --------- ---------------
1     1         0 1
2     1         2 3 4 5 6 7
```

Step 3 Set up Phone Port Queuing

In version 5.5.1, the Catalyst 4000 line does not offer any advanced IP phone queuing features. Therefore, the Catalyst 4000 depends on the default CoS marking and enforcement on the IP phone. See "Connecting the IP Phone" for more details.

Step 4 Configure the Uplink Interface to the Distribution Switch

You do not need to configure any special queuing or scheduling commands on the Catalyst 4000 side of the link from the Access Layer Catalyst 4000 to the Distribution Layer Catalyst 6000. Queuing is on once QoS has been enabled and classification and queue admission is already done.

When using the Catalyst 4000, add the Layer 3 engine, WS-X4232. This engine enables IP, IPX, and multicast routing for the switch so additional uplink configuration can be done. The Layer 3 engine enables the 4000 to support four transmit queues based on IP Precedence for entrance criteria on the two gigabit uplinks. The four queues are scheduled using a user-configurable WRR algorithm.

Transmit Interface

4Q1T—Two standard queues with a single threshold. Scheduling is done on a RR basis. Admission to the queues is based on 802.1p CoS value and is user-configurable in pairs. Once you enable QoS on the Catalyst, you must change the CoS mappings to use the newly created queue.

Note The Layer 3 queue numbering is the reverse of the Layer 2 numbering.

Table 4-8 Default 4000 Layer 3 1000Mbps Uplink Queue Admission Criteria

Queue Number	Queue Admission IP Precedence Value
1	6-7
2	4-5
3	2-3
4	0-1

Catalyst 3500 Access Layer

The IP Telephony features in the Catalyst 2900 and Catalyst 3500 series, with a minimum Cisco IOS of 12.0(5)XU, allow interaction with the IP phone for extending the CoS marking rules. The Catalyst 2900 XL and 3500 XL switches are also able to classify un-tagged packets at the ingress ports by setting a default CoS priority for each port. However, these switches (except for the 3548 XL) cannot re-classify any tagged packets and will only honor the 802.1p priority and place the packets in the appropriate transmit queue. All Catalyst 3500 switches and all Catalyst 2900XLs with 8 MB DRAM support these QoS features. The Catalyst 2900XL with 4 MB DRAM does not support QoS features.

Catalyst 3500 Port Scheduling and Queuing Schemes

Receive Interface

1Q-FIFO—One standard FIFO (First-In, First-out) queue.

Transmit Interface 10/100 Ports

2Q1T—Two standard queues with a single drop threshold. Scheduling is done on a priority-scheduling basis. Admission to the queues is based on 802.1p CoS or port priority CoS value and is not user-configurable.

Table 4-9 Catalyst 3500 Queue Admission Criteria

Queue Number	Queue Admission IP Precedence Value
1	0-3
2	4-7

Transmit Interface Gigabit Ethernet Ports

8Q-FIFO—Eight standard queues with a single drop threshold. Currently, only two queues are used. Scheduling is done on a priority-scheduling basis. Admission to the queues is based on 802.1p or port priority CoS values and is not user-configurable.

The Catalyst 3500 Gigabit Ethernet Queue Admission Criteria are:

Table 4-10 Catalyst 3500 Gigabit Ethernet Queue Admission Criteria

Queue Number	Queue Admission IP Precedence Value
1	0-3
2	4-7
3-8	Not used

The Catalyst 3500 QoS configurations are modeled after the example below:

Figure 4-18 Catalyst 3500 Gigabit Ethernet QoS Configuration Example

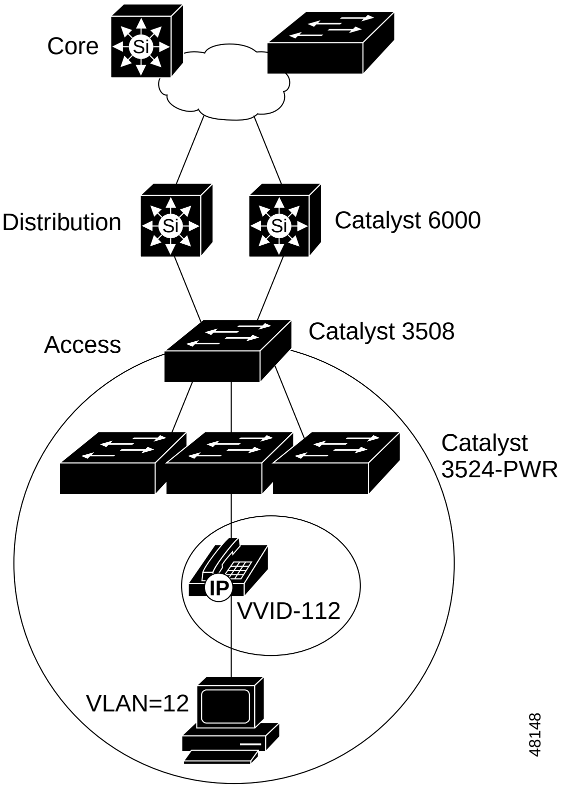

After the IP phone has been connected to the wiring closet switch (See "Connecting the IP Phone"), you must configure the QoS parameters on the switch. This includes the following steps:

1. Set up IP Phone Port Queueing.
2. Configure the Uplink Interface to the Distribution Switch.

Step 1 Set up IP Phone Port Queuing

Following the single cable IP phone installation scenario, this access port is configured to trust the IP phone and not the attached PC. Use the following commands:

```
interface FastEthernet0/1
 switchport trunk encapsulation dot1q
 switchport trunk native vlan 12
 switchport mode trunk
 switchport voice vlan 112
 switchport priority extend cos 0
 spanning-tree portfast
```

Step 2 Configure the Uplink Interface to the Distribution Switch

The recommended design for wiring closet configurations of Catalyst 3500XL series switches is a star topology of 3524 PWR XLs connected to a 3508 that has dual uplinks to the Distribution Layer Catalyst 6000s. These uplinks are Gigabit Ethernet links that are load balancing VLANs across the uplinks and configured with UplinkFast for fast Layer 2 convergence. A Catalyst 3500 series GigaStack configuration cannot provide guaranteed voice QoS because is essentially a shared media access model. Configure the following commands:

```
interface GigabitEthernet0/1
```

```
switchport trunk encapsulation dot1q
switchport mode trunk
```

Catalyst 6000 Distribution Layer

After the access switch is configured and attached to the Distribution Layer, QoS must be set up on the distribution switches. This requires a few simple changes to the configuration:

1. Configure VoIP control traffic transmit queuing.
2. Configure the Distribution Layer with a Layer 3 Access switch:
   a.Trust ToS and DSCP from the Access Layer.
   b.Configure ToS to DSCP mappings.
3. Configure the Distribution Layer with a Layer 2 Access switch:
   a.Trust CoS and DSCP from the Access Layer.
   b.Configure CoS to DSCP mappings.
   c.Configure Layer 3 access lists for VoIP control traffic classification.
4. Configure the connection to the 7200 WAN router.

All Catalyst 6000 Distribution Layer configuration samples in this document are based on the network below:

Figure 4-19 Catalyst 6000 Distribution Layer Configuration

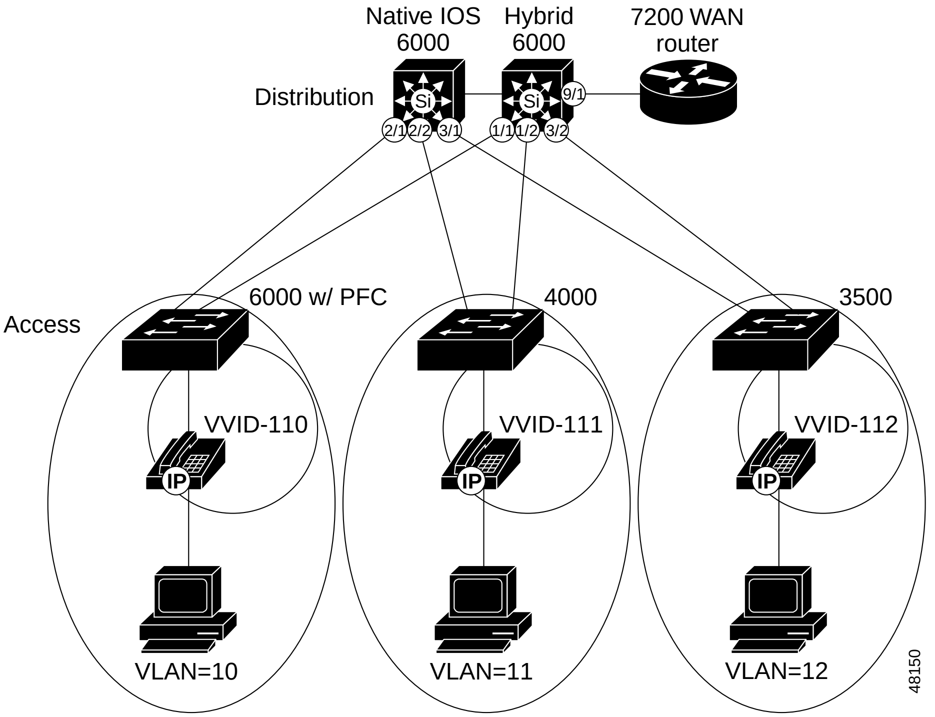

Step 1 Configure the Catalyst 6000 Distribution Layer VoIP Control Traffic Transmit Queues

All VoIP (CoS=5) traffic is placed into the egress interface PQ on 1p2q2t interfaces and Queue #2 on 2q2t (all “1” version of 10/100 linecards) interfaces as soon as QoS is enabled. You must also configure the Catalyst 6000 CoS queue admission rules to ensure CoS=3 traffic flows (VoIP control traffic) are placed into the second queue. Use the following commands:

```
cat6k-distrib> (enable) set qos map 1p2q2t tx queue 2 1 cos 3
cat6k-distrib> (enable) set qos map 2q2t tx queue 2 1 cos 3
```

Step 2 Configure the Catalyst 6000 Distribution Layer with a Catalyst 6000-PFC Access Layer

Once you enable QoS on the Distribution Layer switch and modified the default queue admission, you must:

- Trust DSCP from the Access Layer.
- Configure ToS to DSCP mappings.

These steps are explained below:

a. Trust DSCP from the Layer 3 Access Switch

Turn on trust for DSCP values from adjacent Layer 3 Access switches. Use **port-base qos** on the trunking port and use **trust-dscp** instead of **trust-cos**. This is because **trust-cos** will overwrite the Layer 3 DSCP value with the mapped CoS. There is no need to do this since classification is done at the Access Layer.

```
cat6k-distrib> (enable) set port qos 1/1 port-based
cat6k-distrib> (enable) set port qos 1/1 trust trust-dscp
```

b. Configure the Catalyst 6000 ToS-to-DSCP Mapping

Cisco follows the IETF recommendations for setting the DSCP classification values for both the VoIP Control plane traffic and VoIP Bearer or Media plane traffic. The recommended settings are:

- DSCP=AF31 for VoIP Control plane.
- DSCP=EF for VoIP Bearer plane.

To correctly map the Layer 3 IP Precedence settings to these DSCP values, you must modify the default ToS-to-DSCP mappings as in the following example:

```
cat6k-distrib> (enable) set qos ipprec-dscp-map 0 8 16 26 32 46 48 56

cat6k-distrib> (enable) sh qos map run ipprec-dscp-map
IP-Precedence - DSCP map:
IP-Prec   DSCP
-------   ----
      0   0
      1   8
      2   16
      3   26 -> 26 = AF31
      4   32
      5   46 -> 46 = EF
      6   48
      7   56
```

Step 3 Configuring Catalyst 6000 Distribution Layer with a Layer 2 only Access Switch

Once you enable QoS on the Distribution Layer switch and modify the default queue admission, you must:

- Trust CoS from the Access Layer.
- Configure CoS-to-DSCP mappings.

- Configure Layer 3 access lists for VoIP control traffic classification (See "Marking Control and Management Traffic").

a. Trust CoS from the Layer 2 Access Switch

Turn on trust for CoS values from adjacent Layer 2 Access switches. Use **Port-base qos** on the trunking port and use **trust-cos** instead of **trust-dscp** when the Access Layer switch is a Layer 2-only device performing CoS classification:

```
cat6k-distrib> (enable) set port qos 1/2,3/2 trust trust-cos
```

b. Configure Catalyst 6000 CoS-to-DSCP Mapping

Cisco follows the IETF recommendations for setting the DSCP classification values for both the VoIP Control plane traffic and VoIP Bearer or Media plane traffic. The recommended settings are:

- DSCP=AF31 for VoIP Control plane.
- DSCP=EF for VoIP Bearer plane.

To correctly map the Layer 2 to these DSCP values, you must modify the default CoS to DSCP mappings like the following:

```
cat6k-distrib> (enable) set qos cos-dscp-map 0 8 16 26 32 46 48 56

cat6k-distrib> (enable) sh qos map run cos-dscp-map
CoS - DSCP map:
CoS   DSCP
---   ----
  0   0
  1   8
  2   16
  3   26 -> 26 = AF31
  4   32
  5   46 -> 46 = EF
  6   48
```

c. Configure Layer 3 Access Lists for VoIP Control Traffic Classification

Use the following commands:

```
cat6k-distrib> (enable) set port qos 1/2,3/2 vlan-based
cat6k-distrib> (enable) set qos acl map ACL_IP-PHONES 111
(Use the ACL_IP-PHONES access-list from Marking Control and Management Traffic.)

cat6k-distrib> (enable) sh qos acl map run ACL_IP-PHONES
ACL name                         Type Vlans
-------------------------------- ---- ---------------------------------
ACL_IP-PHONES                      IP 110,111,112
ACL name                         Type Ports
-------------------------------- ---- ---------------------------------
ACL_IP-PHONES                      IP


cat6k-distrib> (enable) sh qos acl info run ACL_IP-PHONES

set qos acl IP ACL_IP-PHONES
----------------------------------------------
1. dscp 26 tcp any any range 2000 2002
2. dscp 26 tcp any any eq 1720
3. dscp 26 tcp any any range 11000 11999
4. dscp 26 udp any any eq 2427
5. trust-cos any
```

Step 4 Configuring the Connection to the 7200 WAN Router

```
cat6k-distrib> (enable) set port qos 9/1 port-based
```

```
cat6k-distrib> (enable) set port qos 9/1 trust trust-ipprec
Current 10/100 "1" linecards must still have trust-ipprec enabled even though the parser
returns an error

cat6k-distrib> (enable) set qos acl ip ACL_TRUST-WAN trust-ipprec any
cat6k-distrib> (enable) commit qos acl ACL_TRUST-WAN
cat6k-distrib> (enable) set qos acl map ACL_TRUST-WAN 9/1

cat6k-distrib> (enable) sh port qos 9/1
QoS is enabled for the switch.
QoS policy source for the switch set to local.

Port  Interface Type Interface Type Policy Source Policy Source
      config         runtime        config        runtime
----- -------------- -------------- ------------- -------------
 9/1      port-based     port-based          COPS          local

Port TxPort Type RxPort Type Trust Type  Trust Type   Def CoS Def CoS
                                config       runtime       config  runtime
--------------------------------------------------------------------
 9/1          2q2t  1q4t    trust-ipprec    trust-ipprec      0      0

Port  Ext-Trust Ext-Cos
----- --------- -------
 9/1  untrusted       0

(*)Runtime trust type set to untrusted.

Config:
Port  ACL name                         Type
----- -------------------------------- ----
 9/1  ACL_TRUST-WAN                    IP

Runtime:
Port  ACL name                         Type
----- -------------------------------- ----
 9/1  ACL_TRUST-WAN                    IP
```

Catalyst 6000 Distribution/Core Running Native Cisco IOS

After you configure the Access switch and attach it to the distribution layer, you must set up QoS on the distribution switches. This requires a few simple changes to the configuration:

1. Configure QoS.
2. Configure VoIP control traffic transmit queuing.
3. Configure the Distribution Layer with a Layer 3 Access switch.
   a.Trust ToS and DSCP from the Access Layer.
   b.Configure ToS to DSCP mappings.
4. Configure the Distribution Layer with a Layer 2 Access switch.
   a.Trust CoS and DSCP from the Access Layer.
   b.Configure CoS to DSCP mappings.
   c.Configure the QoS policies and Layer 3 access lists for VoIP control traffic classification.

All native Cisco IOS Catalyst 6000 Distribution Layer examples are based on Figure 4-19.

Step 1 Configure QoS on the native Cisco IOS Catalyst 6000

Use the **mls qos** Cisco IOS command to enable QoS on the Catalyst 6000 using the native Cisco IOS.

Step 2 Configure Transmit Queue Admission for VoIP control traffic

All VoIP (CoS=5) traffic will be placed into the egress interface Priority Queue on 1p2q2t interfaces and Queue #2 on 2q2t (all "1" version of 10/100 linecards) interfaces as soon as QoS is enabled. You must also configure the Catalyst 6000 CoS queue admission rules to ensure CoS=3 traffic flows (VoIP control traffic) are placed into the second queue. Use the following commands:

```
int range gigabitEthernet 1/1 - 2
 wrr-queue cos-map 1 2 2
 wrr-queue cos-map 2 1 3 4
```

Step 3 Configure the Catalyst 6000 native Cisco IOS distribution layer with a Catalyst 6000-PFC access layer

Once you enable QoS on the native Cisco IOS distribution layer switch and modify the default queue admission, you must:

- Trust DSCP from the Access Layer
- Configure ToS to DSCP Mappings

These steps are described below:

a. Trust DSCP from the Layer 3 Access Switch

Turn on trust for DSCP values from adjacent Layer 3 Access switches. Use **port-base qos** on the trunking port (**port-based qos** is enabled by default when you enable **mls qos**) and use **mls qos trust dscp** instead of the CatOS **trust-dscp**. Please note that the classification has already been accomplished at the Access Layer in this model. See the following example:

```
interface GigabitEthernet2/1
 description trunk port to PFC enabled cat6k-access
 no ip address
 wrr-queue cos-map 1 2 2
 wrr-queue cos-map 2 1 3 4
 mls qos trust dscp
 switchport
 switchport trunk encapsulation dot1q
 switchport mode trunk
```

b. Configure native Cisco IOS ToS-to-DSCP mapping for Layer 3 access switches

Cisco follows the IETF recommendations for setting the DSCP classification values for both the VoIP Control plane traffic and VoIP Bearer or Media plane traffic. The recommended settings are:

- DSCP=AF31 for VoIP Control plane.
- DSCP=EF for VoIP Bearer plane.

To correctly map the Layer 3 IP Precedence settings to these DSCP values, you must modify the default ToS-to-DSCP mappings. Please note that the Catalyst 6000 numerical values of **26** and **46** correlate to **DSCP=AF31** and **DSCP=EF** respectively. This is done in global configuration mode.

```
mls qos map ip-prec-dscp 0 8 16 26 32 46 56 0
```

Step 4 Configure Catalyst 6000 native Cisco IOS distribution layer with a Layer 2 only access switch

Once you enable QoS on the Distribution Layer switch and modify the default queue admission, you must:

- Trust CoS from the Access Layer.
- Configure CoS-to-DSCP mappings.
- Configure Layer 3 access lists for VoIP control traffic classification (See "Marking Control and Management Traffic").

These steps are explained below:

a. Trust CoS from the Layer 2 Access Switch

Enable trust for CoS values from adjacent Layer 2 Access switches. Use **port-base qos** on the trunking port and use the native Cisco IOS **mls qos trust cos** instead of CatOS **trust-cos** when the Access Layer switch is a Layer 2-only device performing CoS classification. See the following example:

```
interface GigabitEthernet2/2
 description trunk port to layer 2-only cat4k
 no ip address
 wrr-queue cos-map 1 2 2
 wrr-queue cos-map 2 1 3 4
 mls qos vlan-based
 mls qos trust cos
 switchport
 switchport trunk encapsulation dot1q
 switchport mode trunk
!
interface GigabitEthernet3/1
 description trunk port to layer 2-only 3500
 no ip address
 wrr-queue cos-map 1 2 2
 wrr-queue cos-map 2 1 3 4
 mls qos vlan-based
 mls qos trust cos
 switchport
 switchport trunk encapsulation dot1q
    switchport mode trunk
```

b. Configure the native Cisco IOS CoS-to-DSCP mapping for Layer 2 access switches

Cisco follows the IETF recommendations for setting the DSCP classification values for both the VoIP Control plane traffic and VoIP Bearer or Media plane traffic. The recommended settings are:

- DSCP=AF31 for VoIP Control plane,
- DSCP=EF for VoIP Bearer plane.

To correctly map the Layer 2 to these DSCP values, you must modify the default CoS-to-DSCP mappings. Note that the Catalyst 6000 numerical values of **26** and **46** correlate to **DSCP=AF31** and **DSCP=EF** respectively. Perform this in global configuration mode.

```
mls qos map cos-dscp 0 8 16 26 32 46 56 0
```

c. Configure the QoS Policies and Layer 3 Access Lists for VoIP Control Traffic Classification

The QoS configuration for native Cisco IOS Catalyst 6000s is very similar to the WAN router Cisco IOS configurations except for using policing for marking traffic flows and applying service policies to VLAN interfaces. The physical Gigabit Ethernet uplink ports are configured to use VLAN-based QoS with the **mls qos** VLAN-based native Cisco IOS interface commands. Finally, the service policy is applied to all VLAN traffic inbound on the uplink.

In the example below, three classes are defined:

- One for the VoIP media stream.
- One for the control traffic.
- One for all other traffic.

Traffic is filtered for these classes based on Layer 3/4 source and destination IP addresses and ports. Each of these classes is referenced in the Voice-QoS policy map. In the policy map statements, a policing function is used to classify all traffic that meets the entrance criteria matched with the class-map access lists.

Note The Catalyst 6000 native Cisco IOS software does not support the **set ip dscp** commands. Instead, the policing algorithm is used for classifying traffic.

In the following example, the policing code tags the traffic flows with DSCP values of AF31, EF, and 0 for VoIP Control traffic, VoIP Media traffic, and all other packets respectively. The “8000” flows size is low enough that any traffic will solicit tagging using the **conform-action set-dscp-transmit 26** syntax.

```
class-map match-all VoIP-Control
  match access-group 100
class-map match-all VoIP-RTP
  match access-group 101
class-map match-all Routine
  match access-group 102
!
!
policy-map Voice-QoS
  class VoIP-Control
     police 8000 8000 8000 conform-action set-dscp-transmit 26 exceed action transmit
  class VoIP-RTP
     police 8000 8000 8000 conform-action set-dscp-transmit 46 exceed-action transmit
  class Routine
     police 8000 8000 8000 conform-action set-dscp-transmit 0 exceed-action transmit
!
! access-list 100 looks for VoIP Control Traffic
access-list 100 permit tcp any any range 2000 2002
access-list 100 permit tcp any any eq 1720
access-list 100 permit tcp any any range 11000 11999
access-list 100 permit udp any any eq 2427
!
! access-list 101 looks for VoIP Bearer Traffic
access-list 101 permit udp any any range 16384 32767
!
! access-list 102 filters for routine traffic
access-list 102 permit ip any any
!
interface GigabitEthernet2/2
 description trunk port to layer 2-only cat4k
 no ip address
 wrr-queue cos-map 1 2 2
 wrr-queue cos-map 2 1 3 4
 ! inform the port that QoS will be VLAN-Based
 mls qos vlan-based
 switchport
 switchport trunk encapsulation dot1q
 switchport mode trunk
!
interface GigabitEthernet3/1
 description trunk port to layer 2-only 3500
 no ip address
 wrr-queue cos-map 1 2 2
 wrr-queue cos-map 2 1 3 4
 ! inform the port that QoS will be VLAN-Based
 mls qos vlan-based
 switchport
 switchport trunk encapsulation dot1q
 switchport mode trunk
!
interface Vlan111
 description voice vlan on cat4k
 ip address 10.1.111.77 255.255.255.0
 ip helper-address 10.1.10.10
 no ip redirects
 ! apply the QoS policy as an inbound policy
 service-policy input Voice-QoS
 standby 111 ip 10.1.111.1
!
interface Vlan112
 description voice vlan on 3500
 ip address 10.1.112.77 255.255.255.0
 ip helper-address 10.1.10.10
 no ip redirects
 ! apply the QoS policy as an inbound policy
 service-policy input Voice-QoS
 standby 112 ip 10.1.112.1
```

```
ios6k#sh mls qos
  QoS is enabled globally
  Microflow policing is enabled globally

  QoS is vlan-based on the following interfaces:
    Vl111 V112 Gi2/2 Gi3/1 Gi3/2 Gi3/3
    Gi3/4 Gi3/5 Gi3/6 Gi3/7 Gi3/8 Gi4/1 Gi4/2 Gi4/3 Gi4/4 Gi4/5
    Gi4/6 Gi4/7 Gi4/8 Fa9/1 Fa9/2 Fa9/3 Fa9/4 Fa9/5 Fa9/6 Fa9/7
    Fa9/8 Fa9/9 Fa9/10 Fa9/11 Fa9/12 Fa9/13 Fa9/14 Fa9/15 Fa9/16 Fa9/17
    Fa9/18 Fa9/19 Fa9/20 Fa9/21 Fa9/22 Fa9/23 Fa9/24 Fa9/25 Fa9/26 Fa9/27
    Fa9/28 Fa9/29 Fa9/30 Fa9/31 Fa9/32 Fa9/33 Fa9/34 Fa9/35 Fa9/36 Fa9/37
    Fa9/38 Fa9/39 Fa9/40 Fa9/41 Fa9/42 Fa9/43 Fa9/44 Fa9/45 Fa9/46 Fa9/47
    Fa9/48

QoS global counters:
    Total packets: 16750372458300
    Packets dropped by policing: 55930847232
    IP packets with TOS changed by policing: 16750372458300
    IP packets with COS changed by policing: 55945330688
    Non-IP packets with COS changed by policing: 16750372458300
```

Building a Branch Office

Section Highlights and Recommendations

- The branch WAN router must support the advanced QoS tools for IP Telephony WAN support.
- Use a switch that supports multiple queues.
- There is currently no way to pass Layer 3 ToS classification to Layer 2 CoS in the routers.
- Layer 3 ToS to Layer 2 CoS mapping will occur in routers with the addition of the Modular CLI in Cisco IOS version 12.1(5)T/12.2(2)T.

Recommended Branch Office Designs for IP Telephony

The traditional branch office design for 5 to 100 users consists of a branch router and an Ethernet switch. The routers handle all IP routing and WAN connectivity. The local PCs are connected to a small Ethernet switch that also connects to the router. There are two areas of concern for voice quality within the branch office: VoIP across the WAN and voice quality within the office. Details of WAN QoS tools for ensuring voice quality across the WAN are covered in "Enabling the Wide Area Network". This section addresses branch office design, IP addressing, and voice quality within the branch.

Figure 4-20 Possible QoS Problem Areas

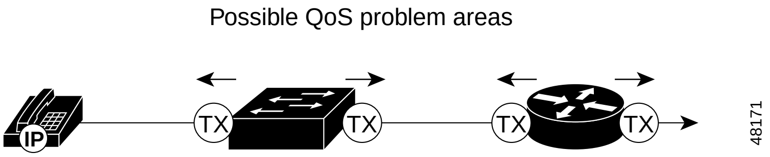

When these branch offices are typically designed, only a single IP subnet is used for each office. Changing this configuration is seldom feasible because it affects the enterprise-wide routing scheme. Therefore, realistic branch office designs must examine three IP addressing options for IP phones:

- Using a Single Subnet at the Branch Office
- Using 802.1Q for Trunking Separate Voice/Data Subnets at the Branch Office
- Using Secondary IP Addressing for Separate Voice/Data Subnets at the Branch Office

Note The IP addresses used for each of these second subnets can be RFC 1918 addresses for ease of management.

Each scenario uses the single cable installation method, which is the most common deployment method (see Figure 4-4).

Using a Single Subnet at the Branch Office

Using a single IP address space for branch offices my be a requirement where it is impractical or impossible to either allocate an additional IP subnet for IP phones or further subnet the existing IP address space used in the remote branch. When this is the case, you still need to prioritize voice over data at both Layer 2 and Layer 3. Layer 3 classification is already taken care of because the phone sets the ToS bits in all media streams to an IP Precedence value of 5 or a DSCP value of 40 (see the "Connecting the IP Phone" section on page 4-8 for more information). However, to ensure that there is Layer 2 classification for admission to the multiple queues in the branch office switches, the phone must also use the User Priority bits in the Layer 2 802.1p header to provide CoS marking. This can only be done by having the switch look for 802.1p headers on the *native* VLAN. The two Ethernet switches primarily used in branch office designs, the Catalyst 3500 and 4000, use different configuration commands to accomplish this.

Cisco 1750 Single Subnet Configuration

The Cisco 1750 series router does not support either ISL or 802.1Q Ethernet trunking. Below is an example of a single subnet 1750 configuration.

```
interface FastEthernet0
 mac-address 0000.1750.0001
 ip address 10.1.40.1 255.255.255.0
 ip helper-address 10.1.10.10
 ip policy route-map Set-IP-QoS
 no ip mroute-cache
 load-interval 30
 speed auto
 full-duplex
```

Catalyst 3500 Single Subnet Configuration

The Catalyst 3500 supports the use of an 802.1p-only option when configuring the Auxiliary VLAN. This allows the IP phone to tag VoIP packets with a CoS of 5 on the Native VLAN while all PC data traffic is sent un-tagged.

```
interface FastEthernet0/2
description Port to IP Phone in single subnet
  switchport trunk encapsulation dot1q
  switchport trunk native vlan 40
  switchport mode trunk
  switchport voice vlan dot1p
  spanning-tree portfast
!
interface FastEthernet0/15
  description Port to 1750 Router in single subnet
  load-interval 30
  duplex full
  speed 100
  switchport access vlan 40
```

Cisco 2600 Single Subnet (no Trunking) Configuration

```
interface FastEthernet1/0
 mac-address 0000.2600.0001
 ip address 10.1.60.1 255.255.255.0
 ip helper-address 10.1.10.10
 ip policy route-map Set-IP-QoS
 no ip mroute-cache
 load-interval 30
 speed 100
 full-duplex
```

Catalyst 4000 Single Subnet Configuration

Unlike the Catalyst 3500 and 6000, the 4000 does not support a dot1p-only option when configuring the Auxiliary VLAN. As an alternative, you should configure the IP phones connected to the Catalyst 4000 with the Auxiliary VLAN ID matching the Port VLAN ID (PVID) or Native VLAN. This will ensure that the phone can still send its packets tagged with a CoS of 5.

```
cat4k> (enable) set vlan 60 name 171.69.60.0_data
cat4k> (enable) set vlan 60 2/1-49
cat4k> (enable) set port host 2/1-49
cat4k> (enable) set port auxiliaryvlan 2/1-48 60
```

Using 802.1Q for Trunking Separate Voice/Data Subnets at the Branch Office

You should always use separate VLANs for voice and data when there is an option to segment the existing branch office IP address space. Ethernet switches that only support Layer 2 services, like the current 3500 and 4000 series, are used in almost every branch office design. When this is the case, the branch WAN routers will trunk the separate VLANs from the Ethernet switch. Use 802.1Q trunking on the router and switch to achieve this.

User Priority bits in the 802.1p portion of the 802.1Q standard header are used to provide prioritization in Ethernet switches. This is a vital component in designing IP Telephony-capable networks. One challenge facing initial IP Telephony branch office deployments is that the Cisco IOS routers cannot currently classify VoIP streams entering from the WAN to a Layer 2 802.1p CoS value for priority queuing within the branch's switched infrastructure.

In other words, if a branch IP phone is communicating with another branch IP phone, both phones use a CoS value of 5 on all Layer 2 packets and the VoIP streams from each phone is given priority in the branch office switched network.

If a branch IP phone is communicating with another IP phone at headquarters, the headquarters phone is classified with a ToS value of 5/EF and given priority throughout the WAN. When the VoIP stream on the headquarters phone hits the branch router, the router sends it to the Ethernet switch with a value of CoS=0 because the Cisco IOS router can not currently re-classify Layer 3 ToS settings to Layer 2 CoS settings. Cisco IOS version 12.2(2)T will address this issue with additions to the Modular CLI QoS code.

At the headquarters, this scenario is not an issue because the Catalyst 6000 correlates all Layer 3 ToS settings to the correct Layer 2 CoS values at the ingress interface.

3600 Branch Office Router using 802.1Q Trunking

```
interface FastEthernet1/0
 description Catalyst 3500 Branch Office Switch
 no ip address
 ip route cache policy
 no ip mroute-cache
 load-interval 30
 speed 100
 full-duplex
!
interface FastEthernet1/0.50
 description native subnet 10.1.50.0 data
encapsulation dot1Q 50
 ip address 10.1.50.1 255.255.255.0
 no ip mroute-cache
!
interface FastEthernet1/0.150
 description native subnet 10.1.150.0 voice
 encapsulation dot1Q 150
 ip address 10.1.150.1 255.255.255.0
 ip helper-address 10.1.10.10
 ip policy route-map Set-IP-QoS
 ip route cache policy
 no ip mroute-cache
```

Catalyst 3500 using 802.1Q Trunking

```
interface FastEthernet0/1
 description DOT1Q port to IP Phone
 switchport trunk encapsulation dot1q
 switchport trunk native vlan 50
 switchport mode trunk
 switchport voice vlan 150
 spanning-tree portfast
!
interface FastEthernet0/15
 description Port to 3640 (supports Dot1q)
 duplex full
 speed 100
 switchport trunk encapsulation dot1q
 switchport trunk native vlan 50
 switchport trunk allowed vlan 1,50,150
 switchport mode trunk
```

Catalyst 4000 using 802.1Q Trunking

```
cat4k> (enable) set vlan 70 name data70
cat4k> (enable) set vlan 170 name voice170
cat4k> (enable) set vlan 70 2/1-48
cat4k> (enable) set port host 2/1-48
cat4k> (enable) set port auxiliaryvlan 2/1-48 170
cat4k> (enable) set port speed 2/1-49  100
cat4k> (enable) set port duplex 2/1-49  full
cat4k> (enable) set trunk 2/49 on dot1q 1-1005
```

Using Secondary IP Addressing for Separate Voice/Data Subnets at the Branch Office

Separate VLANs for voice and data are still preferred in cases where the branch router doesn't support 802.1Q trunking. For example, the 1750 does not support trunking, but you still want a logical separation of voice and data traffic. An alternative to trunking is to use secondary IP addressing on the Cisco router, as in the following example:

1750 Branch Router

```
interface FastEthernet0
 description to Catalyst 3500
 mac-address 0000.1750.0001
 ip address 10.1.40.1 255.255.255.128
 ip address 10.1.40.129 255.255.255.128 secondary
 ip helper-address 10.1.10.10
 ip policy route-map Set-IP-QoS
 no ip mroute-cache
 speed 100
 full-duplex
```

Classifying VoIP Control Traffic at the Branch

The remote branch router also needs to classify VoIP Control traffic leaving the local subnets for a CallManager or VoIP gateway across the WAN. You can accomplish this by using the Policy-Based routing and Route-Maps on the ingress Ethernet interface.

Note With the release of *Encore*, Cisco CallManager will include the ability to configure the CoS and ToS values for all VoIP control and management traffic from the CallManager, IP phones, and Skinny Gateways. When this CallManager-based user-configurable classification is supported, network element access-lists will no longer be required for marking Skinny VoIP control traffic. H.323 and MGCP traffic will still require external, network element marking for the foreseeable future.

```
interface FastEthernet0
 mac-address 0000.1750.0001
 ip address 10.1.60.1 255.255.255.0
 ip helper-address 10.1.10.10
 !  Attach the route-map to the FastEthernet interface
 ip policy route-map Set-IP-QoS
 no ip mroute-cache
 load-interval 30
 speed auto
 full-duplex
!
!  Match all Skinny, H.323 and MGCP Control Traffic
access-list 101 permit tcp any any range 2000 2002
access-list 101 permit tcp any any eq 1720
access-list 101 permit tcp any any range 11000 11999
access-list 101 permit udp any any eq 2427
!
!  Match all VoIP RTP Traffic
access-list 102 permit udp any any range 16384 32767
!
!  Match all other traffic
access-list 103 permit ip any any
!
!  Set all Skinny, H.323 and MGCP Control traffic, matched
!  with ac 101 to IP Precedence 3
route-map Set-IP-QoS permit 10
 match ip address 101
 set ip precedence flash
!
!  Just match VoIP RTP Traffic; don't change the
!  default classification of ToS=5
route-map Set-IP-QoS permit 20
 match ip address 102
!
!  Make sure all data traffic is set to IP Precedence 0
route-map Set-IP-QoS permit 30
 match ip address 103
 set ip precedence routine
```

Enabling the Wide Area Network

Section Highlights and Recommendations

- Use Link Fragmentation and Interleave (LFI) techniques on all WAN connections with speeds below 768 Kbps.
- Use Low-Latency Queuing (LLQ) on all WAN VoIP connections.
- Traffic shaping is required for all Frame Relay and ATM deployments.
- Use cRTP wherever possible.
- ATM WANs operating at speeds below 768kbps must use Multilink PPP (MLPPP) over ATM to reduce frame sizes. MLPPP over ATM is supported in Cisco IOS version 12.1(4)T.
- Frame Relay to ATM Internetworking environments must use MLPPP over ATM and Frame Relay to reduce frame sizes on low-speed connections. MLPPP over ATM and Frame Relay will be supported in Cisco IOS version 12.1(4)T.
- Call Admission Control is a requirement when the number of calls across the WAN can overwhelm the provisioned VoIP bandwidth.

Refer to the following table for recommended Cisco IOS versions with QoS Tools by Supported Media:

Table 4-11 Recommended Cisco IOS Versions with QoS Tools

Media	Minimum Cisco IOS Version	Prioritization	LFI	Traffic Shaping
PPP	12.0(7)T	LLQ	MLPPP	N/A
Frame Relay	12.1(2)T	LLQ	FRF.12	Shape to CIR
ATM	12.1(2)T*	Per VC LLQ	MLPPP over ATM	Shape to Guaranteed Portion of Bandwidth
Frame Relay to ATM Internetworking	12.1(5)T*	Per VC LLQ	MLPP over ATM and Frame Relay	Shape Guaranteed Portion of Bandwidth

* When MLPPP over ATM and Frame Relay is supported, Cisco IOS version 12.1(4)T will be the minimum recommended version.

WAN QoS Overview

A lower total cost of ownership is one of the most compelling reasons for migrating to a converged data, voice, and video network. A converged network can lower overall costs of the enterprise communications infrastructure, but solid planning and design is still required for a successful IP Telephony deployment. This is especially important when running VoIP over a WAN.

You must use these three basic tools to provide an environment that can insure voice quality over a data network:

- Classification
- Queuing
- Network provisioning

When the low bandwidths and slow link speeds of a WAN are introduced into an IP Telephony design, you must also use several additional QoS tools:

- LFI
- Traffic Shaping
- Call Admission Control

Classification

Classification is the method by which certain traffic types are classified, or marked, as having unique handling requirements. These requirements might be a minimum required amount of bandwidth or a low tolerance for latency. This classification can be signaled to the network elements via a *tag* included in:

- IP Precedence/DSCP
- Layer 2 schemes such as 802.1p
- Source and destination IP addresses
- Implicit characteristics of the data, such as the traffic type using the RTP and a defined port range.

In the recommended IP Telephony QoS design model, this classification is done at both Layer 2 and Layer 3 on the IP phone. The phone, which is now the *edge* of the managed network, will set the Layer 2 802.1p CoS value to 5 and Layer 3 IP Precedence/DSCP to 5/EF. For more details on classification, see "Connecting the IP Phone".

Queuing

Interface queuing is one of the most important mechanisms for ensuring voice quality within a data network. This is even more vital in the WAN as many traffic flows are contending for a very limited amount of network resources. Once traffic has been classified, the flow can be placed into an interface egress queue that meets its handling requirements. VoIP, because of its extremely low tolerance for packet loss and delay, should be placed into a Priority queue. However, other traffic types may have specific bandwidth and delay characteristics as well. This is addressed with the Cisco LLQ feature in Cisco IOS.

LLQ combines the use of a priority queue with a class-based weighted fair-queuing scheme. Classes are defined with classification admission schemes. Traffic flows have access to either the PQ, one of the class-based queues, or a default WFQ. LLQ, the recommended queuing scheme for all low-speed links, allows up to 64 traffic classes with the ability to specify priority queuing behavior for voice, a minimum bandwidth for SNA data and IP Telephony control protocols, and WFQ to other traffic types.

When a Priority Queuing class is configured, the PQ has direct access to the tx-ring unless Interleaving is configured. In that case Interleaving occurs prior to placing the PQ traffic onto the TX ring. See the following illustration.

Figure 4-21 Layer 2 and 3 Queueing

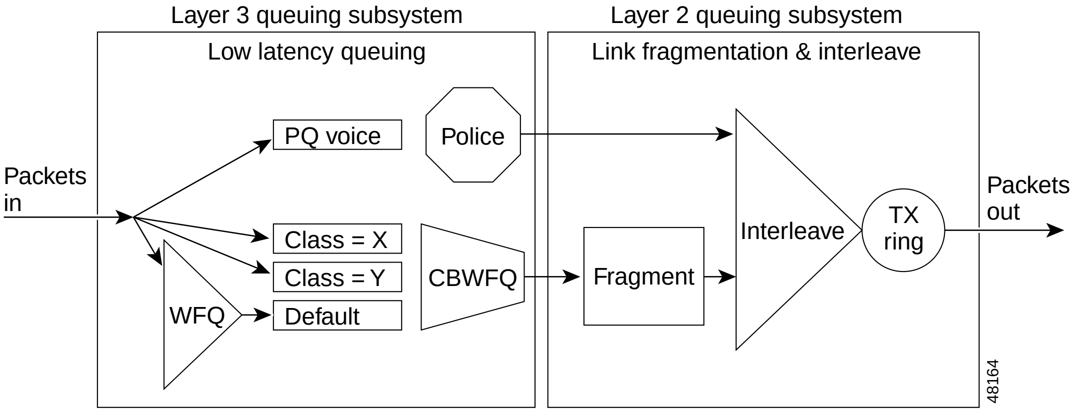

It is important to note that the maximum configured bandwidth in the priority queues and class-based queues cannot exceed the possible minimum amount of bandwidth on the WAN connection. A practical example is a 128 Kbps CIR Frame Relay LLQ scenario. If the priority queue for VoIP is configured for 64 Kbps and both the SNA and IP Telephony control protocol class-based queues are configured for 20 Kbps and 10 Kbps respectively, then the total configured queue bandwidth is 94 Kbps. Cisco IOS defaults to a Minimum CIR (mincir) value of CIR/2. Minimum CIR is the transmit value a Frame Relay router will *rate-down* to when backward explicit congestion notifications (BECNs) are received. In this example, MINCIR=64 Kbps and is lower than the configured bandwidth of the combined queues. For LLQ to work in this example, you should configure a MINCIR value of 128 Kbps.

Link Fragmentation and Interleave (LFI)

For low-speed WAN connections, for which a practical translation is a 768 Kbps clocking speed and below, it is necessary to provide a mechanism for LFI. A data frame can only be sent to the physical wire at the serialization rate of the interface. This serialization rate is the size of the frame divided by the clocking speed of the interface. For example, a 1500 Byte frame takes 214 Msec to serialize on a 56 Kbps circuit. If a delay-sensitive voice packet is behind a large data packet in the egress interface queue, the end-to-end delay budget of 150-200 Msec could be exceeded. Also, even relatively small frames can adversely affect overall voice quality by simply increasing the jitter to a value greater than the size of the adaptive jitter buffer at the receiver.

Table 4-12 Serialization Delay and LFI

		Serialization Delay					
Bytes		64	128	256	512	1024	1500
Link Speed	56kbps	9msec	18msec	36msec	72msec	144msec	214msec
	64kbps	8msec	16msec	32msec	64msec	128msec	187msec
	128kbps	4msec	8msec	16msec	32msec	64msec	93msec
	256kbps	2msec	4msec	8msec	16msec	32msec	46msec
	512kbps	1msec	2msec	4msec	8msec	16msec	23msec
	768kbps	640usec	1.28msec	2.56msec	5.12msec	10.4msec	15msec

LFI tools are used to fragment large data frames into regularly sized pieces and interleave voice frames into the flow so the end-to-end delay can be accurately predicted. This places bounds on jitter by preventing voice traffic from being delayed behind large data frames. The two techniques used for this are FRF.12 for Frame Relay and Multilink PPP for point-to-point serial links.

Figure 4-22 LFI—Before and After

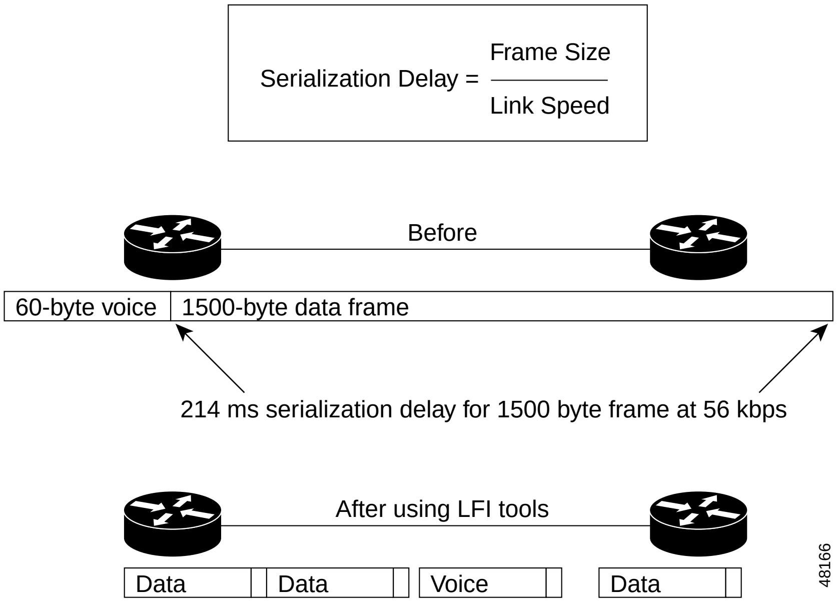
Serialization Delay = Frame Size / Link Speed
Before
60-byte voice
1500-byte data frame
214 ms serialization delay for 1500 byte frame at 56 kbps
After using LFI tools
Data
Data
Voice
Data
48166

A 10 Msec blocking delay is the recommended target to use for setting fragmentation size. Divide the recommended 10 Msec of delay by 1 Byte of traffic at the provisioned line clocking speed to achieve the recommended fragment size. This calculation is shown below:

Fragment Size = (Max_Allowed_Jitter * Link_Speed_in_kbps) / 8

An example is 70 Bytes = (10 Msec * 56) / 8

Table 4-13 Link Speed and Fragment Size

Link Speed	Recommended Fragment Size
56 kbps	70 bytes
64 kbps	80 bytes
128 kbps	160 bytes
256 kbps	320 bytes
512 kbps	640 bytes
768 kbps	960 bytes

Note In Cisco IOS version 12.1(5)T, MLPPP over ATM and Frame Relay is available to support LFI on ATM and ATM/Frame Relay Internetworking WANs.

Traffic Shaping

Traffic shaping is used in ATM and Frame Relay networks, where the physical access speed varies between two endpoints, to prevent excessive delay from congested network interface buffers caused by these speed mismatches. Traffic shaping is a tool that meters the transmit rate of frames from a source router to a destination router. This metering is typically done at a value that is lower than the line or circuit rate of the transmitting interface. This accounts for the circuit speed mismatches that are common in today’s multiple-access, non-broadcast networks. Traffic leaving a high speed interface such as a T1 at a central site often terminates at a remote site that may have a link speed which is much slower (56 Kbps, for example). This is quite common and is one of the major advantages of Frame Relay. In this example, the T1 interface on the central site’s router will send data out at a T1 rate even if the remote site has a clock rate of 56 Kbps. This buffers the frames within the carrier Frame Relay network, increasing variable delay. This same scenario can be applied in reverse. For example, when many remote sites, each with small WAN connections, are added together, they can over-subscribe the provisioned bandwidth or circuit speed at the central site. See the following illustration:

Figure 4-23 Traffic Shaping Scenario

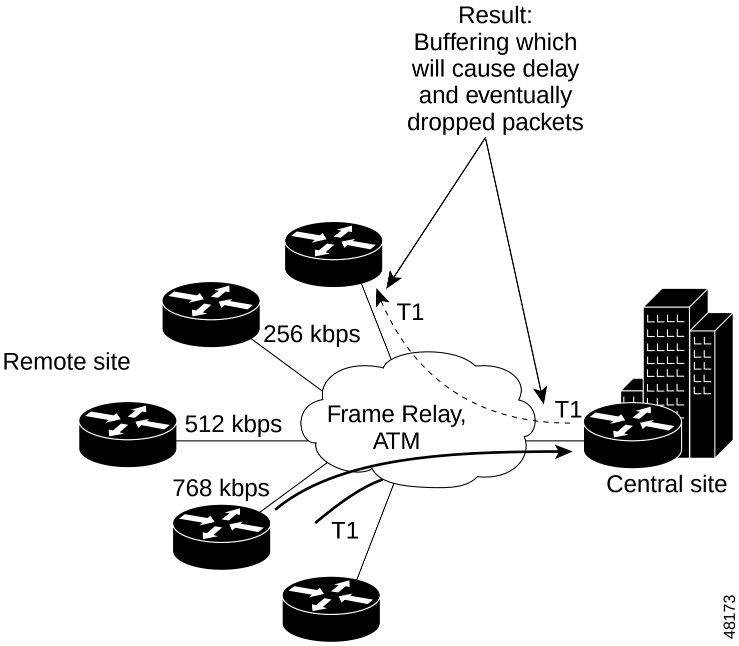

For a more detailed look at traffic shaping and Frame Relay designs, please refer to the Policing and Shaping Overview for more information.

Network Provisioning

A major component of designing a network for a successful IP Telephony deployment is properly provisioning the network bandwidth. You accomplish this by adding the bandwidth requirements for each major application (for example, voice, video, and data). This sum then represents the minimum bandwidth requirement for any given link and should only constitute approximately 75 percent of the link’s bandwidth. This 75 percent rule assumes that some bandwidth is required for overhead traffic, such as routing and Layer 2 keepalives, as well as for additional applications, such as E-mail and HTTP traffic.

Figure 4-24 Bandwidth Provisioning

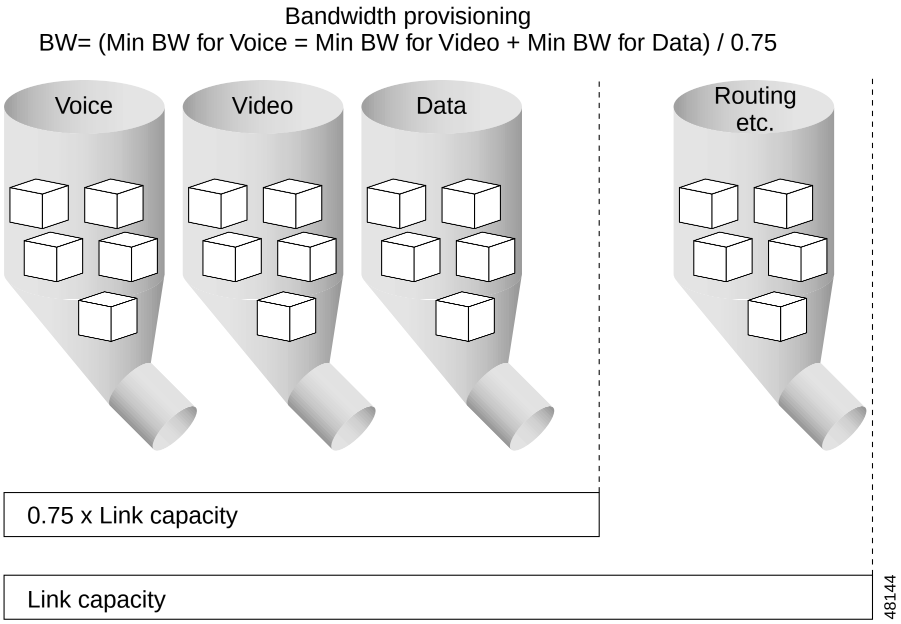

A VoIP packet consists of the payload, IP header, UDP header, RTP header, and Layer 2 Link header. At the default packetization rate of 20 Msec, VoIP packets have:

- 160 Byte payload for G.711.
- 20 Byte payload for G.729.
- IP header of 40 Bytes.
- UDP header of 8 Bytes
- RTP header of 12 Bytes.

The Link header varies according to media. See the following illustration for a sample VoIP packet:

Figure 4-25 VoIP Packet

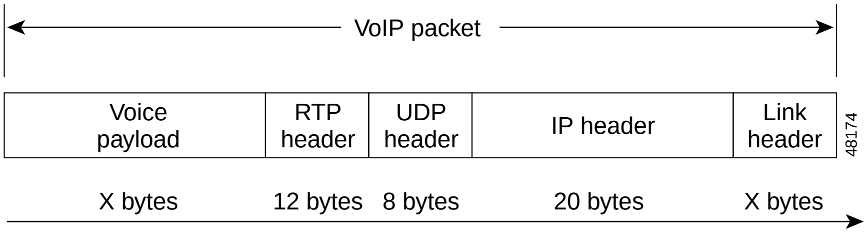

The bandwidth consumed by VoIP streams is calculated by adding the packet payload and all headers (in bits) then multiplying by the packet rate per second (default of 50 pps). The table below details the bandwidth per VoIP flow at a default packet rate of 50 pps. This does not include Layer 2 header overhead and does not take into account any possible compression schemes, such as compressed RTP (cRTP).

Note You can adjust the preferred rate of packets per second in the Service Parameters section of the CallManager configuration page.

Note While it is possible to configure the sampling rate above 30 msecs, this usually results in very poor voice quality.

Table 4-14, Part 1 Bandwidth Provisioning Chart

Codec	Sampling Rate	Voice Payload in Bytes	Packets per Second	Bandwidth per Conversation
G.711	20 Msec	160	50	80 Kbps
G.711	30 Msec	240	33	53 Kbps
G.729A	20 Msec	20	50	24 Kbps
G.729A	30 Msec	30	33	16 Kbps

A more accurate method for provisioning is to include the Layer 2 headers into the bandwidth calculations.

Table 4-14, Part 2 Bandwidth Provisioning Chart

Codec	Ethernet—14 Bytes of Header	PPP—6 Bytes of Header	ATM—53-byte Cells with a 48-byte Payload	Frame Relay—4 Bytes of Header
G.711	20 Msec	160	50	80 Kbps
G.711	30 Msec	240	33	53 Kbps
G.729A	20 Msec	20	50	24 Kbps
G.729A	30 Msec	30	33	16 Kbps

Admission Control

Admission control ensures that voice flows do not exceed the maximum provisioned bandwidth allocated for voice conversations.

After performing the calculations to provision the network with the required bandwidth to support voice, data, and possible video applications, it is important to ensure that voice does not over-subscribe the portion of the bandwidth allocated to it. While most QoS mechanisms are used to protect voice from data, Call Admission Control is used to protect voice from voice. The following illustration shows an environment where the network was provisioned to support two voice calls. However, if a third voice call is attempted, the quality of all calls will degrade.

For more detailed informatin about call admission control, go to the following location: http://www.cisco.com/univercd/cc/td/doc/product/voice/ip_tele/network/dgdistrb.htm.

Figure 4-26 Admission Control

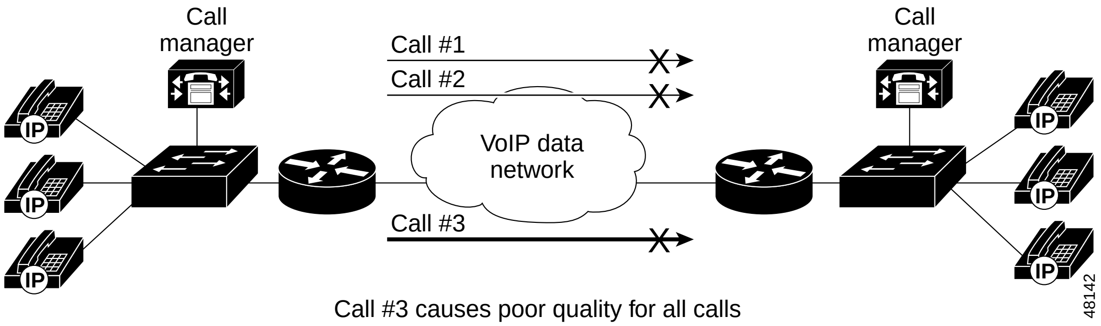

Miscellaneous WAN QoS Tools—VoIP Control Traffic

When allocating bandwidth for the IP WAN, CallManager control traffic is often overlooked. In Centralized Call Processing designs, the IP phones use a TCP control connection to communicate with the CallManager. If there is not enough bandwidth provisioned for these small control connections, the user might be adversely affected. An example where this comes into play is with the Delay to Dial-Tone (DDT) time periods. The IP phones communicate with the CallManager via Skinny Station Protocol over TCP port 2001. When an IP phone goes off-hook, it *asks* the CallManager what to do. The CallManager instructs the IP phone to play Dial-Tone. If this Skinny management and control traffic is dropped or delayed within the network, the user will not get Dial-Tone played out. This same logic applies to all signaling traffic for gateways and phones.

To ensure that this control and management traffic is marked as important (but not as important as voice), ACLs are used to classify these streams on Layer 3/4-aware Catalyst 6000 switches at the central locations. Examples of these configurations are included in "Enabling the High Speed Campus". In the remote offices, a Cisco router might be the first Layer 3/4-aware device a packet encounters before hitting the WAN. To ensure that these control connections are classified as important, but not as important as voice, access lists are used in the branch router. See the following example:

```
class-map VoIP-RTP
  match access-group 100
class-map VoIP-Control
  match access-group 101
!
policy-map QoS-Policy
  class VoIP-RTP
    priority 100
  class VoIP-Control
   bandwidth 8
  class class-default
   fair-queue
!
access-list 100 permit ip any any precedence 5
access-list 100 permit ip any any dscp ef
!
!  Skinny Control Traffic
access-list 101 permit tcp any host 10.1.10.20 range 2000 2002
!
!  MGCP Control Traffic
access-list 101 permit udp any host 10.1.10.20 2427
access-list 101 permit tcp any host 10.1.10.20 2428
!
!  H.323 Control Traffic
access-list 101 permit tcp any host 10.1.10.20 1720
access-list 101 permit tcp any host 10.1.10.20 range 11000 11999
```

Miscellaneous WAN QoS Tools—TX-Ring Sizing

The TX-Ring is the un-prioritized FIFO buffer used to hold frames prior to transmission in order to drive link utilization to 100 percent. In the Cisco 7500 RSP, this is referred to as the *TX-Queue* and you use the **tx-queue-limit** command to modify it. The RSP is a very inefficient QoS platform, especially in modifying the TX-Queue parameters. The 7500 RSP TX-Queue, which refers to the FIFO queue in MEM-D, has to copy the packet from MEM-D to the system buffers in DRAM and then back from the system buffers to MEMD. The TX-Ring, which is much more efficient than the TX-Queue, is used instead on 7500 VIP, 7200, 3600, 2600, and 1750 routers.

While fragmentation and interleaving reduces jitter, a large TX-Ring value can increase jitter when link utilization approaches saturation. Therefore, TX-Ring sizing is related to fragmentation size.

Note TX-Ring buffer sizing is measured in packets, not bits.

Table 4-15 Link Speed and TX-Ring Buffer Size

Link Speed/CIR/PVC	TX-Ring Buffer Sizing (packets)
=<128 Kbps	5
192 Kbps	6
256 Kbps	7
512 Kbps	14
768 Kbps	21

On all PPP and MLPPP, TX-Ring buffer size is automatically configured. These default buffer values cannot be changed. On Frame Relay links, the TX-Ring is for the main interface, which all subinterfaces use. The default value is 64 packets. This may need to be changed when the subinterface is very small or there are many subinterfaces.

Table 4-16 Media and TX-Ring Buffer Size

Media	Default TX-Ring Buffer Sizing (packets)
PPP	6
MLPPP	2
ATM	8192—must be changed for low speed virtual circuits (VCs)
Frame Relay	64 (per main T1 interface)

Miscellaneous WAN Qos Tools—Compressed Voice Codecs

In order to use as much of the limited WAN bandwidth as possible, VoIP uses Codecs to digitize analog voice samples. Many Codecs, such as G.729, can compress a 6 Kbps call down to 8 Kbps. These types of Codecs (low bit-rate Codecs) are commonly used for voice calls across the WAN.

Compressed RTP (cRTP)

cRTP compresses the 40-byte IP/UDP/RTP header of a VoIP packet to approximately 2 to 4 Bytes. cRTP works on a link-by-link basis and is enabled on Cisco routers using the **ip rtp header-compression** command. See the following table:

Note cRTP is currently only supported for leased lines and Frame Relay. Cisco IOS version 12.1(2)T, which greatly enhances performance, is the minimum recommended system software for scalable cRTP.

Table 4-17 cRTP Chart

Codec	PPP—6 Bytes of Header	ATM—53-byte Cells with a 48-byte Payload	Frame Relay—4 Bytes of Header
G.711 at 500 pps	68 Kbps	N/A	67 Kbps
G.711 at 500 pps	44 Kbps	N/A	44 Kbps
G.711 at 500 pps	12 Kbps	N/A	11.2 Kbps
G.711 at 500 pps	8 Kbps	N/A	7.4 Kbps

Miscellaneous WAN QoS Tools—Voice Activity Detection

Voice Activity Detection (VAD) takes advantage of the fact that in most conversations, only a single party is talking at a time. The VAD algorithm in the VoIP code examines the voice conversation, looking for these gaps in conversation. When one is discovered, no packets are sent and the WAN bandwidth can be recovered for use by data applications. We recommend you always turn VAD off throughout the system. It is on by default.

Note In environments that have a large amount of inherent delay, VAD can sometimes cause more voice quality issues than are justified by the bandwidth recovered. This should be examined on a case-by-case basis. However, when troubleshooting clipping at the beginning of conversations in an IP Telephony network, you should first disable VAD.

Point-to-Point WAN

Section Highlights

- Recommended minimum Cisco IOS is version 12.1(5)T.
- Use LFI techniques on all WAN connections with speeds below 768 Kbps.
- Use LLQ with a priority queue for VoIP bearer streams and a class-queue for VoIP control sessions.
- Call Admission Control is a requirement when the number of calls across the WAN can overwhelm the allocated VoIP bandwidth.

Figure 4-27 Point-to-Point WAN—QoS Problem Areas

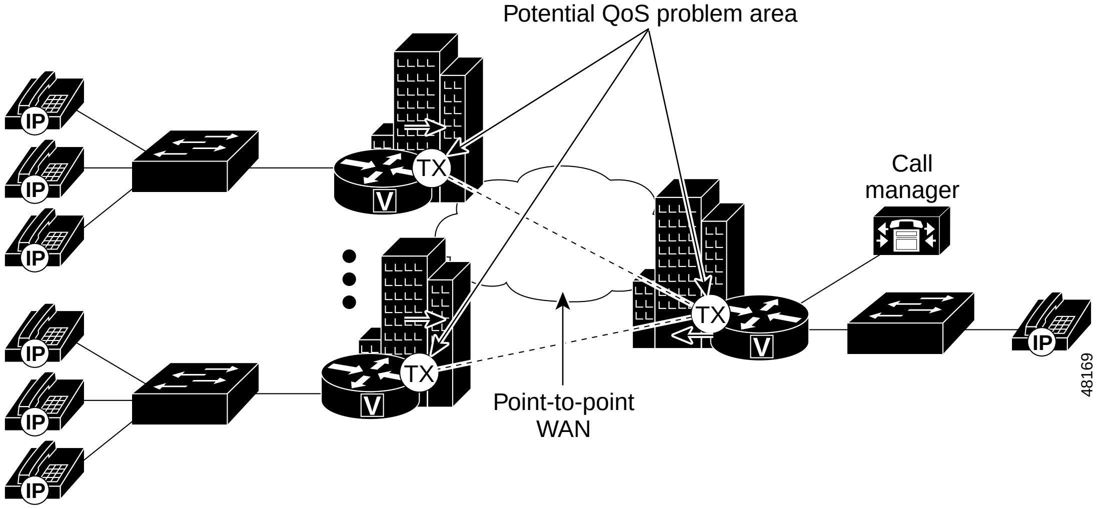

Point-to-point WANs, while not as popular as in the past, are still one of the most popular types of networks today. When designing a point-to-point WAN for an IP Telephony network, you should examine several QoS issues:

LFI on Point-to-Point WANs

If the clocking speed of the connection is below 768 Kbps, you must use LFI. You must also use MLPPP instead of PPP on all point-to-point links where LFI is required. To enable LFI on point-to-point WANs, you must use the MLPPP Cisco IOS command set.

Note When using MLPPP, fragmentation size is configured using the maximum acceptable delay in queue, which is 10 Msec. The TX-Ring is also statically configured at a value of two packets.

```
interface Multilink1
 ip address 10.1.61.1 255.255.255.0
 ip tcp header-compression iphc-format
 no ip mroute-cache
 load-interval 30
 service-policy output QoS-Policy
 ppp multilink
 ppp multilink fragment-delay 10
 ppp multilink interleave
 multilink-group 1
 ip rtp header-compression iphc-format
!
interface Serial0
 bandwidth 256
 no ip address
 encapsulation ppp
 no ip mroute-cache
 load-interval 30
 no fair-queue
 ppp multilink
 multilink-group 1
```

cRTP on MLPPP Connections

cRTP can have a dramatic impact on the amount of bandwidth each voice call uses. It is important to note that prior to Cisco IOS version 12.0(7)T, cRTP was process-switched. In fact, Fast-switching for cRTP didn't actually work on the 2600 and 3600 until a bug fix which was implemented in Cisco IOS version 12.0(7)T. Also, some of the newer versions of Cisco IOS (specifically version 12.1(2.x)T) process-switches cRTP.

```
interface Multilink1
 ip address 10.1.61.1 255.255.255.0
 ip tcp header-compression iphc-format
 no ip mroute-cache
 load-interval 30
 service-policy output QoS-Policy
 ppp multilink
 ppp multilink fragment-delay 10
 ppp multilink interleave
 multilink-group 1
 ip rtp header-compression iphc-format
```

LLQ for VoIP over MLPPP

LLQ is required to support voice over the WAN. When configuring LLQ for MLPPP-enabled interfaces, the service-policy output is placed into the Multilink Interface configuration. In the example below, two classes are defined: one for the VoIP media stream and one for the control traffic. Access to these classes, and therefore the queues they service, is done through access lists that match either Layer 3 TOS classification or source/destination IP addresses and ports. The access lists will look slightly different for the control traffic at the central site because a Catalyst 6000 has already classified VoIP Control sessions with a DSCP value of 26 (AF31, which is backward-compatible with IP Precedence 3).

All VoIP media traffic is placed into the PQ, which is given 100 Kbps. All skinny control traffic is placed-into a class-based queue and given 10 Kbps of bandwidth. All other traffic is queued using WFQ.

```
class-map VoIP-RTP
  match access-group 100
class-map VoIP-Control
  match access-group 101
!
policy-map QoS-Policy-256k
  class VoIP-RTP
    priority 100
  class VoIP-Control
   bandwidth 8
  class class-default
   fair-queue
!
interface Multilink1
 ip address 10.1.61.1 255.255.255.0
 ip tcp header-compression iphc-format
 no ip mroute-cache
 load-interval 30
 service-policy output QoS-Policy
 ppp multilink
 ppp multilink fragment-delay 10
 ppp multilink interleave
 multilink-group 1
 ip rtp header-compression iphc-format
!
!       ToS VoIP Media Stream Classification: either IP Prec or DSCP
!  This access-list is the same at the both the remote and
!  central locations
access-list 100 permit ip any any precedence 5
```

```
access-list 100 permit ip any any dscp ef
!
!  Skinny, H.323 and MGCP VoIP Control Traffic
!  which has already been classified using the
!  route-map in section 4.5.
access-list 101 permit ip any any precedence 3
access-list 101 permit ip any any dscp 26
```

Verifying Queuing, Fragmentation and Interleaving on an MLPPP Connection

```
1750# sh queue multilink1
  Input queue: 1/75/0/0 (size/max/drops/flushes); Total output drops: 8288
  Queueing strategy: weighted fair
  Output queue: 63/1000/64/8288/1967(size/maxtotal/threshold/drops/interleaves)
     Conversations  1/3/256 (active/max active/max total)
     Reserved Conversations 1/1 (allocated/max allocated)

  !  All drops and interleaves are occurring on ToS=0 flows
  (depth/weight/discards/tail drops/interleaves) 63/32384/8288/0/1967
  Conversation 60, linktype: ip, length: 1008
  source: 10.1.60.98, destination: 10.1.10.98, id: 0x0322, ttl: 63,
  TOS: 0 prot: 17, source port 1024, destination port 7


1750# sh policy interface multilink1
 Multilink1
 output : QoS-Policy-256k
  Class VoIP-RTP
   Weighted Fair Queueing
       Strict Priority
       Output Queue: Conversation 264
         Bandwidth 100 (Kbps)
         (pkts matched/bytes matched) 28100/5675882
         (pkts discards/bytes discards) 0/0
  Class VoIP-Control
   Weighted Fair Queueing
       Output Queue: Conversation 265
         Bandwidth 8 (Kbps) Max Threshold 64 (packets)
         (pkts matched/bytes matched) 204/10284
         (pkts discards/bytes discards/tail drops) 0/0/0
  Class class-default
   Weighted Fair Queueing
       Flow Based Fair Queueing
       Maximum Number of Hashed Queues 256
```

Frame Relay WAN

Section Highlights

- Recommended minimum Cisco IOS is version 12.1(5)T.
- You must use traffic shaping.
- Use LFI techniques on all WAN connections with speeds below 768 Kbps.
- Use LLQ with a priority queue for VoIP bearer streams and a class queue for VoIP control sessions.
- Call Admission Control is a requirement when the number of calls across the WAN can overwhelm the allocated VoIP bandwidth.

Figure 4-28 Frame Relay WAN—QoS Problem Areas

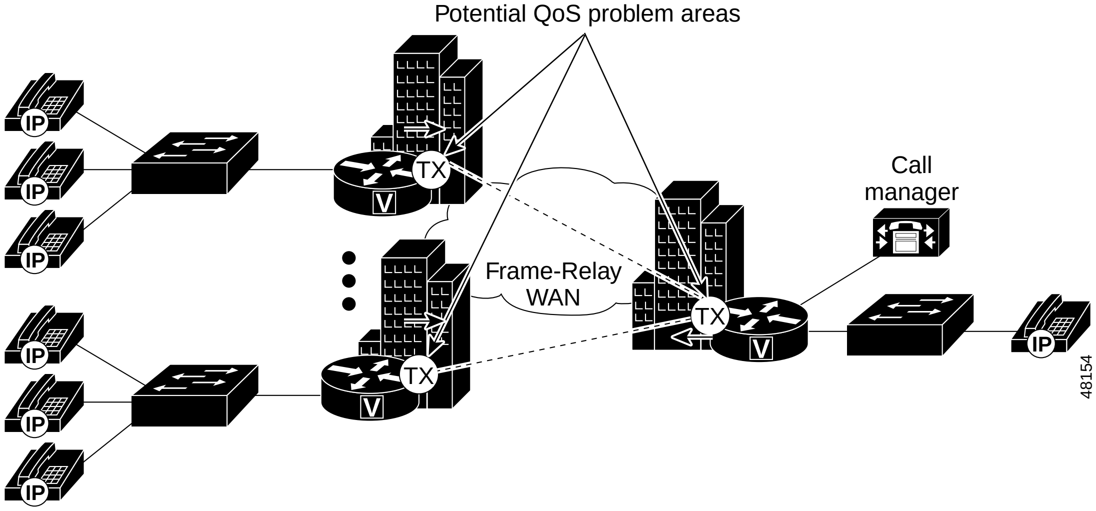

Frame Relay networks are the most popular WANs in use today because of their low cost. However, because Frame Relay is a non-broadcast technology that uses over-subscription to achieve costs savings, it is not always an easy media on which to deploy IP Telephony.

For more information on Frame Relay designs, please refer to the *Policing and Shaping Overview* document. This document can be found at the following location: http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/12cgcr/qos_c/qcpart4/.

Traffic Shaping

Traffic shaping is a requirement in Frame Relay networks for three reasons: over-subscription of sites is part of the nature of Frame Relay networks, bursting over CIR is commonly configured, and the default interval for Cisco Frame Relay devices can add unnecessary delay.

- CIR—in most Frame Relay networks, a central site uses a T1 link or greater to terminate WAN connections from many remote offices. The central site will send data out at 1.536 Mbps while a remote site may only have a 56 Kbps circuit. Additionally, there is typically a many-to-one ratio of remote offices to central hubs. It is quite possible for the remotes to all send traffic at a rate that can overwhelm the T1 at the hub. Both of these scenarios will cause frame buffering in the provider network that can induce delay, jitter, and drops. The only solution is to traffic shape at both the central and remote routers.
- Bc—another problem is the amount of data a Frame Relay node can transmit at any given time. A 56 Kbps PVC can transmit a maximum of 56 Kb of traffic in one second. How this second is divided is called the *interval*. The amount of traffic a node can transmit during this interval is called the Committed Burst (Bc) rate. By default, all Cisco routers set Bc to CIR/8. The formula for calculating the interval is:

  Interval = Bc/CIR or 125 Msec = 7000 / 56,000 for a 56 Kbps CIR

  In the example above, after a router sends its allocated 7000 bits, it must wait 125 Msecs before sending its next traffic. While this is a good default value for data, it is a very bad choice for voice. By setting the Bc value to a much lower number, the interval will decrease, which means the router will send traffic more frequently. An optimal configured value for Bc is 1000.

- Be—if the router doesn’t have enough traffic to send all of its Bc (1000 bits, for example), it can *credit* its account and send more traffic during a later interval. The maximum the router’s account can be credited is set at the Excess Burst (Be) rate. The problem with Be in IP Telephony networks is that this can create a potential for buffering delays within a Frame Relay network because the receiving side can only *pull* the traffic from a circuit at Bc, not Bc + Be.
- MINCIR—Cisco IOS defaults to a Minimum CIR value of CIR/2. Minimum CIR is the transmit value a Frame Relay router will *rate-down* to when BECNs are received. It is important to note that the maximum configured bandwidth in the priority queues and class-based queues cannot exceed the possible minimum amount of bandwidth on the WAN connection. An example of a remote site router connected to a 256 Kbps Frame Relay circuit is shown below:

```
interface Serial1
 no ip address
 encapsulation frame-relay
 load-interval 30
 frame-relay traffic-shaping
!
interface Serial1.71 point-to-point
 bandwidth 256
 ip address 10.1.71.1 255.255.255.0
 frame-relay interface-dlci 71
  class VoIP-256kbs
!
map-class frame-relay VoIP-256kbs
 frame-relay cir 256000
 frame-relay bc 1000
 frame-relay be 0
 frame-relay mincir 256000
 no frame-relay adaptive-shaping
 service-policy output QoS-Policy-256k
 frame-relay fragment 160
```

FRF.12 for LFI on Frame Relay WANs

To enable LFI on Frame Relay WANs, traffic shaping must also be used. Unlike MLPPP, the actual fragment size must be configured when using LFI on Frame Relay. See the following example:

```
map-class frame-relay VoIP-256kbs
 frame-relay cir 256000
 frame-relay bc 1000
 frame-relay be 0
 frame-relay mincir 256000
 no frame-relay adaptive-shaping
 service-policy output QoS-Policy-256k
 frame-relay fragment 160
```

cRTP on Frame Relay Connections

cRTP can have a dramatic impact on the amount of bandwidth each voice call uses. While cRTP fast-switching was enabled with Cisco IOS version 12.0(7)T, some of the newer version of Cisco IOS (specifically, 12.1(2.x)T) process-switches cRTP.

```
interface Serial1
 no ip address
 encapsulation frame-relay
 load-interval 30
 frame-relay traffic-shaping
 ip rtp header-compression iphc-format
```

LLQ for VoIP over Frame Relay

LLQ is required to support voice over the WAN. When configuring LLQ for Frame Relay-enabled interfaces, the service-policy output is placed into the **map-class frame-relay** configuration section. In the example below, two classes are defined: one for the VoIP media stream and one for the control traffic. Access to these classes, and therefore the queues they service, is performed through access lists that match either Layer 3 TOS classification or source/destination IP addresses and ports. The access lists will look slightly different for the control traffic at the central site because a Catalyst 6000 has already classified VoIP Control sessions with a DSCP value of 26 (AF31, which is backward-compatible with IP Precedence 3).

All VoIP media traffic is placed into the PQ, which is given 100 Kbps. All Skinny control traffic is placed into a class-based queue and given 10 Kbps of bandwidth. All other traffic is queued using WFQ.

```
class-map VoIP-RTP
  match access-group 100
class-map VoIP-Control
  match access-group 101
!
policy-map QoS-Policy-256k
  class VoIP-RTP
    priority 100
  class VoIP-Control
   bandwidth 8
  class class-default
   fair-queue
!
interface Serial1
 no ip address
 encapsulation frame-relay
 load-interval 30
 frame-relay traffic-shaping
!
interface Serial1.71 point-to-point
 bandwidth 256
 ip address 10.1.71.1 255.255.255.0
 frame-relay interface-dlci 71
  class VoIP-256kbs
!
map-class frame-relay VoIP-256kbs
 frame-relay cir 256000
 frame-relay bc 1000
 frame-relay be 0
 frame-relay mincir 256000
 no frame-relay adaptive-shaping
 service-policy output QoS-Policy-256k
 frame-relay fragment 160
!
!      ToS VoIP Media Stream Classification: either IP Prec or DSCP
!  This access-list is the same at the both the remote and
!  central locations
access-list 100 permit ip any any precedence 5
access-list 100 permit ip any any dscp ef
!
!  Skinny, H.323 and MGCP VoIP Control Traffic
!  which has already been classified using the
!  route-map in section 4.5.
access-list 101 permit ip any any precedence 3
access-list 101 permit ip any any dscp 26
```

Verifying Frame Relay Queuing, Fragmentation and Interleaving

```
3600# sh policy interface s 0/1.73
Remote Branch 3600
 Serial0/1.73: DLCI 73 -

  Service-policy output: QoS-Policy-256k (1117)

    Class-map: VoIP-RTP (match-all) (1118/2)
      5008 packets, 964953 bytes
      30 second offered rate 0 bps, drop rate 0 bps
      Match: ip precedence 5  (1120)
     Weighted Fair Queueing
       Strict Priority
       Output Queue: Conversation 40
         Bandwidth 100 (Kbps)
         (pkts matched/bytes matched) 4976/955161
        (pkts discards/bytes discards) 0/204

    Class-map: VoIP-Control (match-all) (1122/3)
      53 packets, 3296 bytes
      30 second offered rate 0 bps, drop rate 0 bps
      Match: ip precedence 3  (1124)
     Weighted Fair Queueing
       Output Queue: Conversation 41
         Bandwidth 8 (Kbps) Max Threshold 64 (packets)
         (pkts matched/bytes matched) 53/3296
         (pkts discards/bytes discards/tail drops) 0/0/0

    Class-map: class-default (match-any) (1126/0)
      5329 packets, 985755 bytes
      30 second offered rate 0 bps, drop rate 0 bps
      Match: any  (1128)
        5329 packets, 985755 bytes
        30 second rate 0 bps
     Weighted Fair Queueing
       Flow Based Fair Queueing
       Maximum Number of Hashed Queues 32


HQ_7200# sh frame-relay pvc int s6/0 73
Headquarters 7200

PVC Statistics for interface Serial6/0 (Frame Relay DTE)

DLCI = 73, DLCI USAGE = LOCAL, PVC STATUS = ACTIVE, INTERFACE = Serial6/0.73

  input pkts 114           output pkts 103          in bytes 8537
  out bytes 10633          dropped pkts 0           in FECN pkts 0
  in BECN pkts 0           out FECN pkts 0          out BECN pkts 0
  in DE pkts 0             out DE pkts 0
  out bcast pkts 62        out bcast bytes 5203
  pvc create time 00:04:22, last time pvc status changed 00:04:22
  service policy QoS-Policy-256k

  Service-policy output: QoS-Policy-256k (1099)

    Class-map: VoIP-RTP (match-all) (1100/2)
      0 packets, 0 bytes
      30 second offered rate 0 bps, drop rate 0 bps
      Match: ip dscp 46  (1102)
     Weighted Fair Queueing
       Strict Priority
       Output Queue: Conversation 72
         Bandwidth 100 (Kbps)
         (pkts matched/bytes matched) 0/0
         (pkts discards/bytes discards) 0/0

    Class-map: VoIP-Control (match-all) (1104/3)
      25 packets, 3780 bytes
      30 second offered rate 0 bps, drop rate 0 bps
      Match: ip dscp 26  (1106)
     Weighted Fair Queueing
       Output Queue: Conversation 73
         Bandwidth 8 (Kbps) Max Threshold 64 (packets)
         (pkts matched/bytes matched) 25/3780
         (pkts discards/bytes discards/tail drops) 0/0/0

    Class-map: class-default (match-any) (1108/0)
      163 packets, 15708 bytes
      30 second offered rate 0 bps, drop rate 0 bps
      Match: any  (1110)
        163 packets, 15708 bytes
        30 second rate 0 bps
     Weighted Fair Queueing
       Flow Based Fair Queueing
       Maximum Number of Hashed Queues 64
  Output queue size 0/max total 600/drops 0
  fragment type end-to-end         fragment size 160
  cir 768000    bc   7680      be 0         limit 960    interval 10
  mincir 768000    byte increment 960   BECN response no
  frags 125       bytes 10913     frags delayed 125   bytes delayed 10913

shaping inactive
 traffic shaping drops 0
```

ATM WAN

Section Highlights

- Recommended minimum Cisco IOS is version 12.1(5)T for MLPPP over ATM support.
- For all ATM connections below DS3 speeds, you will need to adjust the TX-Ring Buffer size.
- We recommend using two PVCs if the PVC speed is under 768 Kbps.
- If using a single PVC which is under 768 Kbps, use MLPPP over ATM for LFI.
- If using a single PVC, use LLQ with a priority queue for VoIP bearer streams and a class-queue for VoIP control sessions.
- Call Admission Control is a requirement when the number of calls across the WAN can overwhelm the allocated VoIP bandwidth

Figure 4-29 ATM WAN—QoS Problem Areas

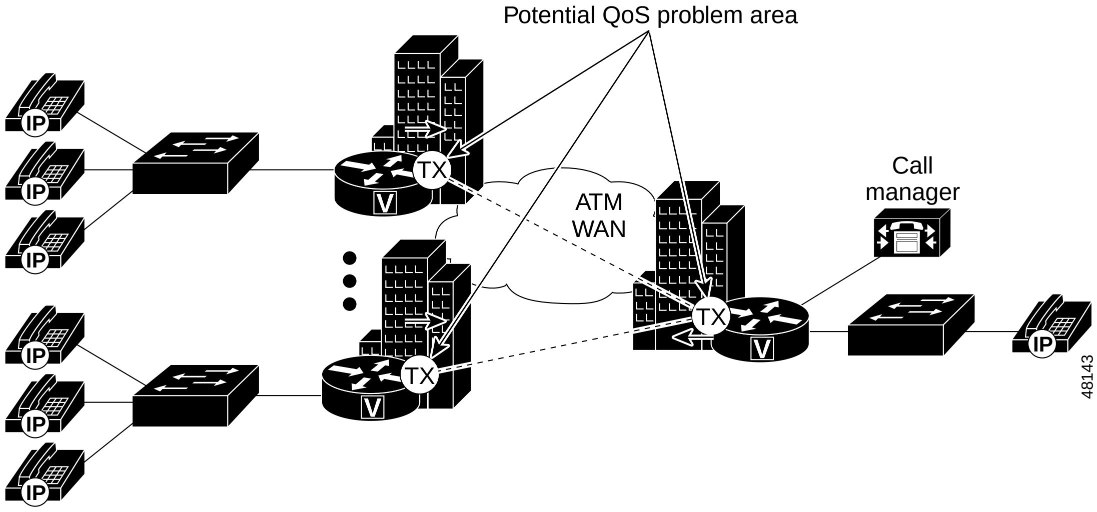

ATM is becoming an increasing popular media for WANs because many service providers have embraced the technology. One of the difficulties with deploying ATM in WANs is that it was designed for high speeds, not low speeds. Many enterprises are attempting to deploy IP Telephony over low-speed ATM connections. This generally results in complications because many of the Cisco IOS QoS tools are not currently supported on ATM interfaces and many of the interface defaults are automatically configured for high-speed ATM circuits.

This is evident in the default sizing of ATM TX-Ring Buffers. For example, by default, the 7200's OC-3 interface, the PA-A3, will set the TX-Ring to 8192. This is a correct setting for an OC-3, but for a 256 Kbps PVC configured on the interface, very large TX-Ring buffer delays can occur. Therefore, you have to configure the TX-Ring to a much lower value on a subinterface level. An example of a remote site router connected to a 256 Kbps ATM PVC is shown below:

```
interface ATM2/0
 no ip address
 no ip mroute-cache
 shutdown
 atm pvc 1 0 16 ilmi
 no atm ilmi-keepalive
!
interface ATM2/0.37 point-to-point
 pvc cisco37 0/37
  tx-ring-limit 7
  abr 256 256
  service-policy output QoS-Policy-256k
  protocol ppp Virtual-Template2
 !
!
```

Two PVCs or LFI on Low-Speed ATM WANs

The best method of designing VoIP over ATM networks when using PVCs lower than 768 Kbps is using separate PVCs for voice and data. This configuration is shown below.

```
interface ATM2/0.38 point-to-point
 bandwidth 256
 ip address 10.1.38.52 255.255.255.0
 pvc cisco38 0/38
  service-policy output Data-Policy-128k
  vbr-nrt 128 128
  encapsulation aal5snap
interface ATM2/0.39 point-to-point
 bandwidth 256
 ip address 10.1.39.52 255.255.255.0
 pvc cisco39 0/39
  tx-ring-limit 5
  service-policy output VoIP-Policy-128k
  vbr-nrt 128 128
  encapsulation aal5snap
```

If two PVCs are not an acceptable design alternative, the other option is to use the new MLPPP-over-ATM (MLPoATM) tools for LFI. Because ATM is a cell technology using a fixed payload size, there are no inherent LFI tools. A new standard, which uses MLPoATM, is available in Cisco IOS version 12.1(4)T. MLPoATM provides a Layer 2 Fragmentation and Interleaving method for low-speed ATM links.

The ideal fragment size for MLPoATM should allow the fragments to fit into an exact multiple of ATM cells. It is important to include MLPPP and AAL5 overhead in all fragmentation calculations. The MLPoATM header is 10 bytes and the AAL5 packet overhead is 8 bytes.

The fragment size for MLPoATM can be calculated as follows:

Fragment_Size = (48 * Number_of_Cells) - 10 - 8

For example, if seven cells per fragment is desirable, the fragment size should be 318 bytes.

There are a few interesting features to note when using MLPPP over ATM. These include the use of *Virtual Templates* instead of Multilink interfaces. Virtual Template configurations will be replaced by Multilink interfaces in later releases of the MLPoATM code. Multilink interfaces provide more scalability and will also provide greater integration into the existing MLPPP installations. The configuration of PPP CHAP is also required if remote sites want to communicate using MLPoATM. MLPoATM requires the MLP bundle to classify the outgoing packets before they are sent to the ATM VC. It also requires that the per-VC queuing strategy for the ATM VC be FIFO. There is some advanced ATM hardware that supports per-VC traffic shaping such as ATM Deluxe PA on 7200 and OC3 NM on 36x0. MLPoATM can only be supported on the platforms that support this ATM hardware. This configuration is shown below:

```
interface ATM2/0
 no ip address
 no ip mroute-cache
 shutdown
 atm pvc 1 0 16 ilmi
 no atm ilmi-keepalive
!
interface ATM2/0.37 point-to-point
 pvc cisco37 0/37
  tx-ring-limit 7
  abr 256 256
  protocol ppp Virtual-Template2
 !
!
interface Virtual-Template2
 bandwidth 254
 ip address 10.1.37.52 255.255.255.0
 service-policy output QoS-Policy-256k
 ppp authentication chap
 ppp chap hostname HQ_7200
 ppp chap password 7 05080F1C2243
 ppp multilink
 ppp multilink fragment-delay 10
 ppp multilink interleave
```

cRTP on ATM Connections

cRTP is not currently supported on ATM interfaces.

LLQ for VoIP over ATM

LLQ is required to support voice over the ATM WAN when a single PVC is deployed. When configuring LLQ for ATM-enabled interfaces, the service-policy output is placed under the subinterface PVC configuration section. In the example below, two classes are defined: one for the VoIP media stream and one for the control traffic. Access to these classes, and therefore the queues they service, is accomplished through access lists that match either Layer 3 TOS classification or source/destination IP addresses and ports. The access lists will look slightly different for the control traffic at the central site because a Catalyst 6000 has already classified VoIP Control sessions with a DSCP value of 26 (AF31, which is backward-compatible with IP Precedence 3).

All VoIP media traffic is placed into the PQ, which is given 100 Kbps. All Skinny control traffic is placed into a class-based queue and given 10 Kbps of bandwidth. All other traffic is queued using WFQ.

```
class-map VoIP-RTP
  match access-group 100
class-map VoIP-Control
  match access-group 101
!
policy-map QoS-Policy-256k
  class VoIP-RTP
    priority 100
  class VoIP-Control
   bandwidth 8
  class class-default
   fair-queue
!
interface ATM2/0
 no ip address
 no ip mroute-cache
 shutdown
 atm pvc 1 0 16 ilmi
 no atm ilmi-keepalive
!
interface ATM2/0.37 point-to-point
 pvc cisco37 0/37
  tx-ring-limit 7
  abr 256 256
  protocol ppp Virtual-Template2
 !
!
interface Virtual-Template2
 bandwidth 256
 ip address 10.1.37.52 255.255.255.0
 service-policy output QoS-Policy-256k
 ppp authentication chap
 ppp chap hostname HQ_7200
 ppp chap password 7 05080F1C2243
 ppp multilink
 ppp multilink fragment-delay 10
 ppp multilink interleave
 !
!
!
!      ToS VoIP Media Stream Classification: either IP Prec or DSCP
!  This access-list is the same at the both the remote and
!  central locations
access-list 100 permit ip any any precedence 5
access-list 100 permit ip any any dscp ef
!
!  Skinny, H.323 and MGCP VoIP Control Traffic
!  which has already been classified using the
!  route-map in section 4.5.
access-list 101 permit ip any any precedence 3
access-list 101 permit ip any any dscp 26
```

Frame Relay to ATM Internetworking WAN

Section Highlights

- Recommended minimum Cisco IOS is version 12.1(5)T for MLPoATM and MLPPP-over-Frame Relay (MLPoFR) support.
- FRF.8—Transparent Mode is the only support method for MLPoATM and Frame Relay Service Internetworking.
- For all ATM connections below DS3 speeds, you need to adjust the TX-Ring Buffer size.
- We recommend you use two PVCs if the ATM and Frame Relay PVC speed is under 768 Kbps.
- If using a single PVC that is under 768 Kbps, use MLPoATM and Frame Relay for LFI.
- If using a single PVC, use LLQ with a priority queue for VoIP bearer streams and a class queue for VoIP control sessions.
- Call Admission Control is a requirement when the number of calls across the WAN can overwhelm the allocated VoIP bandwidth.

Figure 4-30 Frame Relay to ATM—QoS Problem Areas

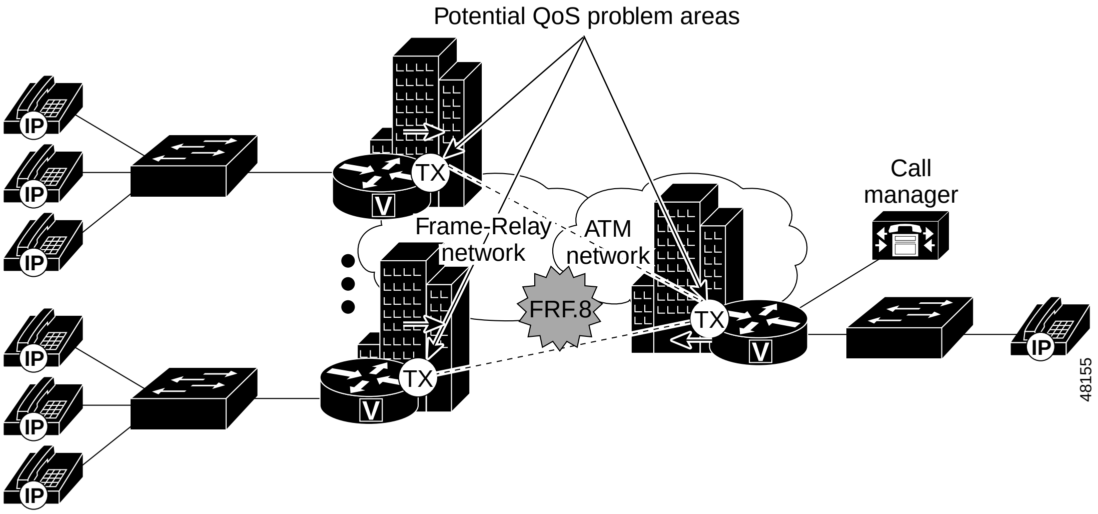

Many enterprises deploy IP Telephony over networks that use Frame Relay at the remote sites and ATM at the central location. The conversion is accomplished through ATM to Frame Relay Service Internetworking (FRF.8) in the carrier network.

Note MLPoATM and Frame Relay for LFI only supports Transparent Mode FRF.8.

LFI on Low-Speed ATM to Frame Relay Internetworking WANs

You cannot use FRF.12 because no service provider currently supports it. Furthermore, no Cisco WAN switching gear supports FRF.12. Tunneling FRF.12 through the service provider's network will do no good because there is no FRF.12 standard on the ATM side. This is a problem because fragmentation is required if any of the remote Frame Relay sites use a circuit of 768 Kbps or below. The best method of designing VoIP over ATM networks when using PVCs lower than 768 Kbps is to use separate PVCs for voice and data.

If using two PVCs is not practical, then you can use the new MLPoATM and Frame Relay tools for link Fragmentation and Interleave available in Cisco IOS version 12.1(5)T. MLPoATM and Frame Relay provide an end-to-end Layer 2 Fragmentation and Interleaving method for low-speed ATM to Frame Relay FRF.8 Service Internetworking links.

FRF.8 Service Internetworking is a Frame Relay Forum standard for connecting Frame Relay networks with ATM networks. Service Internetworking provides a standards-based solution for service providers, enterprises, and end users. In Service Internetworking translation mode, Frame Relay PVCs are mapped to ATM PVCs without the need for symmetric topologies; the paths can terminate on the ATM side. FRF.8 supports two modes of operation of the IWF for upper Layer user protocol encapsulation. These modes are:

- Translation mode—maps between ATM and Frame Relay encapsulation. It also supports internetworking of routed or bridged protocols.
- Transparent mode—does not map encapsulations but sends them unaltered. This mode is used when translation is impractical because encapsulation methods do not conform to the supported standards for Service Internetworking.

MLPPP for LFI on ATM and Frame Relay Service Internetworking networks is only supported in Transparent mode.

In order to make MLPoFR and MLPoATM internetworking possible, the internetworking switch must be configured in transparent mode and the end routers must be able to recognize both MLPoFR and MLPoATM headers. Enable this with the **frame-relay interface-dlci <*dlci*> ppp** command for Frame Relay and **protocol ppp** command for ATM.

When a frame is sent from the Frame Relay side of an ATM to Frame Relay Service Internetworking connection:

1. The sending router encapsulates a packet in the MLPoFR header.
2. The Carrier Switch (in Transparent mode) strips the 2-byte Frame Relay DLCI field and sends the rest of the packet to its ATM interface.
3. The receiving router examines the header of the received packet. If the first two bytes of the received packet are 0x03cf, it treats it as a legal MLPoATM packet and sends it to MLPPP Layer for further processing.

When an ATM cell is sent from the ATM side of an ATM to Frame Relay Service internetworking connection:

1. The sending router encapsulates a packet in the MLPoATM header.

2. The Carrier Switch (in Transparent mode) prepends a 2-byte Frame Relay DLCI field to the received packet and sends the packet to its Frame Relay interface.
3. The receiving router examines the header of the received packet. If the first four bytes after the 2-byte DLCI field of the received packet is **0xfefe03cf**, it treats it as a legal MLPoFR packet and sends it to MLPPP Layer for further processing.

A new ATM to Frame Relay Service Internetworking standard, FRF.8.1, will support MLPoATM and Frame Relay Service Internetworking. But it may be years before all switches are updated to the new standard.

The ideal fragment size for MLPoATM should allow the fragments to fit into an exact multiple of ATM cells. It is important to include MLPPP and AAL5 overhead in all fragmentation calculations. The MLPoATM header is ten bytes and the AAL5 packet overhead is 8 bytes.

The fragment size for MLPoATM can be calculated as follows:

Fragment_Size = (48 * Number_of_Cells) - 10 - 8

For example, if seven cells per fragment is desirable, the fragment size should be 318 bytes.

See the following configuration:

Central ATM Configuration

```
interface ATM2/0
 no ip address
 no ip mroute-cache
 shutdown
 atm pvc 1 0 16 ilmi
 no atm ilmi-keepalive
!
interface ATM2/0.37 point-to-point
 pvc cisco37 0/37
  tx-ring-limit 7
  abr 256 256
  protocol ppp Virtual-Template2
 !
!
interface Virtual-Template2
 bandwidth 254
 ip address 10.1.37.52 255.255.255.0
 service-policy output QoS-Policy-256k
 ppp authentication chap
 ppp chap hostname HQ_7200
 ppp chap password 7 05080F1C2243
 ppp multilink
 ppp multilink fragment-delay 10
 ppp multilink interleave
```

Remote Frame Relay Configuration

```
interface Serial6/0
 description T1 to Frame Relay switch
 no ip address
 encapsulation frame-relay
 load-interval 30
 no arp frame-relay
 frame-relay traffic-shaping
!
interface Serial6/0.73 point-to-point
 description 3640
 no arp frame-relay
 frame-relay interface-dlci 73 ppp Virtual-Template2
  class VoIP-256kbs
```

```
!
interface Virtual-Template2
 bandwidth 254
 ip address 10.1.37.51 255.255.255.0
 service-policy output QoS-Policy-256k
 ppp authentication chap
 ppp chap hostname R72HQ
 ppp chap password 7 05080F1C2243
 ppp multilink
 ppp multilink fragment-delay 10
 ppp multilink interleave
```

There are a few useful features to note when using MLPoATM. For instance, you can use Virtual Template interfaces instead of Multilink interfaces. However, Multilink interfaces will replace Virtual Template configurations in later releases of the MLPoATM code. Multilink interfaces provide more scalability and greater integration into the existing MLPPP installations. You must also configure PPP CHAP if remote sites need to communicate using MLPoATM. MLPoATM requires the MLP bundle to classify the outgoing packets before they are sent to the ATM VC. It also requires that the per-VC queuing strategy for the ATM VC be FIFO. There is limited advanced ATM hardware that supports per-VC traffic shaping such as ATM Deluxe PA on the 7200 and OC3 NM on 36x0. MLPoATM can only be supported on the platforms that support this ATM hardware.

Note MLPoFR relies on the Frame Relay traffic shaping (FRTS) engine to control the flow of packets from the MLP bundle to the FR VC.

cRTP on ATM to Frame Relay Connections

ATM interfaces do not currently support cRTP.

LLQ for Voice over ATM and Frame Relay

When using Service Internetworking, the LLQ configurations for Frame Relay and ATM links are identical to end-to-end MLPoATM. See "ATM WAN" for more detailed information.

Summary

You can accomplish guaranteed VoIP quality when you combine proper VoIP network design with the new Catalyst products, latest Cisco IOS images, and CallManager Call Admission Control technologies. When building an IP Telephony network, you should adhere to a few core principles:

- Use 802.1Q/p connections for the IP phones with the Auxiliary VLAN used for voice.
- Classify voice RTP streams as EF/IP Precedence 5 and place it into a second queue or a PQ on all network elements.
- Classify voice control traffic as AF31/IP Precedence 3 and place it into a second queue on all network elements.
- You must enable QoS within the campus if LAN buffers reach 100 percent utilization.
- Always properly provision the WAN by estimating that 25 percent of the bandwidth will be used for overhead, routing protocols, Layer 2 link information, and other miscellaneous traffic.
- You should use LLQ on all WAN interfaces.
- Use LFI techniques for all link speeds below 768 Kbps.

Designing Cisco CallManager Clusters

http://www.cisco.com/univercd/cc/td/doc/product/voice/ip_tele/network/dgclustr.htm

Selecting Gateways

http://www.cisco.com/univercd/cc/td/doc/product/voice/ip_tele/network/dggatewy.htm

Dial Plan Architecture and Configuration

http://www.cisco.com/univercd/cc/td/doc/product/voice/ip_tele/network/dgdialp.htm

Designing a Multi-site WAN with Distributed Call Processing

http://www.cisco.com/univercd/cc/td/doc/product/voice/ip_tele/network/dgdistrb.htm

Designing a Multi-site WAN with Centralized Call Processing

http://www.cisco.com/univercd/cc/td/doc/product/voice/ip_tele/network/dgcentrl.htm

Catalyst DSP Provisioning

http://www.cisco.com/univercd/cc/td/doc/product/voice/ip_tele/network/dgcatdsp.htm

Cisco Packet Fax and Modem Support Guidelines

This section details the various fax and modem transport solutions for IP Telephony packet telephony networks. New fax-relay and modem pass-through configurations and designs are available with the release of Cisco CallManager 3.0.1 and the new Catalyst 6000 gateway modules. However, the new fax-relay technology used in the Catalyst 6000 gateways is not compatible with the existing fax-relay code in the Cisco IOS router gateway products. This section describes which gateway fax and modem technologies can work together to provide fax and modem solutions.

It is important to note that faxes and modems are not supported over VoIP packet networks configured with the G.711 Codec for the following reasons:

- **Packet Loss**—A packet lost during the initial training interval will cause a re-train. Also, if enough packets are lost during the data transfer, the fax machines disconnect. This can be particularly problematic because modems can only tolerate a maximum number of re-trains before dropping the call. A packet lost during data transfer causes the modems to re-train.
- **Variable Delay and Jitter**—Because of the variable delay inherent in packet networks, the receiving VoIP gateway’s jitter buffer can be overrun or filled. If this occurs, packets containing the fax information can be dropped.

- **Clock Slew**—Clock slew refers to the difference in clock speeds at the ingress and egress gateways. If the difference is great enough, the jitter buffers within the gateways can be overrun, resulting in data loss and fax/modem re-trains.

Single switch configurations using a priority queue (PQ) for voice traffic can minimize the loss, delay, and clock synchronization problems. However, it is better to use the fax-relay, fax pass-through, and modem pass-through configurations. See the following sections for a detailed description of these three items.

Cisco IOS VoIP Router Gateways

Cisco IOS router gateways provide a Cisco proprietary fax-relay solution for fax over IP networks. This algorithm only works with other Cisco IOS router gateway endpoints. For example, a Cisco 1750 using fax-relay to transmit faxes across an IP network can communicate with any of the following VoIP endpoints: 1750, 3810v3, 2600, 3600, 7200, 5300, and the Catalyst 4000 VoIP gateway module using H.323v2.

Only the 5300 currently supports any type of modem-over-IP capabilities. See the "Cisco IOS Modem Support" section on page 4-74 for more information.

Cisco IOS Fax-Relay

The Cisco IOS VoIP router gateways offer fax-relay as a solution for fax over packet networks since Cisco IOS version 11.3. The fax-relay algorithm in Cisco IOS takes the analog fax transmission, de-modulates it, and encapsulates the resulting information into a digitally encapsulated packet. This packet is then re-modulated as an analog transmission at the far end and sent to the receiving fax machine. See the following sample configuration of a Cisco IOS VoIP dial-peer with fax-relay configured.

```
dial-peer voice 6743 voip
 destination-pattern  +3036636347
 prefix 6347
 Fax-rate 9600
 no vad
 ip precedence 5
 session target ipv4:172.247.70.134
```

Cisco IOS Modem Support

The Cisco 5300 is the only Cisco IOS VoIP gateway that supports modem transport over a VoIP network. Modem pass-through was created to address the possibility of packet loss, delay, and jitter introduction in a packet network. With the 5300 Cisco IOS version 12.1(3)T), Cisco announced a new method: modem pass-through.

This new modem pass-through code uses a variety of changes in the gateway to successfully handle modem transmission:

- Turns off echo cancellation
- Turns off VAD
- Eliminates the high pass filter
- Creates RTP payload redundancy
- Creates a static jitter buffer of 200ms

However, a highly available, low-loss infrastructure must be in place for modem pass-through to work. You must also use the G.711 Codec.

Figure 4-31 Cisco IOS Gateway Fax-Relay

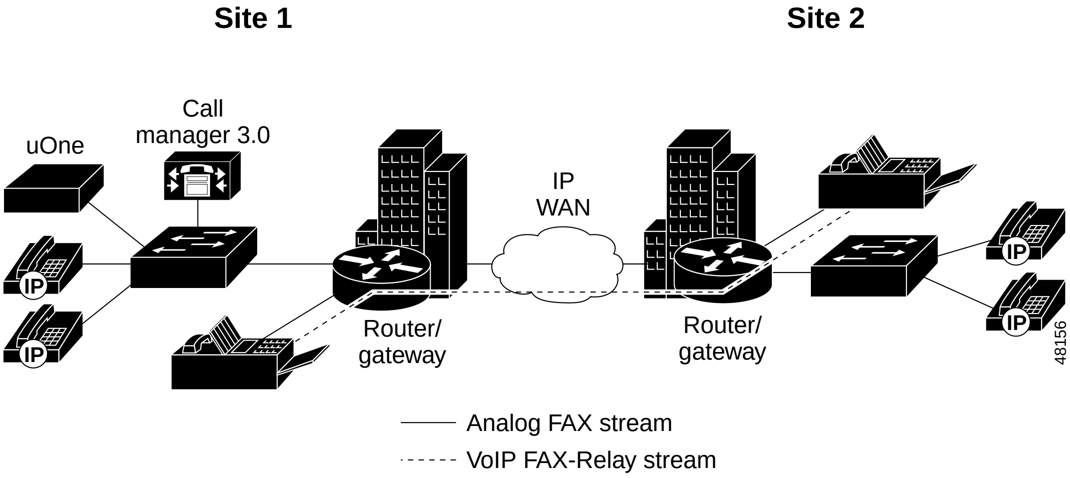

Cisco VG200

The Cisco VG200 is an IP Telephony gateway that supports either MGCP or H.323v2 for Cisco CallManager communication.

VG200 Fax Support

There are two modes:

MGCP Mode:

- The VG200 does not offer any fax-relay functionality when configured in MGCP mode.
- Although this design is not a supported configuration, it meets the following design conditions:
  - Both VG200s are attached to the same Catalyst 6000 or 3500 Ethernet switch. You should configure this switch with all voice sessions using a PQ.
  - VAD is disabled in the fax dial-peers.
  - The jitter buffer on the receiving fax port is set to the highest value.
  - Echo Cancellation is turned off on the receiving fax port.

H.323v2 Mode:

- Because the FXS interface has not yet been tested on the VG200, the Cisco IOS fax-relay option is not supported.
- Once the FXS interface and Cisco IOS fax-relay code are tested on the VG200 running H.323v2, the VG200 will be able to work with any Cisco IOS VoIP router gateway.

VG200 Modem Support

Modems over VoIP networks are not supported using the VG200 gateways.

Catalyst 6000 VoIP Gateways

With the release of Cisco CallManager 3.0(1), two new VoIP IP Telephony gateway modules are available on the Catalyst 6000. These two modules, a 24-port analog FXS gateway and an 8-port T1/E1 gateway, support both fax and modem transmission over IP Telephony VoIP networks. See Figure 4-32 for a sample network design using the Catalyst 6000 VoIP gateways for fax and modem transport.

Catalyst 6000 Fax-Relay (G.729A)

When using G.729A voice compression, both gateways support fax over IP using a proprietary fax-relay algorithm. This fax-relay algorithm does not currently work with the fax-relay code on the Cisco IOS gateways. However, this problem should be corrected soon. Fax-relay between Catalyst 6000 gateway modules is supported.

Catalyst 6000 Fax Pass-through (G.711)

When using the G.711 Codec for VoIP sessions, the Catalyst 6000 gateways use fax pass-through for fax transmission. The fax pass-through functionality is similar to the modem pass-through (see Catalyst 6000 Modem Pass-Through (G.711)) and to the 5300 modem pass-through algorithm used in Cisco IOS version 12.1(3)T. Fax pass-through is enabled on a G.711 connection by default.

For fax pass-through to work, a highly available, low-loss infrastructure must be in place. You must also configure G.711.

Catalyst 6000 Modem Pass-Through (G.711)

The new modem pass-through code available on the Catalyst 6000 VoIP gateways is similar to the modem pass-through technology on the 5300 introduced in Cisco IOS version 12.1(3)T. It uses a variety of changes in the gateway code to successfully handle modem transmission. The changes performed on the incoming port receiving the modem transmission is dynamically configured to do the following:

- Turn off echo cancellation.
- Turn off voice activity detection.
- Eliminate the high pass filter.
- Create RTP payload redundancy (RFC 2198).
- Create a static jitter buffer of 200ms.

However, for modem pass-through to work, a highly available, low-loss infrastructure must be in place. Also, G.711 is the only supported Codec.

Figure 4-32 Catalyst 6000 Fax and Modem

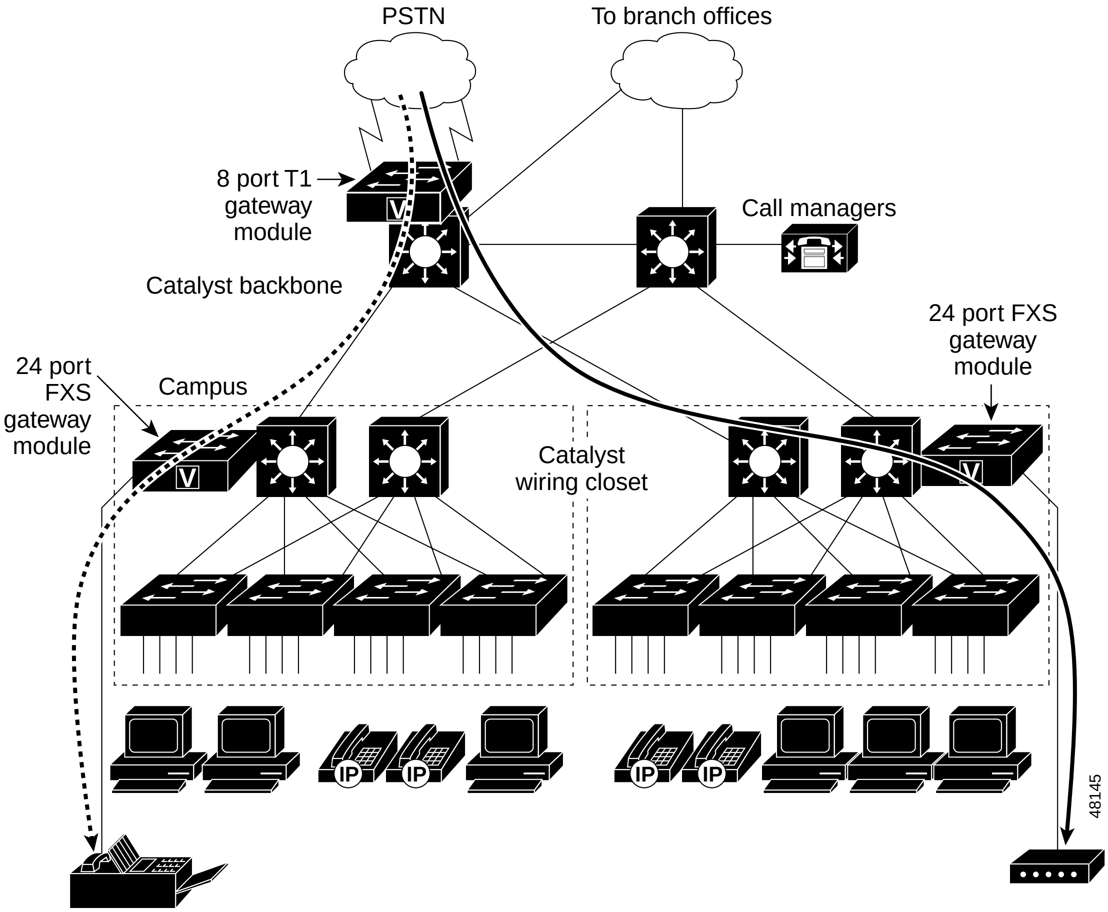

DT-24+/DT-30+ Gateways

- DT-24+/DE-30+ Fax Support—the DT-24+ or DE-30+ gateways do not support fax over VoIP networks.
- DT-24+/DE-30+ modem Support—the DT-24+/DE-30+ gateways do not support modems over VoIP networks.

Future T.38 Fax-relay Support

All Cisco VoIP implementations are moving toward compliance with the T.38 standard for fax-relay over IP networks. The Cisco 3810, 2600, 3600 and 5300 products will have T.38 capability in Cisco IOS version 12.1.(3)T. T.38 capability is a future feature for the 1750, 7200, 7500, Catalyst 6000, and VG200 gateways.

E911 and 911 Emergency Services

Cisco telephony systems architects should read this section to properly deploy IP Telephony solutions in regards to lifeline issues.

As a rule of thumb, a systems architect should clearly plan how a system handles emergency calls. You know what call processing to apply when a user dials an emergency number (e.g., 9-1-1 in North America).

The general responsibility of the Enterprise Telephony System is to:

- Route the call to the appropriate point (either on-net or off-net).
- When applicable, deliver calling party identification, typically the caller's phone number (either as an extension number or PSTN number).

The implications of properly accounting for emergency calls extend beyond the proper configuration of the IP Telephony telephony solution. It will often impact the actual equipment type and count, as well as the required type and quantity of interfaces to LEC systems. Furthermore, this document can only provide technical guidelines. There are no two service areas where Emergency call implementations are identical, and many of the factors that drive the design of emergency call handling are matters of policy rather than technology. These political factors are outside the scope of this document.

Today's E9-1-1 Service

Use this section as a quick reference on the existing North American telephony infrastructure supporting 9-1-1 calls. It explains the typical topology of the E9-1-1 network in North America, which is required to properly plan how Enterprise Telephony handles 9-1-1 calls.

For areas outside of North America, the handling of emergency calls is generally very similar to that of non-emergency calls. In most cases, the same switching and routing equipment is used for both call types, and the special treatment of emergency calls is limited to having a dedicated number (the equivalent of North America's 9-1-1) reserved for the public to dial. The caller's phone number may be delivered to the answering point, but the caller's location typically is not. You should inquire with the proper authorities to establish the governing rules and best practices in handing off emergency calls to the LEC.

Service Provider Perspective

In North America, E9-1-1 is an integral part of the phone service. All LEC telephony service providers have to provide the functionality, as mandated by federal, state, provincial, or local laws. To do this, the vast majority of North American telephony service providers rely on an overlay network to handle the voice and data traffic outside of the switched dial tone network.

Note As of December 1999, only two major RBOCs rolled out an E9-1-1 solution that uses SS7 for part of the routing of 9-1-1 calls.

This network is composed of analog MF signaling trunks dedicated to handling 9-1-1 traffic (for the aggregation and routing of the calls), and either analog MF signaling trunks or ISDN lines (PRI or BRI) to deliver the calls to the Public Safety Answering Point (PSAP).

There is a misconception that 9-1-1 call handling is a function of the PSTN. In most cases, the only part of the PSTN used for handling 9-1-1 calls is the line side of the CLASS 5 switch serving the local customer. The telephony service provider's E9-1-1 responsibilities are to:

- Automatically route 9-1-1 calls to the appropriate PSAP.
- Forward the calling party phone number to the appropriate PSAP.
- Provide an address database link to PSAP for retrieving Automatic Location Identification (ALI) (this may not be provided by the telephony Service Provider).

Unlike PSTNs, 9-1-1 calls are routed on the calling number, not the called number. Figure 4-33 highlights the 9-1-1 call flow through a likely E9-1-1 architecture.

Note The information is provided solely to support our examples and does not represent the actual E9-1-1 topology for the San Francisco bay area.

It depicts the voice and information flows associated with a PSTN subscriber dialing 9-1-1 from a phone line directly connected to a CLASS 5 switch. Also, the following examples do not highlight the specifics of an Enterprise Telephony subscriber making a 9-1-1 call.

The top part of the drawing shows the signaling and bearer connections between our two CLASS 5 switches for PSTN traffic. We include this detail to reinforce the fact that the PSTN SS7 fabric and associated switched bearer traffic trunks are *not* involved in routing 9-1-1 calls.

Figure 4-33 Typical E9-1-1 Network Topology

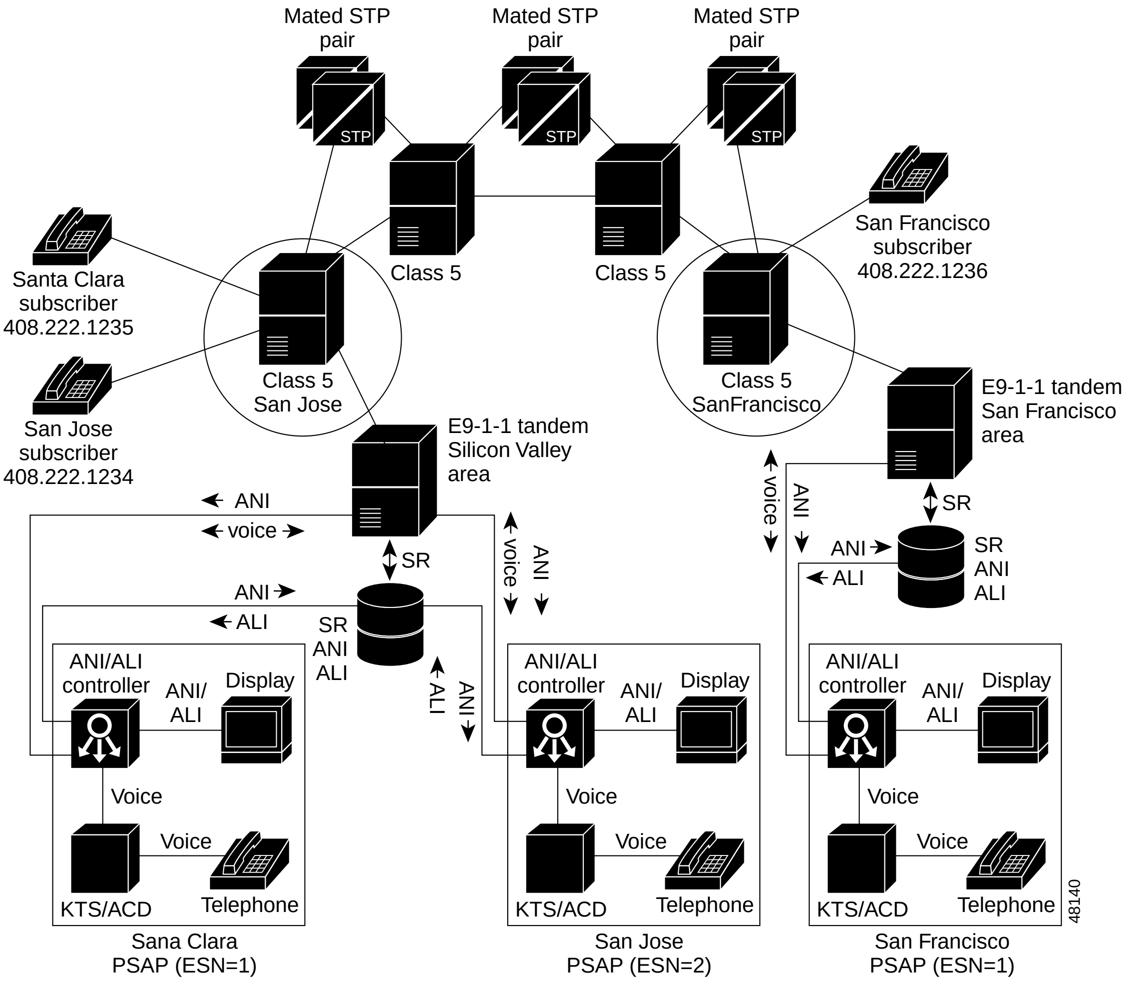

Example 1:

In this example a Santa Clara subscriber dials 9-1-1. We will follow the call until it is connected to the call taker.

Assumptions:

- The line's number is 408.222.1235.
- The serving switch is physically located in San Jose, even though it is also serving customers in Santa Clara.
- For each phone number in the geographical area served by the Silicon Valley area E9-1-1 tandem switch, a database provides an association between the caller's ANI, Emergency Service Number (ESN), and ALI. This association is created through collaboration between the Local Exchange Carrier (using the service records showing where that phone line is installed) and civil authorities (using the Master Street Address Guide (MSAG)). In other words, based on the phone number, the tandem will learn which PSAP receives the call. In our case, the ESN associated with the Santa Clara PSAP is 1.

Process:

1. The Santa Clara subscriber goes off-hook and obtains dial tone from San Jose CLASS 5 switch.
2. The subscriber dials 9-1-1.
3. The San Jose switch routes the call to outgoing trunks connected to the E9-1-1 tandem Silicon Valley area. Please note that the PSTN call path is not taken.
4. The San Jose switch sends the ANI (408.222.1235) through MF signaling on the CAMA trunk.
5. The tandem queries the Selective Routing (SR) database for the ESN associated with the ANI.
6. The tandem receives ESN=1. This prompts the routing of the call toward the Santa Clara PSAP.
7. The tandem sends the ANI information to the Santa Clara ANI/ALI controller.
8. At this point, the voice path is cut through in both directions. The ANI/ALI controller provides ringing progress tone.
9. The ANI/ALI controller queries the database for the ALI corresponding to the ANI.
10. As the Santa Clara call taker answers the call, the voice, ANI, and ALI of the caller are presented.

Example 2:

In this example, a San Jose subscriber places the 9-1-1 call.

Assumptions:

- The line's number is 408.222.1234.
- The serving switch is physically the same as in Example 1.
- The ESN associated with the San Jose PSAP is ESN2.

Process:

1. The San Jose subscriber goes off-hook and obtains a dial tone from the San Jose CLASS 5 switch.
2. The subscriber dials 9-1-1.
3. The San Jose switch routes the call to outgoing trunks connected to the E9-1-1 tandem Silicon Valley area. Please note that the PSTN call path is not taken.
4. The San Jose switch sends the ANI (408.222.1234) through the MF signaling on the CAMA trunk.
5. The tandem queries the Selective Routing (SR) database for the ESN associated with the ANI.
6. The tandem receives ESN=2. This prompts the routing of the call toward the San Jose PSAP.

7. The tandem sends ANI information to the San Jose ANI/ALI controller.
8. At this point, the voice path is cut through in both directions. The ANI/ALI controller provides ringing progress tone.
9. The ANI/ALI controller queries the database for the ALI corresponding to the ANI.
10. As the San Jose call taker answers the call, the voice, ANI, and ALI of the caller are presented.

We can see that even though the callers in examples 1 and 2 are served out of the same switch using the same trunks from the serving switch to the E9-1-1 tandem, the routing of this call is different from example 1.

Note The E9-1-1 tandem trunks are dedicated to the handling of 9-1-1 calls. No switched traffic is sent on these trunks.

Rather than cover the San Francisco customer, just note that the customer is served by a different E9-1-1 tandem, a different database (although it could be the same), and that it could *not* be routed to the San Jose or Santa Clara PSAP. This last point is important: E9-1-1 tandems are typically *not* interconnected! From a practical point of view, this means that if a 9-1-1 call is delivered to the wrong PSAP, it could not be transferred back to the appropriate PSAP if they are served by different E9-1-1 tandems. For WAN deployment of a VoIP solution, this may require that the gateways to the LEC be evaluated for their E9-1-1 tandem accessibility.

IP Telephony Emergency Call Support

IP Telephony offers two basic approaches to handling emergency calls:

- Provide automated on-net routing of emergency calls to the appropriate attendant station. This is generally accompanied by the calling party’s extension number. A typical case would be a Campus environment, where emergency calls are routed to Campus security, who then may conference-in Public Safety parties as appropriate. This allows Campus security to assist Public Safety personnel with the caller’s location details.
- Provide automated off-net routing of emergency calls to the Local Exchange Carrier’s (LEC) Point Of Presence (POP). The POP is generally a group of outgoing trunks connected to the PSTN on which IP Telephony dials the appropriate emergency call number (e.g. 9-1-1 in most North American locations). In North America, the call’s Calling Party Number (CPN) is used by the LEC to route the call to a PSAP. Once the CPN is used as part of a 9-1-1 call, it is referred to as Automatic Number Identification, or ANI. The PSAP will then use the ANI to retrieve the Automatic Location Identification (ALI) for the call.

Caution In some cases, the CPN presented to the LEC may not be the calling party’s DID number. Also, the associated entry in the Public Safety Automatic Location Identification (PSALI) database may not appropriately represent the caller’s location.

The choice between the on-net or off-net handling of emergency calls is based on policy.

To route emergency calls to an on-net location, the system designer may use IP Telephony features such as:

- Digit Translation Tables

In this case, an emergency number (such as 9-1-1) would match a CallManager translation table entry, and the corresponding on-net DN would then be used as a destination for the call (e.g., users dialing 911 could be sent to extension 5911).

- On-Net Route Pattern.

  This may be used when emergency calls are answered at an on-net location served by a remote call manager, reachable through a WAN. In this case, the number of digits presented to a remote CallManager must match the dialed digit length that the remote site uses for internal calls. An on-net route pattern could then be used to prepend digits to the dialed number for presentation to the remote CallManager (e.g., prepend “55” to “911”, effectively sending 9-1-1 calls to extension 55911). This approach also allows contingency planning for cases where the WAN resources are not available; the underlying Route List could then allow for using the PSTN as a second choice if the WAN route was not available.

Please note that the caller’s number may not be available to the called party once the call has been directed through the PSTN.

For off-net routing of emergency calls, the following IP Telephony tools may be used:

- Dial Plan Groups

  Using Dial Plan Groups, IP Telephony offers the flexibility for allowing gateway selection for LEC hand-off of emergency calls based on the calling phone’s membership in a particular calling search space. Calling search spaces contain partitions that define the reachability characteristics of PSTN gateways. This approach may prove useful in WAN IP Telephony deployment where the off-net destination of emergency calls varies depending on the caller’s location.

The following illustration presents a WAN configuration requiring two gateways to serve emergency calls.

Note Although our example is based on the North American E9-1-1 network model, the gateway selection mechanisms described here are applicable to international markets.

Remember that the San Francisco users need to use a different LEC gateway than the Oakland users.

Figure 4-34 WAN Emergency Call Configuration

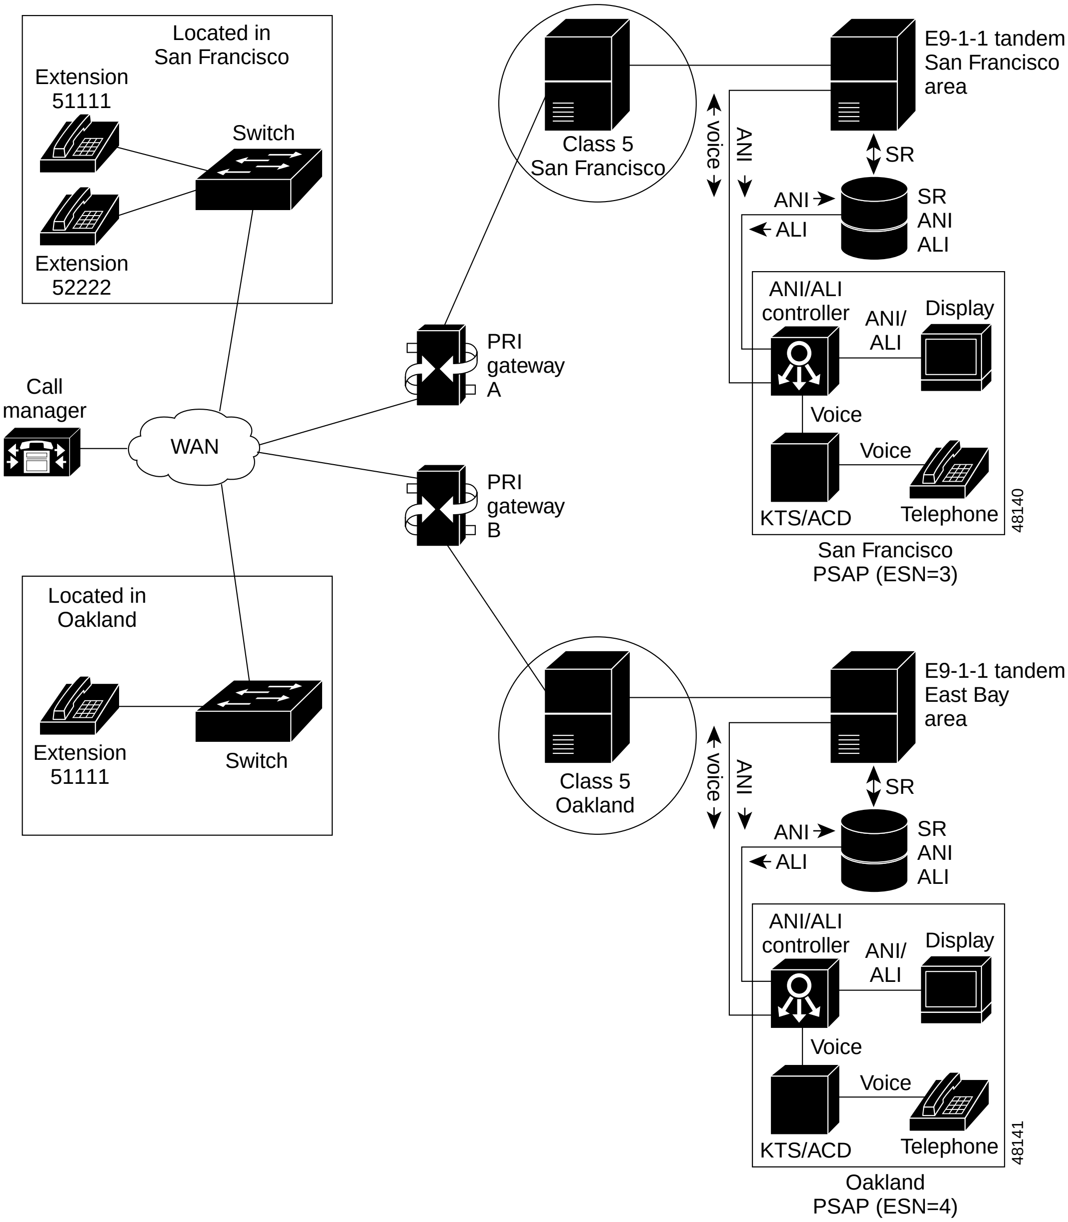

The intent is to have the following call behavior:

- All users can reach all users.
- All users use Gateway A for preferred PSTN access.
- All users use Gateway B for alternate PSTN access if Gateway A is not available.
- SF users use Gateway A for 9-1-1 access.
- Oakland users use Gateway B for 9-1-1 access.

Gateway Selection

The following discussion assumes a working knowledge of CallManager 3.0 routing. For more information regarding route patterns, route lists, route groups, partitions, and calling search spaces, refer to the *Dial Plan Architecture and Configuration* document. This document can be viewed from the following location: http://www.cisco.com/univercd/cc/td/doc/product/voice/ip_tele/network/dgdialp.htm.

The following example highlights the various configuration elements available to the IP Telephony systems architect in planning the appropriate routing of emergency calls.

Route Groups:

- SanFranciscoGW, which contains PRI gateway A.
- OaklandGW, which contains PRI gateway B.

Route Lists:

- NonEmergencyCalls, which contains route groups SanFranciscoGW, OaklandGW.
- SanFranciscoEmergency, which contains route group SanFranciscoGW.
- OaklandEmergency, which contains route group OaklandGW.

Route Filters:

- 911calls: SERVICE=911.

Route Patterns:

- Route Partition: NonEmergencyCalls.
- Route Pattern: 9.@.
- Route List: NonEmergencyCalls.
- DiscardDigits: PreDot.

- Route Partition: SanFranciscoEmergency.
- Route Pattern: 9.@.
- Route Filter: 911calls.
- Route List: SanFranciscoEmergency.
- DiscardDigits: PreDot.

- Route Partition: OaklandEmergency.
- Route Pattern: 9.@.
- Route Filter: 911calls.
- Route List: OaklandEmergency.
- DiscardDigits: PreDot.

Partitions:

- BayUsers.
- SanFranciscoEmergency.

- OaklandEmergency.
- NonEmergencyCalls.

Calling Search Spaces:

- UsersSanFrancisco, which contains BayUsers, SanFranciscoEmergency, and NonEmergencyCalls.
- UsersOakland, which contains BayUsers, OaklandEmergency, and NonEmergencyCalls.

We have assumed that users would dial 9,911. To add support for direct dialing of 911, we would need to create an extra route pattern for San Francisco and for Oakland:

- Route Partition: SanFranciscoEmergencyno9.
- Route Pattern: @.
- Route Filter: 911calls.
- Route List: SanFranciscoEmergency.

and

- Route Partition: OaklandEmergencyno9.
- Route Pattern: @.
- Route Filter: 911calls.
- Route List: OaklandEmergency.

And add them to the appropriate calling search spaces.

Calling Party Number

Once a call is routed off-net through a gateway, the CPN handed off to the CLASS 5 triggers the automated processes of delivering ANI and ALI to the PSAP.

Note LEC CLASS 5 switches may not use the CPN as the ANI presented to the E9-1-1 tandem. In some cases, the trunk group’s Listed Directory Number (LDN) may be used as the number on which ESN and ALI retrieval will be effected.

For DID phones (phones that have a fully qualified E.164 address), the DID number would typically be used as the CPN delivered to the LEC for 9-1-1 calls. This would allow the proper authorities to call the phone back directly and for a unique ALI record to be specified for each DID. As a rule of thumb, banks of DID numbers are associated with an ALI record equivalent to the Listed Directory Number (LDN) of the LEC’s trunk group. Customizing each ALI database entry associated with each individual DID number is typically a manual process that must be initiated with the LEC and/or the local Public Safety authorities.

For non-DID phones, there are two approaches that will allow control over the CPN presented to the LEC gateway upon delivery of the call:

- Set the external phone number mask on each line appropriately and then check the **Use External Phone Number Mask** on the transformations for the 911 route patterns.
- Use the **calling party transformation mask** associated with the partitions. This allows you to use the same appropriate E.164 number for a group of phones (e.g., all phones on the first floor).

Note By *appropriately* and *appropriate*, we imply using a fully qualified E.164 phone number which has associated ESN and ALI entries acceptable for that phone (or group of phones). This could mean that the associated ESN entry will route the call to the right PSAP *and* the associated ALI entry describes the phone's location within a given tolerance (e.g., 7000 ft2).

Some states have enacted legislation that poses requirements on:

- the precision of the ALI records for 9-1-1 calls coming from Enterprise Telephony systems.
- the availability of a call back number.

The are seven states where such requirements exist: Illinois, Mississippi, Tennessee, Texas, Vermont, Washington, and Kentucky. There are third party vendors that offer Enterprise Telephony systems adjunct equipment that allows for the proper handling of emergency calls.

The are seven states that have these requirements: Illinois, Mississippi, Tennessee, Texas, Vermont, Washington and Kentucky. There are third party vendors that offer Enterprise Telephony systems adjunct equipment that allows for the proper handling of emergency calls.

Warning **IP Telephony may require third party equipment to satisfy legal requirements in certain states. You may need to verify specific local requirements.**

IP Telephony configuration to accommodate the proper handling of emergency calls does not support dynamic user mobility. With each phone location change, you must properly revise the CallManager's configuration for that phone (such as applicable calling search space, gateway selection, and CPN manipulation).

TDD Support:

TDD (Telecommunications Device for the Deaf) devices are used to allow telephony communications for the speech/hearing impaired. A TDD device is a keyboard-driven modem, which allows text information to be exchanged over a standard phone line.

Most TDDs use Baudot. However, some of the newer TDD machines allow users to select either Baudot or ASCII modulation.

At this time, IP Telephony does not support TDD communications devices. If the customer requires lifeline support for TDD calls, you must have separate telco POTS lines.

Security Considerations for IP Telephony Networks

Voice security is a much more sensitive topic than data security. Users often expect that all voice communications are confidential, even when they don't have the same expectations of an E-mail containing the same information.

Many security teams spend most of their time preventing outside attackers from penetrating a corporate firewall or Internet-accessible bastion servers. However, some sources indicate that insiders commit approximately 80 percent of all corporate hacking incidents. Most companies spend little or no effort protecting the internal network infrastructure or servers from inside attack. In the context of voice communications, a prime example is an employee listening to another employee's personal or company-confidential phone calls.

You can prevent the vast majority of attempts to compromise the security of an IP Telephony network by implementing reasonable security measures on network infrastructure components and servers. Future product enhancements will address several complicated attack scenarios that these measures cannot protect against; however, those enhancements are not within the scope of this document.

This section discusses the following issues:

- Infrastructure Security Best Practices—explains general network infrastructure security that is not specific to IP Telephony solutions. Most network administrators will benefit from some of these security guidelines. However, those very experienced with security concerns may wish to skip directly to the following section.
- Securing CallManager Servers—discusses IP Telephony-specific network security issues.

Infrastructure Security Best Practices

The following sections contain general guidelines for protecting the security of your network:

- Physical Security
- Authentication, Authorization, and Accounting
- Secure Cisco IOS Devices
- Secure CatOS Devices
- Cisco IOS Security/Firewall Features
- Private Address Space and NAT

Physical Security

As with most computing devices, Cisco routers, switches, servers, and other infrastructure components are not designed to provide any protection against penetration or destruction by an attacker with direct physical access. You must take reasonable steps to prevent physical access by unauthorized personnel.

Common precautions include restricting access to wiring closets and data centers to staff that are considered *trusted*; generally, such staff already have direct or indirect logical access to the devices being protected and therefore gain no advantage from physical access. In data centers where untrusted staff may be present, use separate locking cabinets for individual items or racks which have more stringent security requirements.

When using keyed or electronic locks on doors, be sure to consider any facilities, security, and janitorial staff that may have the ability or clearance to bypass the locks.

Authentication, Authorization, and Accounting

RADIUS and TACACS+

Using either RADIUS or TACACS+ to perform per-user AAA functions is vital to enhance security and accountability for infrastructure devices. These features enable centralized administration of account and password information, eliminating per-device maintenance efforts (and errors) when trusted users are added or removed. You can also secure CiscoSecure ACS to require *strong* passwords, password aging, and intrusion lockout.

Access and Enable Passwords

When using RADIUS or TACACS+ for AAA, each user has their own password for accessing network devices. After gaining access to the device, another password (either per-user or network-wide) is necessary to *enable* commands that allow configuration changes or viewing of sensitive configuration information. We do not recommend setting users to default to level 15 (enabled) since it reduces the number of security safeguards.

Choosing Good Passwords

Ideally, you should use one-time password systems such as SoftToken, SecurID, or DES Gold Cards to prevent an attacker from reusing trusted users' passwords. However, many organizations consider these tools to be cumbersome or too expensive. If you use normal passwords, they should be chosen so that individuals cannot guess them by trying one or more dictionary words in sequence or by observing the user type. For maximum security, you should include numbers or punctuation symbols, as well as mixed case letters.

Secure Cisco IOS Devices

Restrict Virtual Console Access

Limiting virtual console access to the IP address range(s) of operations staff and network management hosts is a useful method to prevent unauthorized users from accessing network devices, even if a password is discovered.

The following example defines an access list that permits only hosts on network 192.89.55.0 to connect to the virtual console:

```
access-list 12 permit 192.89.55.0  0.0.0.255
line vty 0 4
 access-class 12 in
```

For more information, refer to the following website:
http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/12cgcr/np1_r/1rprt2/1rip.htm.

Restrict SNMP Access

Nearly all of the information viewable or configurable via a virtual console can also be accessed via SNMP. Since an SNMP community is essentially a password that does not require a user name, it is essential that you restrict this method of access as completely as possible. Only those hosts with a verified need to perform SNMP writes should have full access.

The following example defines an access list that permits only hosts on network 192.89.55.0 to perform SNMP reads with the community *foobar* and only the host 192.89.55.132 to perform SNMP writes with the community *foobaz*:

```
access-list 12 permit 192.89.55.0  0.0.0.255
snmp-server community foobar RO 12
access-list 13 permit host 192.89.55.132
snmp-server community foobaz RW 13
```

For more information, refer to the following website:
http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/12cgcr/fun_r/frprt3/frmonitr.htm#xtocid1998360.

Enable Session Timeouts

Sometimes operations staff can become distracted or be called away from their systems while logged in to network devices. Automatically disconnecting idle users helps prevent accidental access by unauthorized users.

The following example sets the real and virtual consoles to automatically disconnect the user after five minutes of inactivity:

```
line con 0
 session-timeout 5
line vty 0 4
 session-timeout 5
```

For more information, refer to the following website:
http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/12cgcr/dial_r/drprt1/drtermop.htm#4907.

Encrypt-configured Passwords

Some passwords, such as those for dialup links or local users, must be stored in the device's configuration file. Encrypting the passwords stored in the configuration file makes it difficult for a casual observer to determine or remember these passwords if they come into possession of the configuration file.

The following example enables encryption of static passwords in the configuration file:

```
service password-encryption
```

For more information, refer to the following website:
http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/12cgcr/secur_r/srprt5/srpass.htm#4899.

Disable minor host services

By default, several services are enabled which either allow an attacker to more easily consume device resources, indirectly attack other hosts, or gain information about operations staff currently accessing the network devices. You can disable these services to prevent malicious use of these services or the information they may provide.

The following example disables these services:

```
no service tcp-small-servers
no service udp-small-servers
no service finger
```

For more information, refer to the following website:
http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/12cgcr/fun_r/frprt3/frgenral.htm.

Disable or Restrict the HTTP Server

Web configuration is disabled on most platforms; however, novice network administrators often enable it. If HTTP configuration is not necessary, you should disable it entirely. If disabling the service is not feasible, restrict HTTP access to management addresses.

The following example defines an access list that permits only hosts on network 192.89.55.0 to connect to the HTTP server:

```
access-list 12 permit 192.89.55.0  0.0.0.255
ip http access-class 12
```

For more information, refer to the following website:
http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/12cgcr/fun_r/frprt1/frui.htm.

Disable Forwarding of Directed Broadcasts

Directed broadcasts are unicast packets that are addressed to another subnet's broadcast address. While forwarding these packets has a limited diagnostic value, there is a significant risk in becoming an amplifier in various types of Denial of Service attacks. Cisco IOS versions 12.0 and later disable directed broadcasts by default, but they should be manually disabled on all prior versions.

The following example disables directed broadcasts on interfaces Ethernet0/0 and Serial1/1:

```
interface Ethernet 0/0
 no ip directed-broadcast
interface Serial 1/1
 no ip directed-broadcast
```

For more information, refer to the following website:
http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/12cgcr/np1_r/1rprt2/1ripadr.htm.

Disable Forwarding of Source-routed Packets

Source-routed packets contain additional hop-by-hop routing information that can supersede what is present in routing tables. Although it was initially intended as a diagnostic tool for network operators, it is not very valuable and can be used to exploit security vulnerabilities.You should disable it on all routers.

The following example disables forwarding of packets containing source-route information:

```
no ip source-route
```

For more information, refer to the following website:
http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/12cgcr/np1_r/1rprt2/1rip.htm.

Disable RCP and RSH Services

Use the Berkeley Remote Copy (RCP) command to copy files to a device and the Remote Shell (RSH) command to execute commands without logging in. However, be aware that these services have extremely weak authentication and should not be enabled unless no other option (such as SSH support in Cisco IOS version 12.1T) is available.

The following example disables RCP and RSH services:

```
no ip rcmd rcp-enable
no ip rcmd rsh-enable
```

For more information, refer to the following website:
http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/12cgcr/fun_r/frprt2/fraddfun.htm.

Enable Neighbor Authentication

Most common networking protocols provide a means for neighbors to authenticate each other to ensure that unauthorized devices are not allowed to affect the stability or security of the network. These authentication mechanisms prevent casual attempts at disrupting proper operation, but should not be expected to stop a determined attacker.

HSRP

The following example enables the authentication string *foobar* for HSRP group 21 on interface Ethernet 2/1:

```
interface Ethernet 2/1
 standby 21 authentication foobar
```

For more information, refer to the following website:
http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/12cgcr/np1_c/1cprt2/1cip.htm.

Enhanced IGRP

The following example enables the authentication string *foobar* for EIGRP AS 1 on interface Ethernet 2/1:

```
key chain baz
 key 1
  key-string foobar
interface Ethernet 2/1
 ip authentication mode eigrp 1 md5
 ip authentication key-chain eigrp 1 baz
```

For more information, refer to the following website:
http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/12cgcr/np1_c/1cprt1/1ceigrp.htm.

OSPF

The following example enables the authentication string *foobar* for OSPF Area 2 on interface Ethernet 2/1:

```
interface Ethernet 2/1
 ip ospf message-digest-key 1 md5 foobar
router ospf 1
 area 2 authentication message-digest
```

For more information, refer to the following website:
http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/12cgcr/np1_c/1cprt1/1cospf.htm.

IS-IS

The following example enables the authentication strings *foobar* for Level 1 routes and “foobaz” for Level 2 routes:

```
router isis
 area-password foobar
 domain-password foobaz
```

For more information, refer to the following website:
http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/12cgcr/np1_c/1cprt1/1cisis.htm.

BGP

The following example enables the TCP MD5 authentication string *foobar* for the connection to neighbor 192.89.55.12:

```
router bgp 1
 neighbor 192.89.55.12 password foobar
```

For more information, refer to the following website:
http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/12cgcr/np1_c/1cprt1/1cbgp.htm.

PPP (leased lines)

PPP CHAP authentication requires neighboring routers to verify each others' identity over a link before the link can be used for traffic. PPP is typically used over leased lines; however, recent Cisco IOS features allow PPP to be used over ATM and Frame Relay links as well.

The following example enables the PPP CHAP authentication string *foobar* for routers SJ-WAN and DALLAS-WAN connected via a leased line:

```
hostname dallas-wan
username sj-wan password foobar
interface Serial 2/1
 encapsulation ppp
 ppp authentication chap

hostname sj-wan
username dallas-wan password foobar
interface Serial 4/3
 encapsulation ppp
 ppp authentication chap
```

For more information, refer to the following website: http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/12cgcr/dial_c/dcppp.htm.

Configure Accurate Time Stamping

Many troubleshooting tasks, such as determining the nature of a Denial of Service attack or tracing attempts to pass through firewalls, involve correlating logs from several devices. Unless the clocks of these devices are synchronized, it is much more difficult to correlate different logs.

Note The details of enterprise-wide NTP design and upstream time sources are beyond the scope of this document.

The following example enables millisecond-precision time stamps in log and debug messages and configures an upstream NTP server at 192.89.55.132:

```
service timestamps debug datetime msec
service timestamps log datetime msec
ntp server 192.89.55.132
```

Refer to the following documents for more information:

- Troubleshooting Commands

  http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/12cgcr/fun_r/frprt3/frtroubl.htm
- Performing Basic System Management

  http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/12cgcr/fun_c/fcprt3/fcgenral.htm

Syslog server

Logging all system notices and error messages often provides valuable insight into the operational status of network devices. If access list violations are logged, the logs may also be correlated between devices to determine that the network is being probed or that a device has been compromised.

The following example enables logging to a syslog server at 192.89.55.132:

```
logging 192.89.55.132
```

Refer tothe following document for more information:
http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/12cgcr/fun_r/frprt3/frtroubl.htm.

IP Accounting

The IP accounting feature provides basic IP traffic statistics functions. By enabling IP accounting, users can see the number of bytes and packets switched through the Cisco IOS software based on a source and destination IP address.

The following example enables IP Accounting on interface Ethernet 2/1:

```
interface Ethernet 2/1
 ip accounting
```

Refer to the following document for more information:
http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/12cgcr/np1_c/1cprt2/1cip.htm.

Secure CatOS Devices

Restrict virtual console and SNMP access

You can limit virtual console and SNMP access to the IP address range(s) of operations staff and network management to prevent unauthorized users from accessing network devices, even if a password is discovered.

The following example defines an access list that permits only hosts on network 192.89.55.0 to connect to the virtual console or make SNMP requests:

```
set ip permit 192.89.55.0 255.255.255.0 all
```

For more information, refer to the following website:
http://www.cisco.com/univercd/cc/td/doc/product/lan/cat6000/sw_5_5/cnfg_gd/ip_perm.htm.

Configure SNMP communities

Nearly all of the information that you can view or configure via a virtual console can also be accessed via SNMP. Since an SNMP community is essentially a password that does not require a user name, it is essential that this method of access is restricted as completely as possible. You should only allow access to those hosts with a verified need to perform SNMP writes.

The following defines RO, RW, and RWA communities of *foo*, *bar*, and *baz*, respectively:

```
set snmp community read-only foo
set snmp community read-write bar
set snmp community read-write-all baz
```

For more information, refer to the following website:
http://www.cisco.com/univercd/cc/td/doc/product/lan/cat6000/sw_5_5/cnfg_gd/snmp.htm.

Enable session Timeouts

Sometimes operations staff become distracted or are called away from their systems while logged in to network devices. Automatically disconnecting idle users helps prevent accidental access by unauthorized users.

The following example sets the switch to automatically disconnect the user after 5 minutes of inactivity:

```
set logout 5
```

For more information, refer to the following website:
http://www.cisco.com/univercd/cc/td/doc/product/lan/cat6000/sw_5_5/cmd_refr/set_m_pi.htm.

Enable Port Security

Many infrastructure attacks require that the attacker assume the MAC address of the victim's device or replace the victim's device with a counterfeit clone. Enabling port security on Catalyst switches will automatically cause the switch to disable any port which changes MAC addresses or uses another port's MAC address. You should use this feature with caution in environments where laptops or other devices are commonly connected to different switch ports.

The following example enables port security on ports 2/1 through 2/48 and configures the disablement to last 30 minutes after each violation:

```
set port security 2/1-48 enable
set port security 2/1-48 shutdown 30
```

For more information, refer to the following website:
http://www.cisco.com/univercd/cc/td/doc/product/lan/cat6000/sw_5_5/cnfg_gd/sec_port.htm.

Enable VTP Server Authentication

VTP servers are capable of adding or removing VLANs within a VTP domain. Since you can potentially configure any Catalyst switch as a VTP server for a given domain, setting a password on the intended VTP servers is the only way to prevent an attacker from sending false VTP commands.

The following example enables the password *foobar* for the current VTP domain:

```
set vtp passwd foobar
```

For more information, refer to the following website:
http://www.cisco.com/univercd/cc/td/doc/product/lan/cat6000/sw_5_5/cnfg_gd/vtp.htm.

Cisco IOS Security/Firewall Features

Anti-spoofing Filters

Many widespread Denial Of Service attacks rely on the ability of the attacker to send packets with forged (spoofed) source addresses, which makes tracking the true source of the attack very difficult. To help prevent your site from sourcing these types of attacks, you should block any outbound packets outside of your own address space.

The following example prevents any outbound packets that do not come from the site's address block of 192.168.123.0/24:

```
access-list 101 permit ip 192.168.123.0 0.0.0.255 any
access-list 101 deny ip any any
interface Serial 2/1
 ip access-group 101 out
```

For more information, refer to the following website:
http://www.cisco.com/warp/public/707/21.html#spoofing.

TCP Intercept

TCP Intercept is a Cisco IOS feature that is specifically designed to protect vulnerable hosts from SYN Attacks. These attacks abuse a common flaw in TCP implementations to render the host temporarily unable to accept incoming connections.

The following example enables TCP Intercept for all connections going to any host on network 192.168.1.0:

```
access-list 101 permit tcp any 192.168.1.0 0.0.0.255
ip tcp intercept list 101
```

For more information, refer to the following website: http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/12cgcr/secur_c/scprt3/scdenial.htm.

Context-Based Access Control (CBAC)

Cisco's Context-Based Access Control (CBAC) extends the functionality of a traditional access list to include firewall functionality by monitoring stateful connections.

For more information, refer to the following website: http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/12cgcr/secur_c/scprt3/sccbac.htm.

Intrusion detection

The Cisco Intrusion Detection System (IDS) is the industry's first real-time network intrusion detection system that protects the network perimeter, extranets, and the increasingly vulnerable internal network. The system uses sensors, which are high-speed network appliances, to analyze individual packets to detect suspicious activity. If the data stream in a network exhibits unauthorized activity or a network attack, the sensors detect the misuse in real time, forward alarms to an administrator, and remove the offender from the network. Many Cisco routers, switches, and the Cisco Secure PIX Firewall also have IDS functionality.

Please contact your Cisco account team for additional information about Cisco's Intrusion Detection System products and how you can apply them to your specific environment.

Private Address Space and NAT

Using RFC 1918 private address space (10/8, 172.16/12, and 192.168/16) for the internal parts of a corporate network is often considered a security measure since these addresses cannot be directly reached from the public Internet.

While using private addressing is useful in many instances for other purposes, it is not ideal. Any penetration from the public Internet is likely to first step through one of the company's exposed Internet servers that can usually reach private addresses behind a NAT or proxy device.

Also, since most attacks come from inside attackers, the ability to hide devices from attackers on the Internet does not pose any protection against the most significant threat.

Securing CallManager Servers

A CallManager server is built from several distinct components (Windows 2000 Server, SQL Server, and Internet Information Server), and security for each of these components must be addressed separately.

Microsoft Windows 2000 Tasks

Install all available security patches

Microsoft operating systems, particularly Windows NT and 2000, are subject to intense scrutiny by the hacking community. When a new vulnerability is discovered, Microsoft responds to security vulnerabilities with patches, usually within days. Due to the frequency at which bugs are discovered, it is essential that an administrator make sure that all possible security patches are applied to all Microsoft products in a timely fashion.

Stop Unnecessary Service

One of the fundamental principles of securing a server is that each service running will expose potential security vulnerabilities. Since securing every service is difficult and time-consuming, it is a logical task to disable all services that are not mandatory, even if those services aren't immediately known to have a hole.

Unless otherwise needed on the system, all of the following services should be stopped and set to Manual Start status:

- Alerter Service
- Computer Browser
- Distributed File System
- DHCP Client
- DHCP Server
- DNS server
- FTP Server
- Fax Service
- License Logging Service
- Net Meeting Remote Desktop Sharing

Disk Security

FAT-style file systems do not provide any inherent means to restrict access to specific users or encrypt entire disks. Since unauthorized users being able to see or alter privileged information could compromise the security of a CallManager, you must use NTFS. This is the default for new installations of CallManager.

Sanitize Accounts

There is no logical reason for unknown users to be logging into the CallManager system. Disable the **Guest** account and remove any users from the **Guests** group. This is done via the Users and Passwords control panel.

Many attacks rely on blindly trying to execute commands as Administrator; changing the name of the Administrator account to **CallmgrAdmin** or some other logical name will prevent these attacks from working properly even if the system is compromised. This is done via the Security Options in the Local Security Policy.

Many attacks require access as a non-privileged user on the system. These attacks can be prevented by only allowing Administrators to log on to the CallManager locally. This is done via User Rights Assignments in the Local Security Policy.

Since Administrator accounts are privileged, it is essential to ensure that unauthorized users cannot guess the passwords for these accounts. Requiring that all local passwords on the CallManager to be at least eight characters long prevents casual observation of the password and an automatic lockout after five incorrect attempts will prevent an attacker from finding a correct password by entering random guesses. These tasks can be done via Local Security Settings in the Local Security Policy.

System Auditing

Auditing allows usage tracking for many privileged tasks in Windows 2000. When auditing is enabled, regularly reviewing the Event Viewer may help determine if the system has been compromised.

This is the suggested Auditing Scheme:

Table 4-18 Auditing Scheme

Description	Log Success	Log Failure
Audit Account logon events	Yes	Yes
Audit account management	Yes	Yes
Audit directory service access	Yes	Yes
Audit logon events	Yes	Yes
Audit object access	No	Yes
Audit Policy Change	Yes	Yes
Audit privilege use	Yes	Yes
Audit process tracking	No	Yes
Audit system events	Yes	Yes

Microsoft Internet Information Server (IIS) Tasks

The Microsoft Internet Information Server has been plagued with security holes. At least one to two security bulletins are issued for IIS every month.

The following Security Bulletins represent the most recent security patches that Microsoft has released to harden the IIS product.

- Microsoft Security Bulletin (MS00-018)
  Chunked encoding transfers
- Microsoft Security Bulletin (MS00-019)
  Patch Available for “Virtualized UNC Share” Vulnerability
- Microsoft Security Bulletin (MS00-023)
  Patch Available for “Myriad Escaped Characters” Vulnerability
- Microsoft Security Bulletin (MS00-030)
  Patch Available for “Malformed Extension Data in URL” Vulnerability
- Microsoft Security Bulletin (MS00-031)
  Patch Available for “Undelimited .HTR Request” and “File Fragment Reading via .HTR” Vulnerabilities

The frequency and quantity of the security patches illustrate the importance of applying security patches on a regular basis.

Securing Microsoft IIS

The reference material on IIS recommends the following actions to maximize the security of the IIS Application:

Note You can add any recommended registry changes to the SecEdit script so that a single script can tighten the security on the server.

Step 1 Disable unnecessary services.

Required services:

- Event Log
- License Logging Service
- Windows NTLM Security Support Provider
- Remote Procedure Call (RPC) Service
- Windows NT Server or Windows NT Workstation
- IIS Admin Service
- MSDTC
- World Wide Web Publishing Service
- Protected Storage
- Server
- Workstation

May Be Required:

- FTP Publishing Service (required for FTP Service)
- NNTP Service (required for NNTP News service)
- SMTP Service (required for SMTP service)
- Content Index (required if using Index Server)
- RPC Locator (required if performing remote administration)
- Server Service (can be stopped, but will have re-start if you need User Manager)
- Telephony Service (required if access is via dialup)
- Remote Access Service (required if dialup access is used)

- Certificate Authority (required if you plan on issuing certificates)
- Plug & Play (recommended but not required)
- Workstation (optional; important if UNC virtual roots are used)
- UPS (optional; it is recommended to us a UPS)

Not Required by Most Installations:

- Alerter
- Clipbook Server
- Computer Browser
- DHCP Client
- Messenger
- Net Logon
- Network DDE & Network DDE DSDM
- Network Monitor Agent
- Simple TCP/IP Services
- Spooler
- Netbios Interface
- TCP/IP NetBIOS Helper
- WINS Client (TCP/IP)
- NWLink Netbios
- NWLink IIPX/SPX Compatible transport (Not required if you have TCP/IP)

Step 2 Disable RDS support.

The RDS Datafactory (a single component of RDS) allows implicit remoting of data access requests by default. Therefore, it can be exploited to allow unauthorized Internet clients to access OLE DB data sources available to the server. To accomplish this, remove the following registry keys and any subkeys:

```
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\W3SVC\Parameters\ADCLaunch\RDSServer.
DataFactory
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\W3SVC\Parameters\ADCLaunch\AdvancedDa
taFactory
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\W3SVC\Parameters\ADCLaunch\VbBusObj.V
bBusObjCls
```

Step 3 Enable W3C Extended Logging Format.

The default logging mechanism does not record enough information to help determine whether a server is under attack.

Step 4 Clear Indexing.

By indexing source code, it is possible for an attacker to view the content of the web pages. To clear indexing, follow these steps:

a. Start the IIS Microsoft Management Console (MMC) and go to the Web Site Properties by right-clicking on the web site entry and selecting **Properties**.

b. Select the **Home Directory** tab.

c. Clear the **Index this Directory** and the **Directory browsing allowed** options.

Step 5 Disable Parent Paths.

This prevents the use of the "." in calls to MapPath. This option is enabled by default, but you should disable it. Follow these steps:

a. Start the IIS MMC and go to the Web Site Properties by right-clicking on the Web Site entry and selecting **Properties**.

b. Select the **Home Directory** tab.

c. Select the **Configuration** tab.

d. Select **App Options**.

e. Clear the **Enable Parent Paths** option.

Step 6 Remove unused script mappings.

IIS is pre-configured to support various common filename extensions such as .ASP, .SHTML, and .HTR. Processing of these requests are handled by various DLLs located on the system. By removing the mappings to extensions that are not used, you minimize the potential attack points. Follow these steps:

a. Start the IIS MMC and go to the Web Site Properties by right-clicking on the Web Site entry and selecting **Properties**.

b. Select the **Home Directory** tab.

c. Select the **Configuration** tab.

d. Select the **App Mappings** tab.

e. Remove the necessary mappings.

Step 7 Remove IIS virtual directories.

IIS contains several virtual directories that need to be removed. They are:

- IISAMPWD
- IISSAMPLES
- IISADMIN
- IISHELP

Step 8 Remove all sample application directories.

The following directories contain sample files that you should remove from the system. This will prevent an attacker from exploiting a vulnerability in one of the sample files to gain access to the system:

- \Inetpub\iisamples
- \Inetpub\scripts\samples
- \Inetpub\wwroot\samples
- \Program Files\Common Files\System\msadc\Samples
- \WINNT\system32\inetsrv\adminsamples
- \WINNT\system32\inetsrv\iisadmin
- \WINNT\system32\inetsrv\iisadminpwd

Step 9 Set appropriate virtual directory permissions/web application space.

It is important to ensure you apply the correct permissions to the files available on the web server. These permissions vary depending on the type of files being accessed. The following table provides a rough guideline to follow:

File Type	ACL
CGI and related files .EXE, .DLL, .CMD, .PL	Everyone (Execute) Administrators and System (Full Control)
Script files .ASP etc.	Everyone (Execute) Administrators and System (Full Control)
Include files .INC, .SHTML, .SHTM	Everyone (Execute) Administrators and System (Full Control)
Static Content .HTML, .GIF, .JPEG	Everyone (Execute) Administrators and System (Full Control)

Step 10 Set appropriate IIS Log file ACLs.

To prevent malicious users from deleting log files to cover their activities, the file permissions on the IIS generated log files (%systemroot%\system32\LogFiles) should be as follows:

File Type	ACL
.LOG files	Everyone (Execute) Administrators and System (Full Control)

Step 11 Install Microsoft MDAC 2.1.2.4202.3.

On web sites that have both IIS and certain versions of MDAC, a visitor could perform privileged actions on the system. You can remove this vulnerability by installing MDAC 2.1 and configuring it to operate in Safe Mode. Change the following registry key:

- **Hive**: HKEY_LOACL_MACHINE\SOFTWARE
- **Key**: \Microsoft\DataFactory\HandlerInfo
- **Name**: HandlerRequired
- **Type**: REG_DWORD

A value of 0 = unsafe mode and 1 = safe mode.

Microsoft SQL Server Tasks

Default Configuration Settings

STAT performed a cursory analysis of the security settings using the default database installation instructions. The SQL database comprises the heart of the CCM. Information on devices, configuration, and call data are all stored in this database and corruption of the database can disrupt the entire operation of the CCM cluster. This analysis was intended to discover any glaring problems with the database configuration. A more detailed investigation is required to thoroughly examine the security of the SQL configuration, including data replication and remote database access.

Our examination uncovered the following security concerns:

- Default access is set to Mixed Mode.
- The audit level is set to none.

- The group **Everyone** has full access to the database file.

The following sections explain these concerns.

Default Access is Mixed Mode

There are two access modes Microsoft SQL 7.0 supports: **Mixed** mode and **Windows NT Only** mode. When specifying the Windows NT Only mode, the operating system handles all authentication. This implies that the NT challenge-response mechanism is used during the logon process. Mixed mode, on the other hand, allows both OS authentication and SQL Server authentication. With SQL server authentication, logins are authenticated against a table of username/password combinations stored by the SQL server.

Audit Level is Set to None

Microsoft SQL provides the capability to track logons to the database server. This information can be extremely useful when attempting to monitor for attacks against the database, especially with respect to failed logon attempts. SQL provides the following levels of auditing:

- None—logs no auditing information.
- Success—logs only successful logins.
- Failure—logs only failed logins.
- All—logs successful and failed logins.

At a minimum, you should audit failed logon attempts.

Everyone Has Full Access to the Database File

A Windows 2000 server contains a default group called **Everyone**. This group represents any valid logon to the system, including default logons such as guest and anonymous. By providing everyone with full access to the database file, anyone who has logon access to the Windows 2000 server can read, alter, or delete the information stored in the database. Since this database represents the heart of the CCM cluster, this is an unacceptable condition. You should restrict full access to the database file to administrators. Furthermore, you should provide any access to the database on an as-needed basis.

Setting Up a Secure SQL Server 7.0 Installation

The section applies to SQL Server 7.0 installed on Windows 2000 only since the Windows 95 and Windows 98 environments do not provide these security features.

Note The following section assumes that you configured SQL Server 7.0 has with Windows 2000 Authentication Mode to provide the highest level of security.

Step 1 Do Not Use the sysadmin (sa) Account

We recommend that all SQL Server administrators be granted access to SQL Server through Windows 2000 group membership and that you make this same group a member of the sysadmin server role. This approach has one minor drawback. Windows 2000 administrators can give anyone sysadmin privileges on SQL Server 7.0 since they can add any user to the appropriate Windows 2000 group.

If a site does not want to give Windows 2000 administrators the ability to give others (or themselves) sysadmin access to SQL Server, only individual Windows 2000 accounts should be assigned to the role of sysadmin.

In each case, we strongly recommend that you not use the sa account for day-to-day administration, but rather that you assign a password. You should then lock the password in a safe for emergency access only.

If you are running SQL Server 7.0 with Windows 2000 Authentication Mode (as recommended in this document) you cannot log on using the sa account, which allows trusted connections only .

Note Even though you cannot use the sa account to log in to SQL Server 7.0 when it is running in Authentication Mode, it is still important to assign an sa password. This small change in the registry can change the security mode from Authentication Mode to Mixed Mode. The registry key where the login mode is configured is:
HKLM/Software/microsoft/MSSQLServer/MSSQLServer/LoginMode
If the value is 0, then Mixed Mode has been configured. If it is a 1, then Windows 2000 Authentication Mode has been selected. If the sa password is blank (as in a default installation), an intruder (or the Windows 2000 Administrator) could gain access to the server.

Step 2 Set up the Service Accounts.

SQL Server 7.0 runs as three Windows 2000 services:

- MSSQLServer—the engine that provides the core functionality of SQL Server.
- SQLServerAgent—provides the capability to schedule regular commands and replication, provide a method for dealing with errors, and contact SQL Server operators when errors occur, in addition to other support functions.
- Microsoft Search—provides the full-text search capability and must always be configured to use the local system account.

You can configure the MSSQLServer and SQLServerAgent services to use one of the following types of Windows 2000 accounts:

- Local service account
- Local user account
- Domain user account

The selection depends on the functionality that is required for SQL Server 7.0. You can configure both services to use the same Windows 2000 user account. If you need to change the service account after the server has been installed, you should use the SQL Server Enterprise Manager. While it is also possible to change the service account for the MSSQLServer and SQLServerAgent services in Control Panel, we do not recommend this because the configuration details for the Microsoft Search service are not synchronized.

The changes to account information take effect the next time the service is started. You can configure the MSSQLServer and SQLServerAgent services to use different Windows 2000 user accounts, although this is not usually recommended. When changing the service account, the changes must be made to both services since they are configured separately. One consideration that can reduce administrative overhead in a multi-server environment is the use of one domain user account for all SQL Server 7.0 servers in the enterprise.

Local System Account

You can run SQL Server 7.0 using the local system account if the SQL Server is not configured for replication and does not require access to network resources.

You must set the following permissions for the local system account for SQL Server 7.0 to properly perform its tasks:

- Full Control on the SQL Server directory (by default \MSSQL7)
- Full Control on all .mdf, .ndf, and .ldf database files
- Full Control on the registry keys at and under:

```
HKEY_LOCAL_MACHINE\Software\Microsoft\MSSQLServer
HKEY_LOCAL_MACHINE\System\CurrentControlset\Services\M
SSQLServer
```

Local User Account

If SQL Server 7.0 is configured to use a Windows 2000 local user account, the same restrictions apply as the local system with the following addition: the user account must be granted the **Log On As A Service** right.

Domain User Account

Configuring SQL Server 7.0 with a domain user account provides the greatest level of flexibility. The following are some examples of functionality available when you only use a domain user account:

- Replication
- Backing up to and restoring from network drives
- Performing heterogeneous joins that involve remote data sources
- SQL Server Agent mail features and SQL Mail

For SQL Server 7.0 to perform its tasks, the domain user account must be set up with the same configuration as the local user account discussed earlier. However, some extended functionality is available only if further permissions are considered. See the following table:

Table 4-19 SQL Server 7.0 Extended Functionality

Service	Permission	Functionality
MSSQLServer	Network write permissions	Ability to read/write to remote backups, data loads, and so on.
MSSQLServer	Act as part of the operating system and replace process level token.	Run xp_cmdshell for a user other than a SQL Server administrator.
SQLServerAgent	Member of the administrators local group.	Create CMDExec and ActiveScript jobs belonging to someone other than a SQL Server administrator.
SQLServerAgent	Member of the Administrators local group.	Use the autorestart feature.
SQLServerAgent	Member of the Administrators local group	Use run-when-idle jobs.

To provide maximum functionality to SQL Server 7.0, we recommend that the domain user account be a member of the Administrators local group.

Step 3 Set the File Permissions.

Windows 2000 provides an excellent security framework for securing operating system objects such as files. Microsoft recommends that you apply NT file system (NTFS) file permissions to the data and log files of all databases. The user account to which SQL Server 7.0 is configured must be given Full Control permissions on the database files.

Further, you should configure all SQL Server 7.0 files, including executables and dynamic link libraries (DLLs), so that users cannot manipulate them. You should set permissions on these files to allow the user account that SQL Server uses, the Administrators group and local system accounts, Full Control permissions. You should not set any other permissions.

Step 4 Set the Registry Permissions.

To secure the SQL Server 7.0 installation from attacks on the physical server by users who have logon rights, you should set Windows 2000 permissions on the registry keys that are used to configure SQL Server 7.0. Specifically, you should secure all the keys under the following registry key:

```
HKEY_LOCAL_MACHINE\SOFTWARE\MICROSOFT\MICROSOFT SQL SERVER 7.0\
```

You should remove the Everyone group permissions on this key and add Full Control permissions for the Administrators group, the local system account, or the SQL Server service account.

Setting permissions on the registry keys is particularly important if the SQL Server administrators want to stop the Windows 2000 administrators from accessing the SQL Server. In this case, the SQL Server administrators should also take ownership of the registry key and remove permissions from the Administrators group. It is then imperative that the SQL Server service account has Full Control permissions. Although this does not stop administrators from gaining access, it allows SQL Server administrators to know when the Windows 2000 administrators have compromised security. Administrators can always take ownership, but they cannot give it.

Step 5 Set the Audit Level.

When SQL Server 7.0 uses Windows 2000 Authentication Mode, it provides the capability to audit logons to the server. You can configure the audit level using SQL Server Enterprise Manager or by using the sp_loginconfig stored procedure. Possible auditing settings are:

- None—logs no auditing information.
- Success—logs only successful logins.
- Failure—logs only failed logins.
- All—logs successful and failed logins.

The auditing information is written to the SQL Server 7.0 error log.

Step 6 Use the Profiler for Auditing.

SQL Server 7.0 provides a very powerful profiler, which allows the analysis of many internal events that occur within SQL Server.

SQL Server Profiler works by capturing all the actions performed on the SQL Server and sending them to the SQL Server Profiler where they can be analyzed. The capture can be viewed in realtime on the screen, saved to a text file, or inserted into a SQL Server table.

The SQL Server Profiler allows capturing of virtually all events that take place within SQL Server, including:

- Login Failed
- Locking: Deadlock
- Object: Closed
- Stored Procedure: Statement Starting
- Session: Disconnect

- RPC: Completed

This information can provide excellent support to establish event time and origin.

Step 7 Secure the Backup Files and Media.

The most secure method for backups is to use SQL Server 7.0 to back up to data files and then to use the Windows 2000 backup program to back the data files up to backup media using the password feature. This ensures that only those who know the password can restore the files. The backup data files should be on an NTFS partition with directory permissions set to prevent the ordinary user from gaining access to the files.

If backup media can be physically secured, the standard SQL Server 7.0 backups will not pose any security risks.

Step 8 Restore the Database to Another Server.

There are three situations related to restoring the database to another server:

- Mixed Mode
- Windows NT Authentication (same domain)
- Windows NT Authentication (different domain)

Mixed Mode

When restoring a database to a server using Mixed Mode for security authentication, the database security breaks. This is because the logons are maintained in the sysxlogins table in the master database, while the user's rights to access a database are stored in the sysusers table of the respective database; a logical link is maintained between the user's entry in the sysxlogins table and the user's entry in the sysusers table. This link is a generated 16-byte GUID.

The net effect of the GUID implementation for Mixed Mode authentication is that when a database is restored to a computer running SQL Server 7.0, other than the one where the database access is granted, the link between the sysxlogins table and the sysusers table breaks, thereby effectively granting access to the database to no one. Members of the sysadmin group are an exception to this. You would have to re-create all role memberships and user permissions.

Restoring a database to another server exposes one of the weaknesses of Mixed Mode authentication.

Windows NT Authentication (Same Domain)

If the database is restored to another computer running SQL Server 7.0 in the same domain, the permissions in the database remain intact. The only consideration here is whether users are granted permission to log on to the server. The permission to log on to the server is implemented at each instance of SQL Server.

For example, Bob is a member of the Sales group and the Sales group is granted login permissions at SQLSERVER1. Bob is granted database access rights to the sales database. When the sales database is restored to SQLSERVER2, Bob's permissions still exist in the sales database. However, because the SALES group is not granted login rights to the server, Bob cannot use the database. If the administrator grants the Everyone group login rights to the server, Bob can use the database. This is because the only restriction stopping Bob from using the sales database was logging in to the server.

When restoring a database to another server in the same domain, the permissions within the database remain intact, but the permissions to log in to this specific server may need to be granted.

- Users from a Trusted Domain

  If a Windows 2000 trust relationship has been established between the old and the new domain such that the new domain trusts the old domain, the users from the old domain may use the database with all permissions intact, provided that they have been granted the right to log in to SQL Server.

Users from other trusted domains would not have rights to access the database, much like the users from the new domain.

- Users from the New Domain

  None of the users from the new domain will have access because their SIDs do not exist in the sysusers table of the database.

  The only exceptions to this are the BUILTIN accounts of Windows 2000. Since these accounts always have the same SIDs on all servers, any permissions granted to a BUILTIN account, such as the local Administrators group, remain intact. This assumes that the BUILTIN accounts have logon rights, and that SQL Server is installed on a domain controller.

- Users from Any Domain with Same Username and Password

  In most Windows 2000 security implementations, when access is required to a resource that is not in the user's own domain, the user is able to access the resource providing that a user account exists with the same username/password combination. This behavior is transparent.

  Provided that the user is using named pipes to connect to the server, this will work if the user establishes a connection to a file share first. This method also works if a user wants to use an account of another name, providing that he/she is running Windows 2000 as the computer operating system. If a user is denied access when connecting to a file resource from a computer running the Windows 2000 operating system (and the user is not currently using any other credentials on the computer being connected to), the opportunity is given to provide a username and password for logon purposes.

Step 9 Attach/Detach Database Files.

The issues associated with attaching and detaching database files are identical to those discussed in Step 8. An exception is the requirement to create the database before restoring the data.

Step 10 Disable Windows 2000 Guest Account

When running SQL Server 7.0 in Windows 2000 Authentication Mode, the server relies on Windows 2000 to perform all authentication of clients. This brings with it the security framework that applies to Windows 2000, both the strengths and the weaknesses. Fortunately, there are not many of the latter. However, one issue that has been adequately documented in many Microsoft and third-party security papers is the use of the Windows 2000 Guest account. It is strongly recommended that the Guest account be disabled, if this has not already been done.

Integrating Voice Mail

This section discusses IP Telephony voice mail integration, including:

- Voice Messaging with Cisco uOne 4.1E
- Integrating SMDI Voice Mail
- Integrating SMDI Voice Mail Over IP WAN

Voice Messaging with Cisco uOne 4.1E

http://www.cisco.com/univercd/cc/td/doc/product/voice/uone/index.htm

Integrating SMDI Voice Mail

With SMDI integration, call information is transmitted over an RS-232 serial link between CallManager and the voice mail system. Voice communications are provided through a separate path created by a hunt group. This group consists of analog stations (FXS ports derived from either the VG200 gateway(s) or Catalyst 6000 24-port FXS Gateway(s) WS-X6624) between CallManager and the voice mail system. When the hunt group receives an incoming call, it is accompanied by a digital message in standard SMDI format from CallManager. This digital message contains all of the necessary call information. The voice mail system then answers the call on the specified port and plays the appropriate greeting. To set or cancel message-waiting notification, the voice mail system sends a digital message over the RS-232 serial link to CallManager.

The following sections are based on this topology:

Figure 4-35 Voice Mail Topology

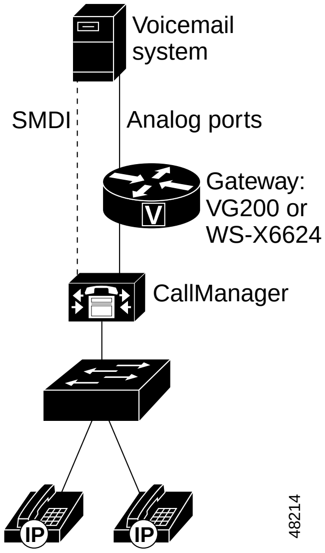

Configuration Overview

Follow these steps to configure the SMDI voice mail system:

Step 1 Verify Cisco Messaging Interface (CMI) is installed.

You enable SMDI through the CMI application. This application is *not* part of the default installation so you must manually select it under **Optional Components** if SMDI is required. Once installed, you can select **Cisco Messaging Interface** from the **Service** menu.

Figure 4-36 Cisco Messaging Interface Command

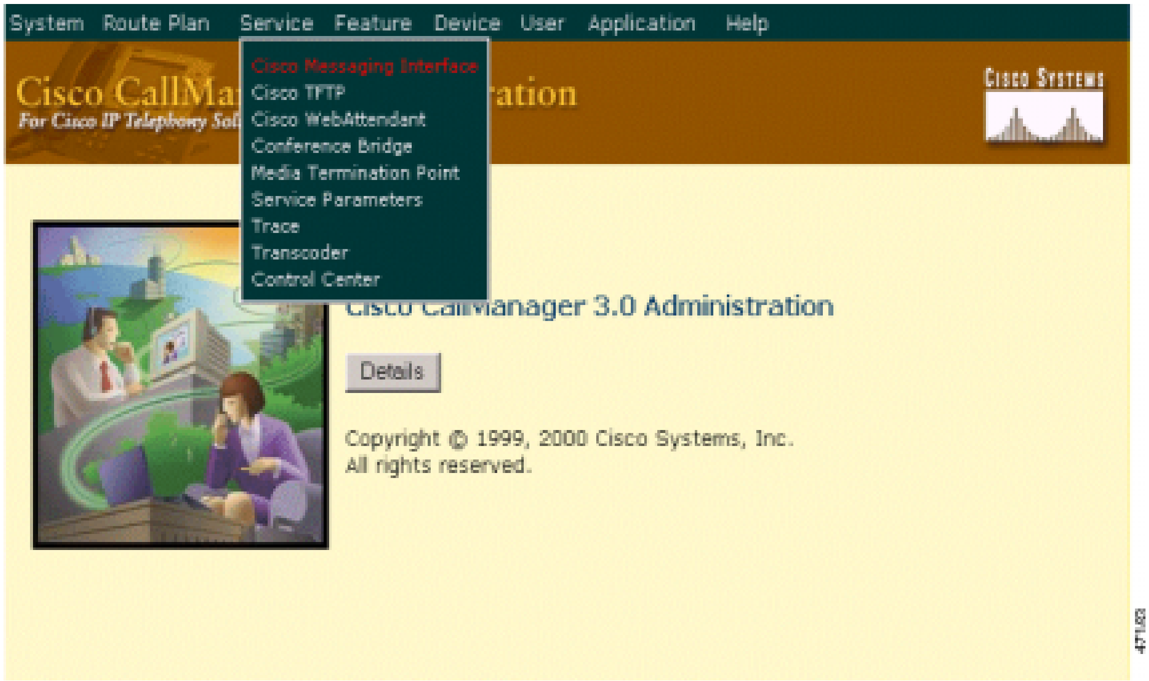

CMI will function in a cluster but can only reside on a single server within a cluster. For example, you cannot have two CMI applications connected to the same voice mail system running at the same time.

Additional voice mail systems will require additional CMI applications. For instance, there could be a cluster with four active CallManagers and each CallManager could have its own voice mail system attached. Each voice mail system would then need it's own CMI application running.

Step 2 Install and configure FXS ports – VG200(s) or Catalyst 6000 24-port Gateway(s).

a. Configure the MGCP (Skinny) gateway. In this case, it is a 4-port VG200 MGCP gateway.

Figure 4-37 MGCP Configuration Screen

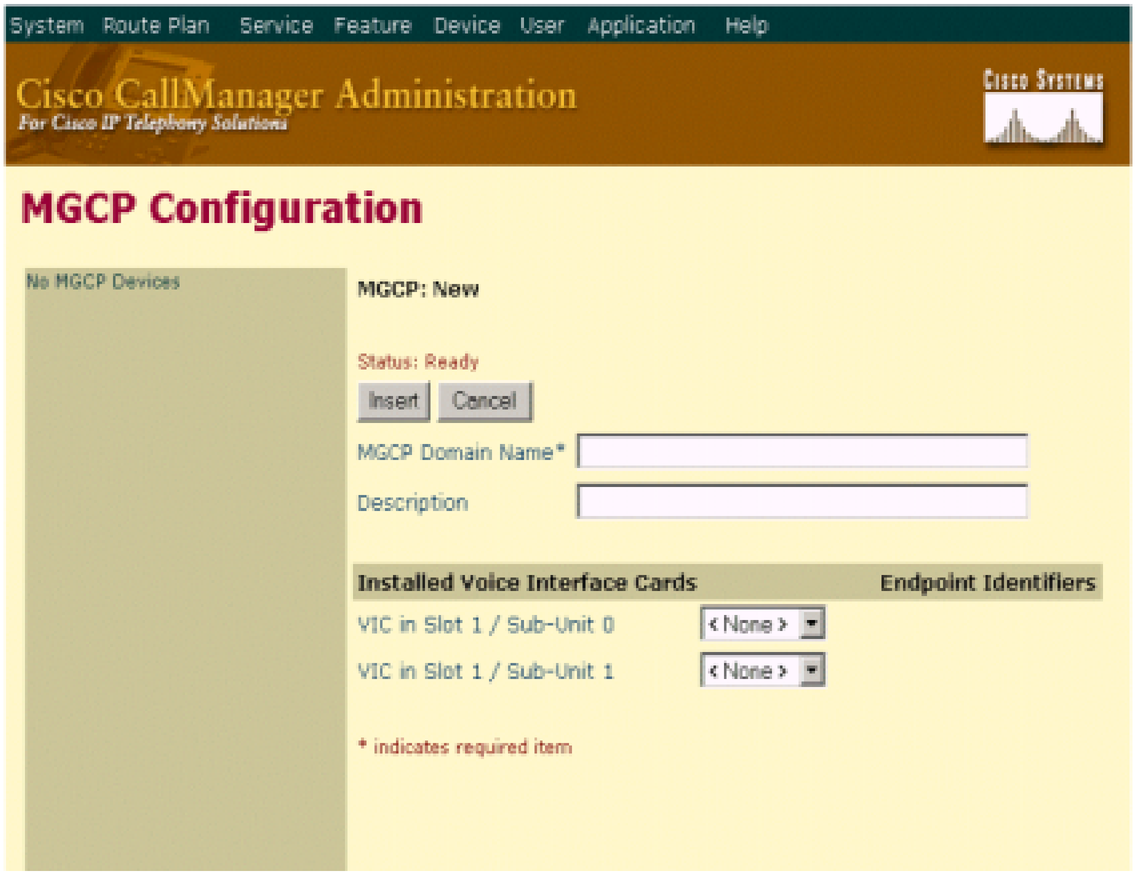

MGCP Configuration

No MGCP Devices

MGCP: New

Status: Ready

MGCP Domain Name*

Description

Installed Voice Interface Cards	Endpoint Identifiers
VIC in Slot 1 / Sub-Unit 0	< None >
VIC in Slot 1 / Sub-Unit 1	< None >

* indicates required item

b. Ensure that you select **FXS** for the Installed Voice Interface Cards.

Figure 4-38 Selecting the Voice Interface Cards

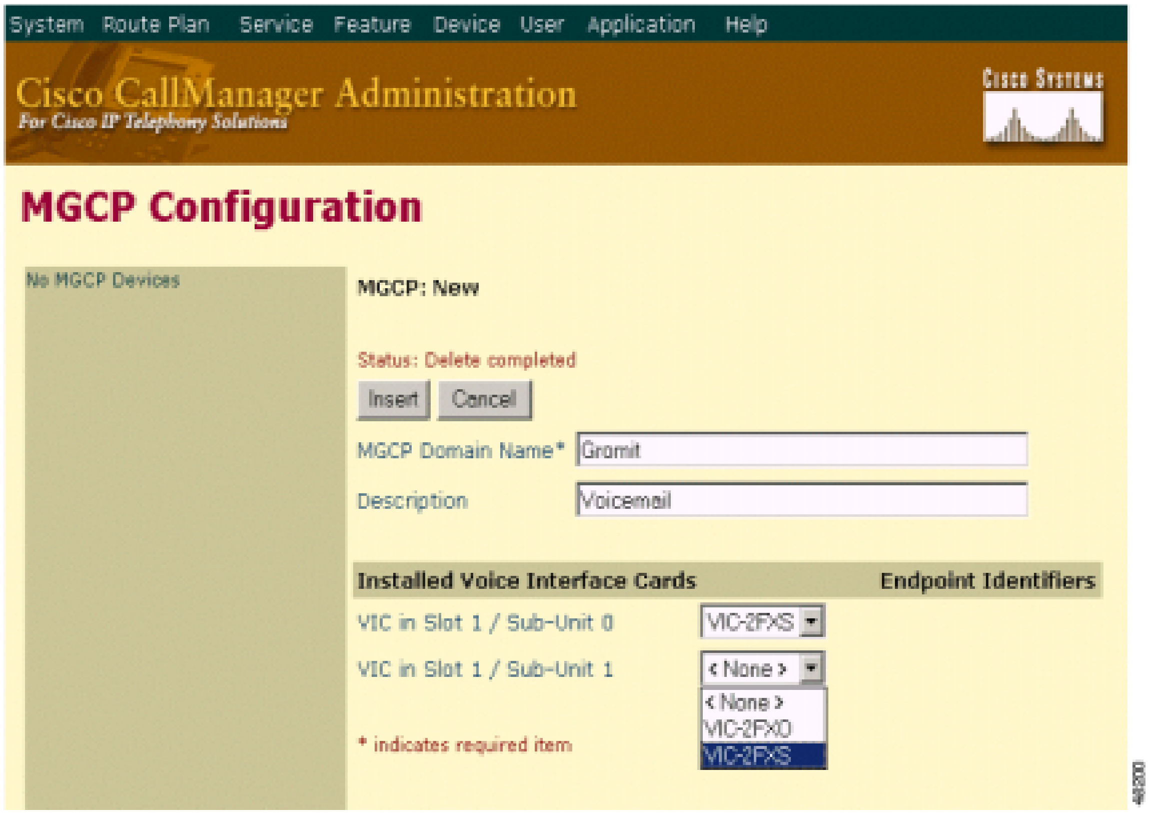
System Route Plan Service Feature Device User Application Help
Cisco CallManager Administration
For Cisco IP Telephony Solutions
Cisco Systems
MGCP Configuration
No MGCP Devices
Insert
Cancel
Installed Voice Interface Cards
Endpoint Identifiers
< None >
VIC-2FXO
VIC-2FXS
48200

MGCP: New

Status: Delete completed

MGCP Domain Name* Gromit

Description Voicemail

VIC in Slot 1 / Sub-Unit 0 VIC-2FXS

VIC in Slot 1 / Sub-Unit 1 < None >

* indicates required item

The gateway now has 4 ports or *Endpoint Identifiers*.

Note Analog ports will always be configured as FXS ports when connecting to voice mail systems.

Figure 4-39 Endpoint Identifiers

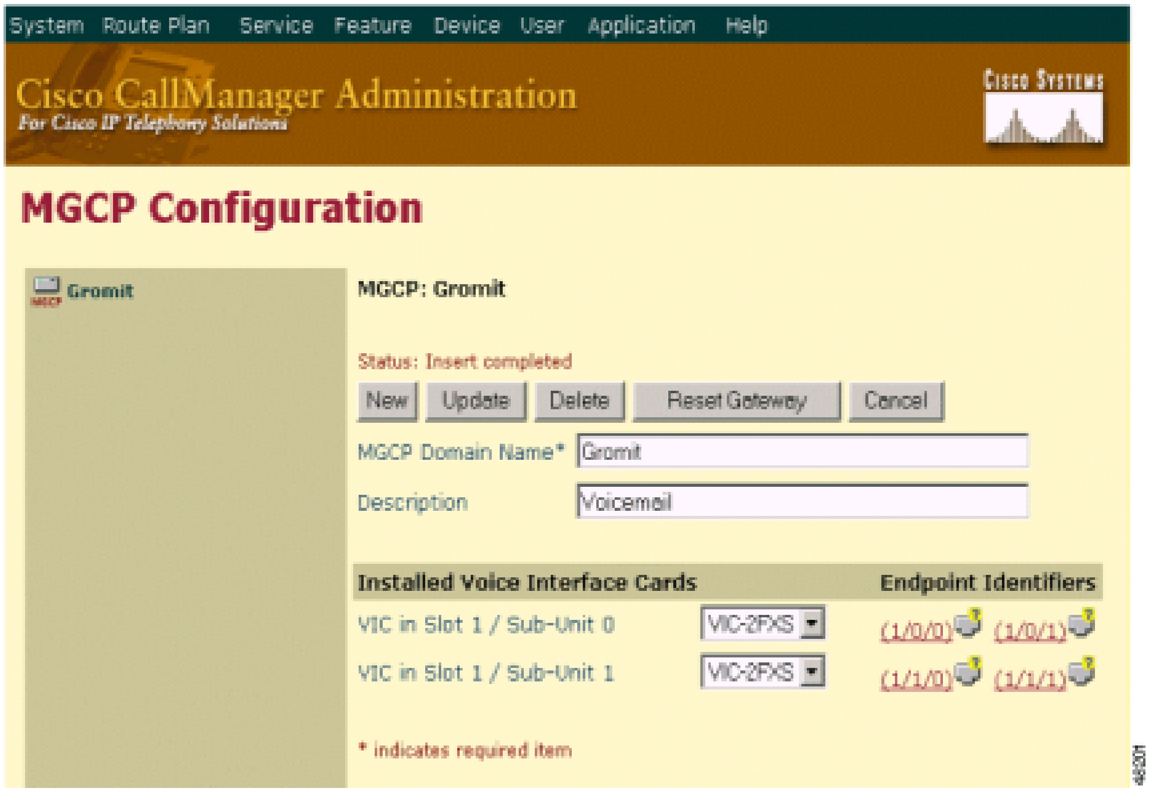

MGCP Configuration

MGCP: Gromit

Status: Insert completed

MGCP Domain Name* Gromit

Description Voicemail

Installed Voice Interface Cards		Endpoint Identifiers
VIC in Slot 1 / Sub-Unit 0	VIC-2FXS	(1/0/0) (1/0/1)
VIC in Slot 1 / Sub-Unit 1	VIC-2FXS	(1/1/0) (1/1/1)

* indicates required item

c. Configure each port as a POTS line.

Figure 4-40 Select Port Type: POTS

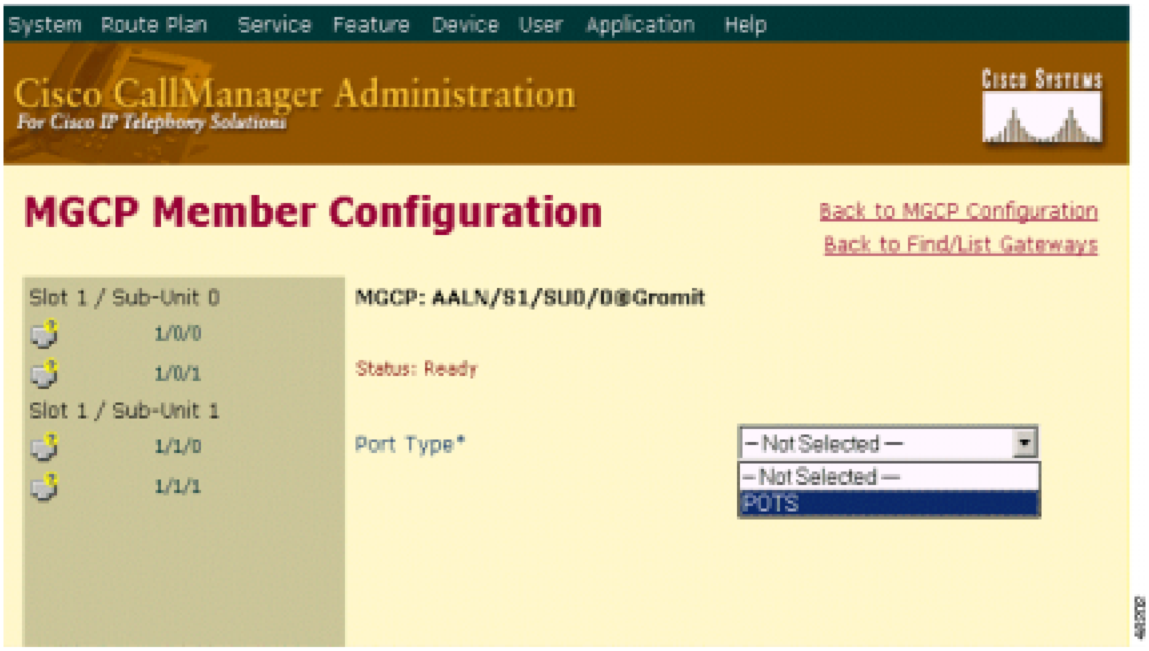

MGCP Member Configuration

MGCP: AALN/S1/SU0/0@Gromit

Status: Ready

Port Type* -- Not Selected --

-- Not Selected --
POTS

The only parameter you need to change is the **Device Pool**; in this case, select **Default**. However, in practice, select the appropriate Device Pool according to the customer's configuration. This will provision the CallManager Redundancy Group, Time Zone, and Region.

Figure 4-41 Configuring the First Port

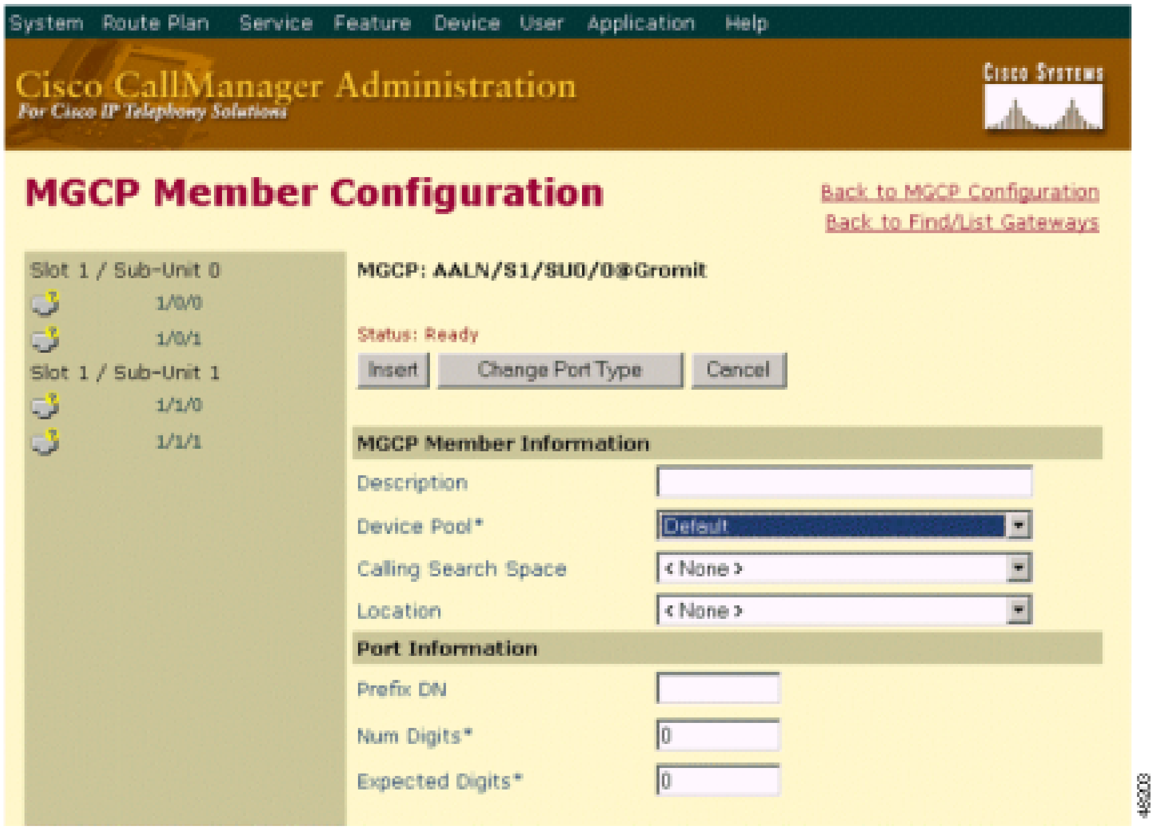

MGCP Member Configuration

MGCP: AALN/S1/SU0/0@Gromit

Status: Ready

MGCP Member Information

Description	
Device Pool*	Default
Calling Search Space	< None >
Location	< None >

Port Information

Prefix DN	
Num Digits*	0
Expected Digits*	0

d. Now that we configured one port as **POTS**, we can configure the remaining three.

Figure 4-42 Configuring the Remaining Ports

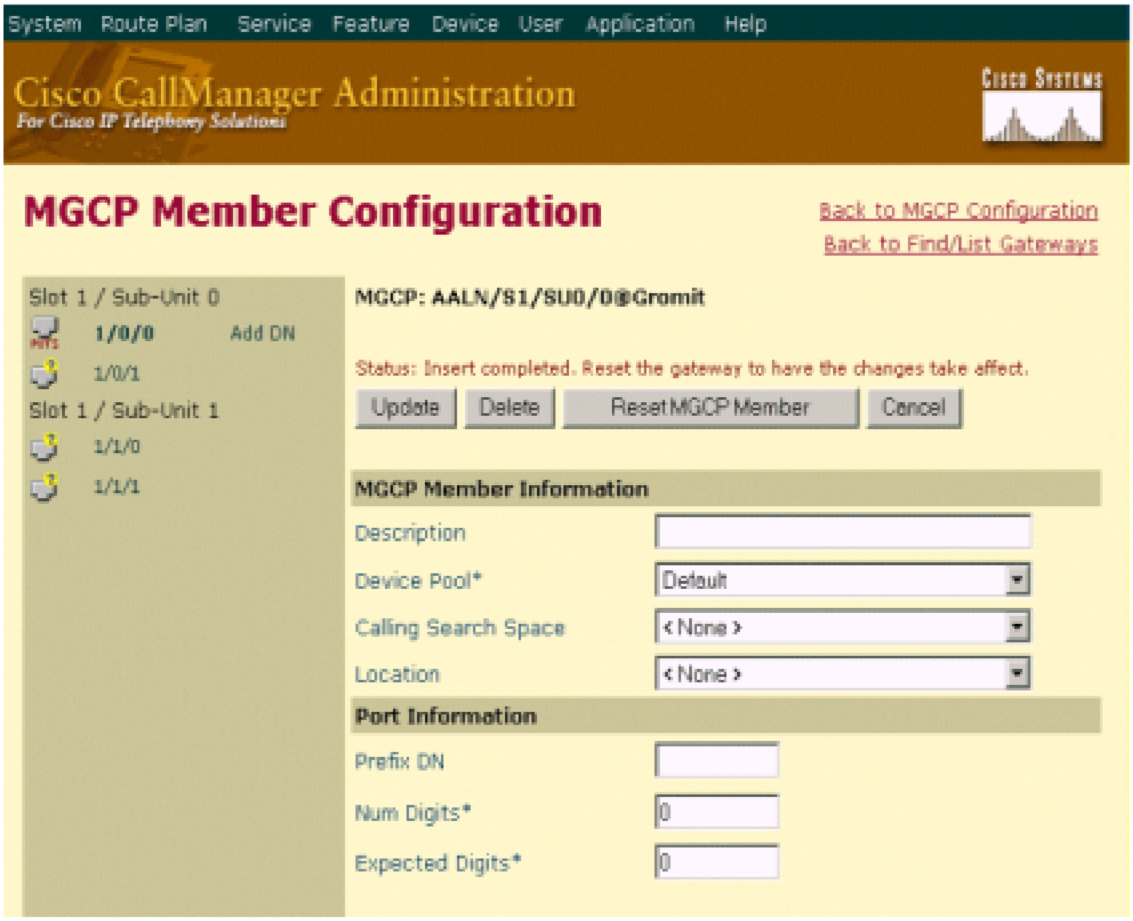

MGCP Member Configuration

MGCP: AALN/S1/SU0/0@Gromit

Status: Insert completed. Reset the gateway to have the changes take affect.

MGCP Member Information

Description	
Device Pool*	Default
Calling Search Space	< None >
Location	< None >

Port Information

Prefix DN	
Num Digits*	0
Expected Digits*	0

Note **Add DN** appears next to the first POTS line. This is not required since the POTS lines will later form a Route Pattern from which the DN will be derived.

We have now configured all four ports as POTS lines.

Figure 4-43 Port Configuration is Complete

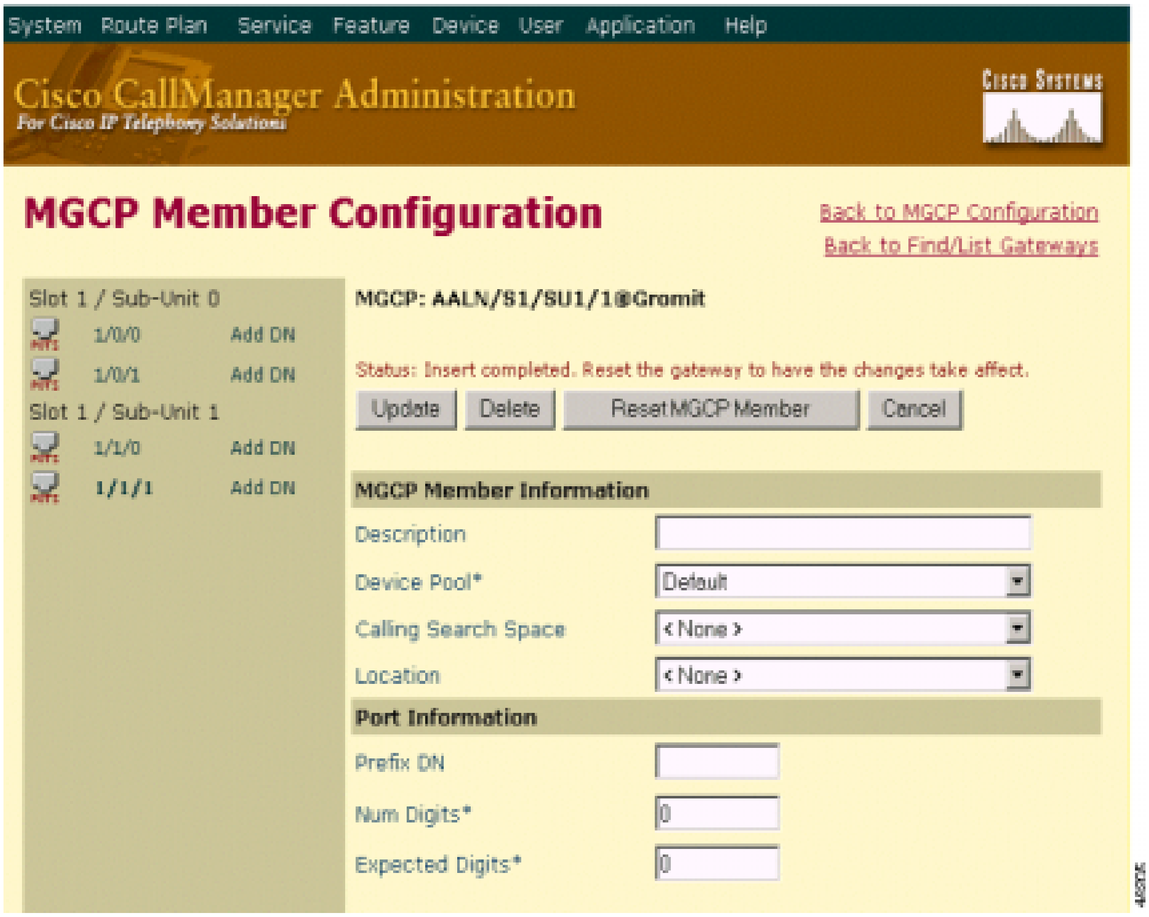
System Route Plan Service Feature Device User Application Help
Cisco CallManager Administration
For Cisco IP Telephony Solutions
Cisco Systems
Back to MGCP Configuration
Back to Find/List Gateways
Slot 1 / Sub-Unit 0
1/0/0 Add DN
1/0/1 Add DN
Slot 1 / Sub-Unit 1
1/1/0 Add DN
1/1/1 Add DN
Update
Delete
Reset MGCP Member
Cancel

MGCP Member Configuration

MGCP: AALN/S1/SU1/1@Gromit

Status: Insert completed. Reset the gateway to have the changes take affect.

MGCP Member Information

Description	
Device Pool*	Default
Calling Search Space	< None >
Location	< None >

Port Information

Prefix DN	
Num Digits*	0
Expected Digits*	0

Step 3 Create Route Group.

a. From the window in Figure 4-43, click on **Update**.

The following window displays.

Figure 4-44 New Route Group Configuration

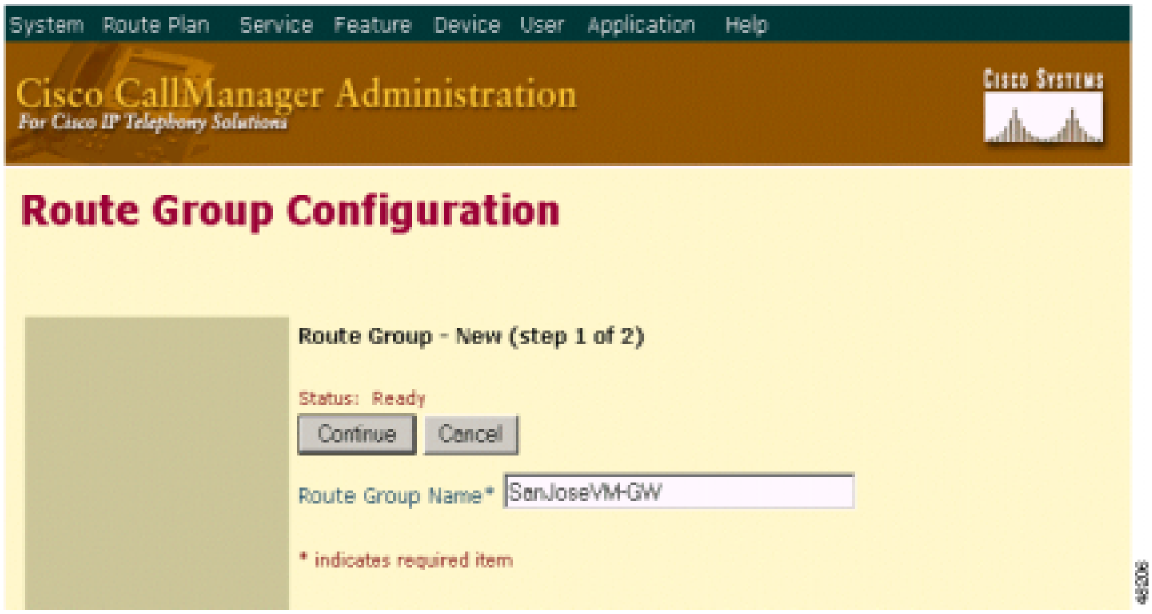

System Route Plan Service Feature Device User Application Help
Cisco CallManager Administration
For Cisco IP Telephony Solutions
Cisco Systems
Continue
Cancel
48206

Route Group Configuration

Route Group - New (step 1 of 2)

Status: Ready

Route Group Name* SanJoseVM-GW

* indicates required item

b. Select the first port from the **Device Name** drop list.

Figure 4-45 Selecting the First Port

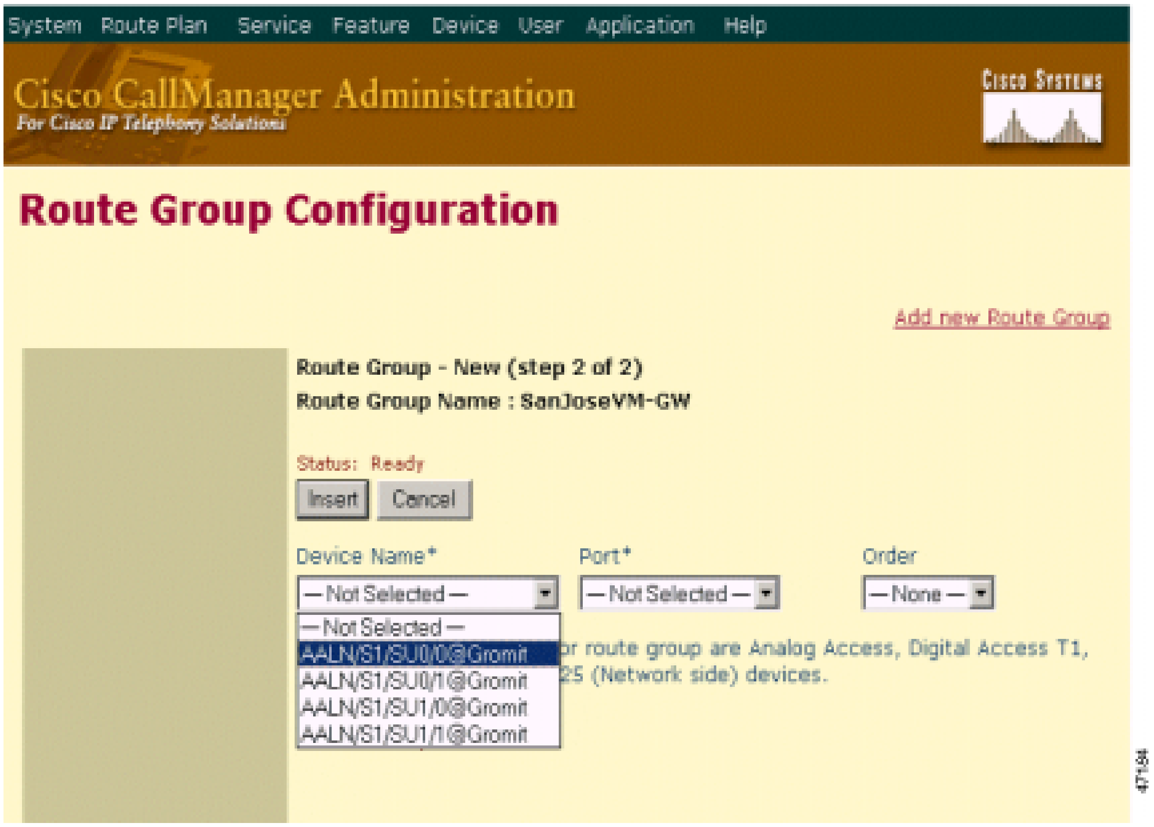

System Route Plan Service Feature Device User Application Help
Cisco CallManager Administration
For Cisco IP Telephony Solutions
Cisco Systems
Add new Route Group
Insert
Cancel
47184

Route Group Configuration

Route Group - New (step 2 of 2)

Route Group Name : SanJoseVM-GW

Status: Ready

Device Name* — Not Selected —

Port* — Not Selected —

Order — None —

— Not Selected —
AALN/S1/SU0/0@Gromit
AALN/S1/SU0/1@Gromit
AALN/S1/SU1/0@Gromit
AALN/S1/SU1/1@Gromit

…r route group are Analog Access, Digital Access T1, …25 (Network side) devices.

c. Select **1** from the **Port** drop list.

Note If you select **All**, you would not be able to assign the correct **Order** to the ports. For example, selecting **All** would always send a Logical Terminal Number (LTN) of 1 within the SMDI packet, regardless of which port was actually seized. This would result in integration failure.

Figure 4-46 Selecting the Port

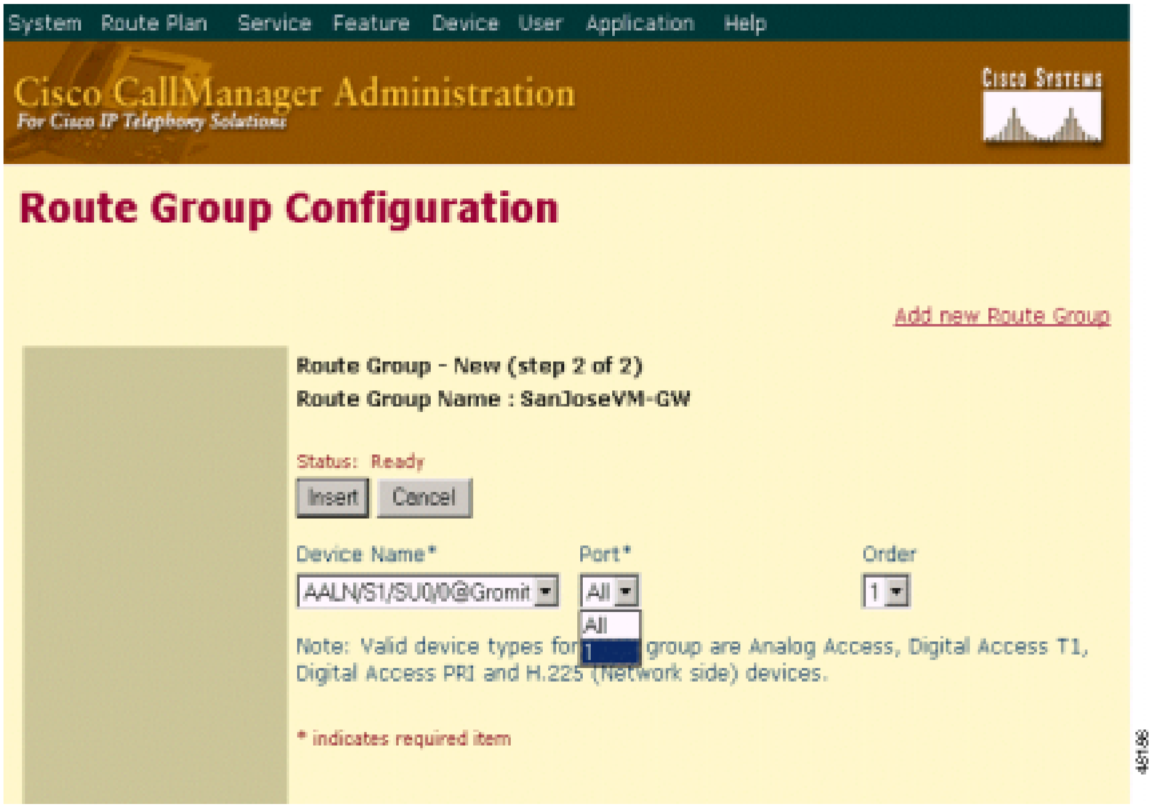

Route Group Configuration

Add new Route Group

Route Group - New (step 2 of 2)

Route Group Name : SanJoseVM-GW

Status: Ready

Device Name*: AALN/S1/SU0/0@Gromit

Port*: All (All, 1)

Order: 1

Note: Valid device types for group are Analog Access, Digital Access T1, Digital Access PRI and H.225 (Network side) devices.

* indicates required item

d. Select **1** from the **Order** drop list. This is the order the port will have within the Route Group.

Figure 4-47 Selecting the Order

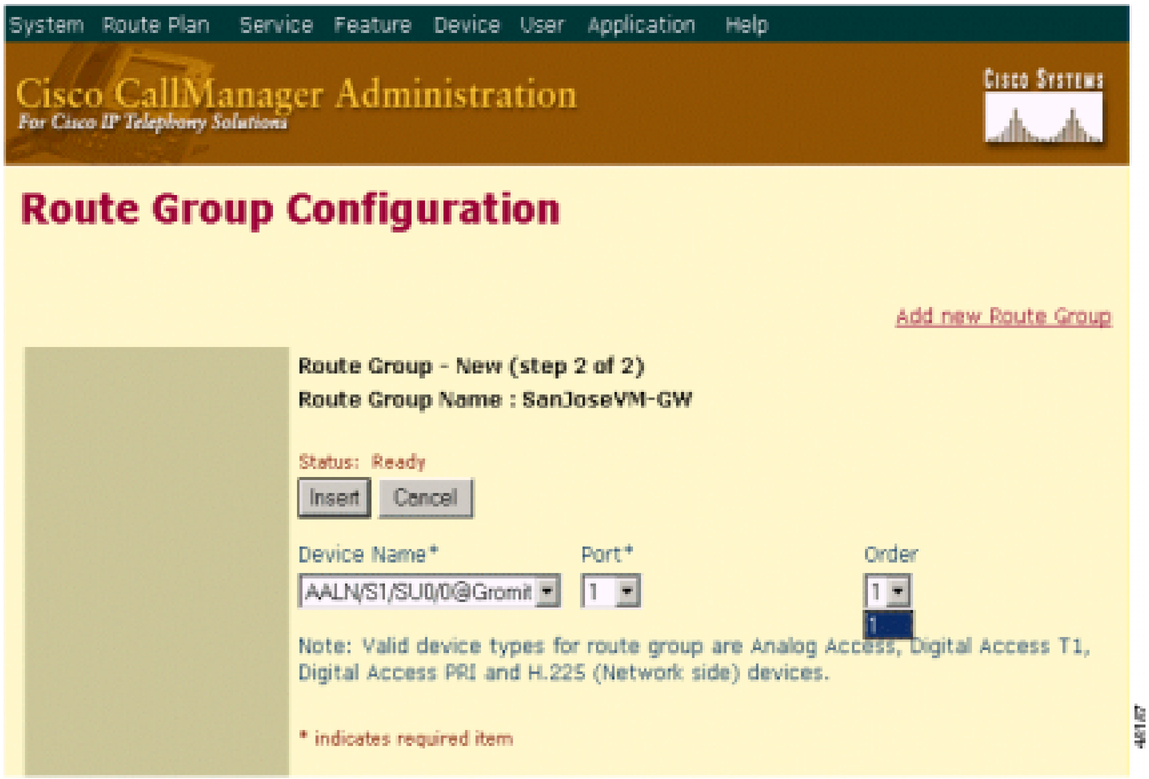

Route Group Configuration

Route Group - New (step 2 of 2)

Route Group Name : SanJoseVM-GW

Status: Ready

Device Name*: AALN/S1/SU0/0@Gromit

Port*: 1

Order: 1

Note: Valid device types for route group are Analog Access, Digital Access T1, Digital Access PRI and H.225 (Network side) devices.

* indicates required item

We have now added one port to the Route Group.

Figure 4-48 Port Added to Route Group

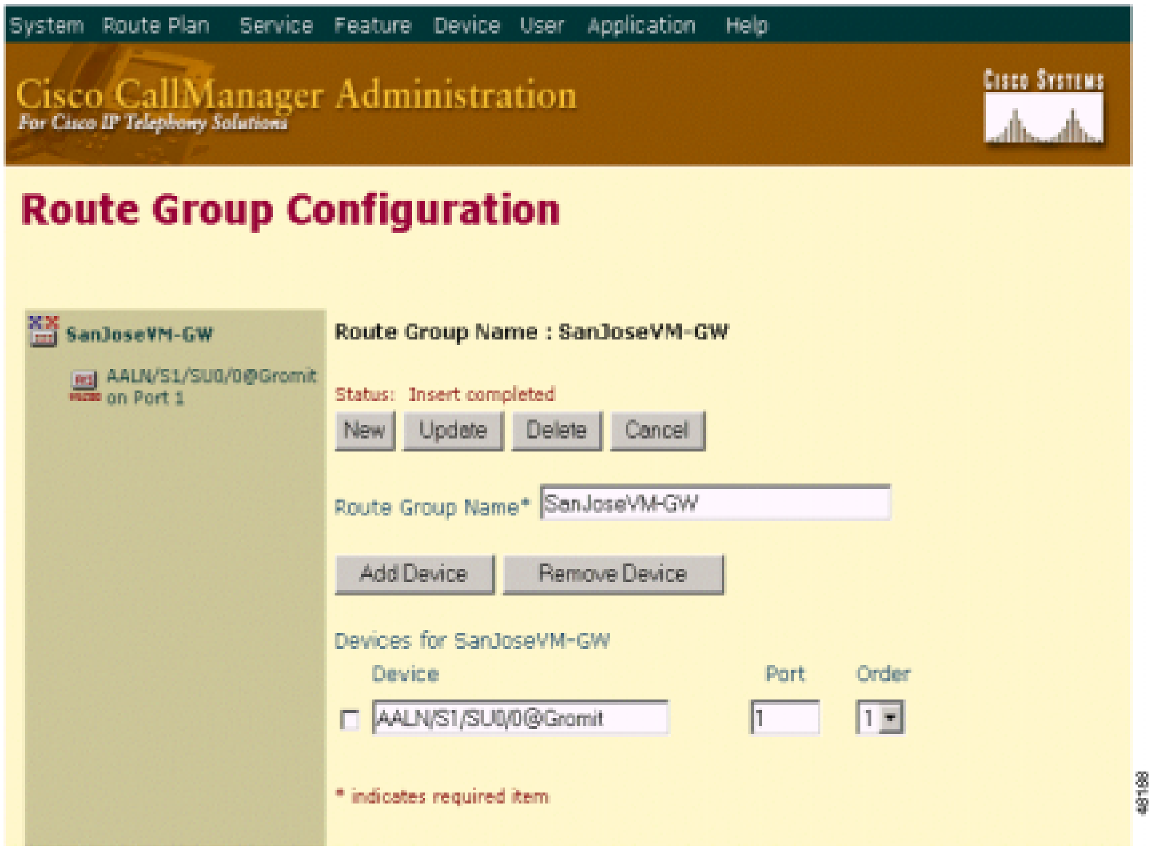

Route Group Configuration

SanJoseVM-GW
AALN/S1/SU0/0@Gromit on Port 1

Route Group Name : SanJoseVM-GW

Status: Insert completed

Route Group Name* SanJoseVM-GW

Devices for SanJoseVM-GW

	Device	Port	Order
□	AALN/S1/SU0/0@Gromit	1	1

* indicates required item

e. Configure the next port by repeating Steps B through E, but select the second port from the **Device Name** drop list and select **2** from the **Order** drop list. Remember to select **1** from the **Port** drop list.

The Route Group now consists of two ports.

Figure 4-49 Second Port Configured

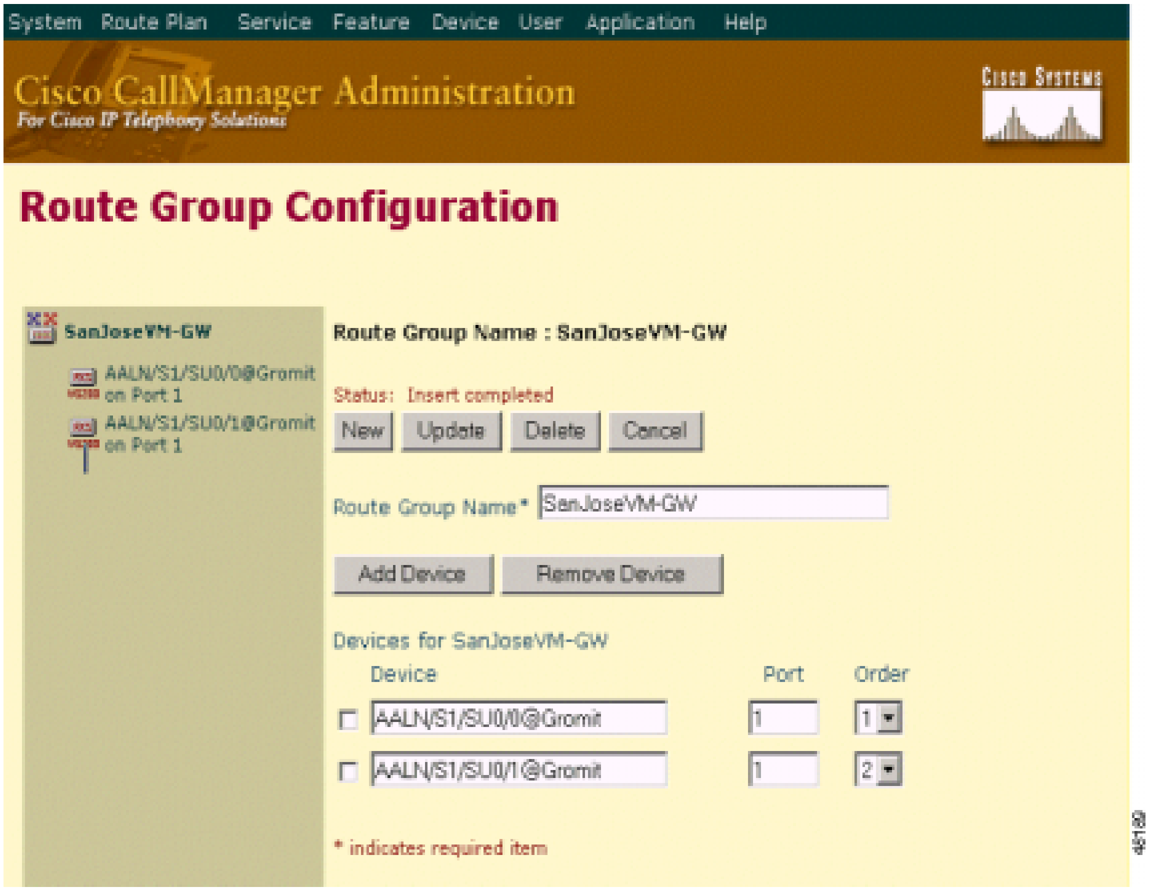

Route Group Configuration

Route Group Name : SanJoseVM-GW

Status: Insert completed

Route Group Name* SanJoseVM-GW

Devices for SanJoseVM-GW

	Device	Port	Order
□	AALN/S1/SU0/0@Gromit	1	1
□	AALN/S1/SU0/1@Gromit	1	2

* indicates required item

f. Configure the remaining ports.

The Route Group is complete.

Figure 4-50 Completed Route Group

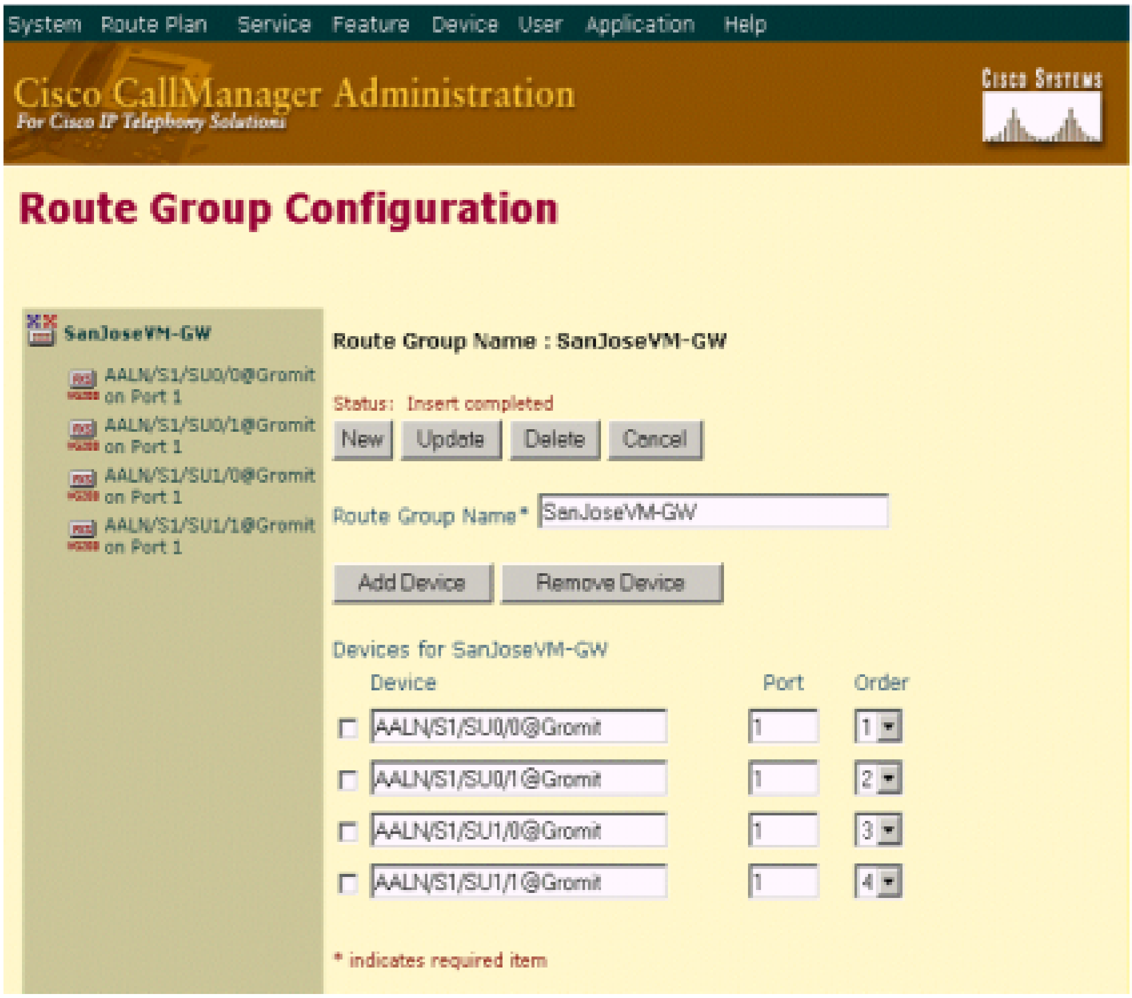

Route Group Configuration

SanJoseVM-GW
- AALN/S1/SU0/0@Gromit on Port 1
- AALN/S1/SU0/1@Gromit on Port 1
- AALN/S1/SU1/0@Gromit on Port 1
- AALN/S1/SU1/1@Gromit on Port 1

Route Group Name : SanJoseVM-GW

Status: Insert completed

Route Group Name* SanJoseVM-GW

Devices for SanJoseVM-GW

	Device	Port	Order
□	AALN/S1/SU0/0@Gromit	1	1
□	AALN/S1/SU0/1@Gromit	1	2
□	AALN/S1/SU1/0@Gromit	1	3
□	AALN/S1/SU1/1@Gromit	1	4

* indicates required item

Step 4 Create the Route List.

a. From the **Route Plan** menu, select **Route List**.

The following window displays.

Figure 4-51 Adding a New Route List

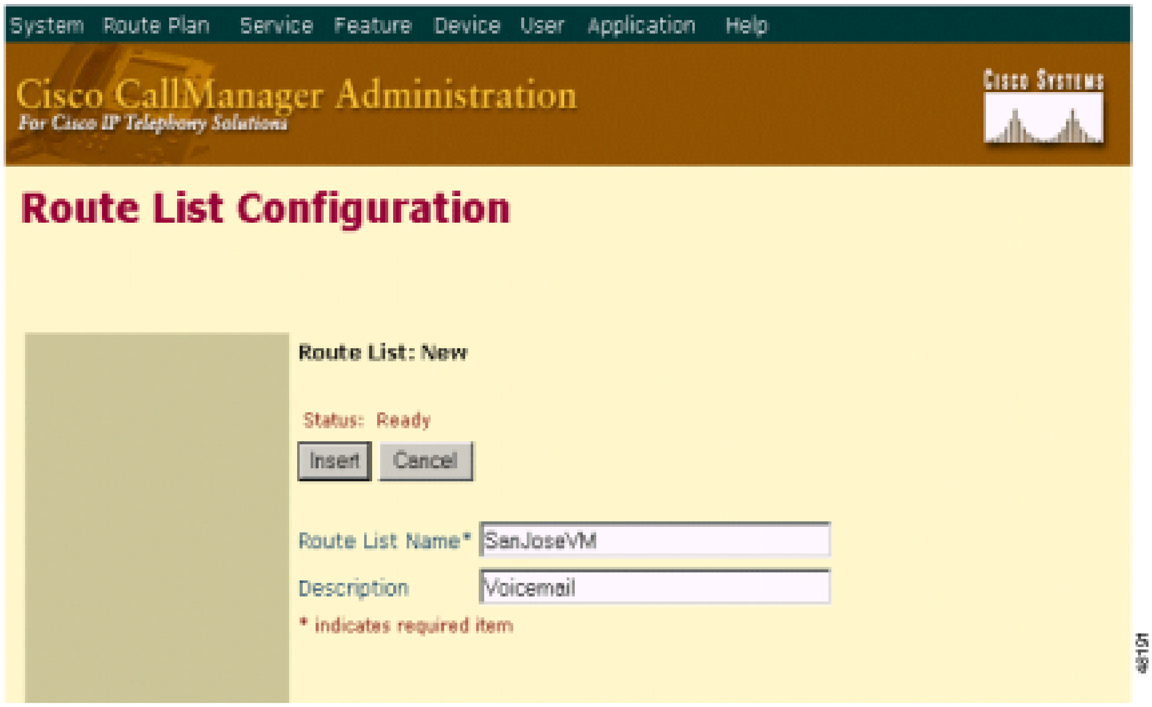

Route List Configuration

Route List: New

Status: Ready

Route List Name* SanJoseVM

Description Voicemail

* indicates required item

b. Enter the **Route List Name** and click on **Insert**.

The following window displays.

Figure 4-52 Adding the Route Group

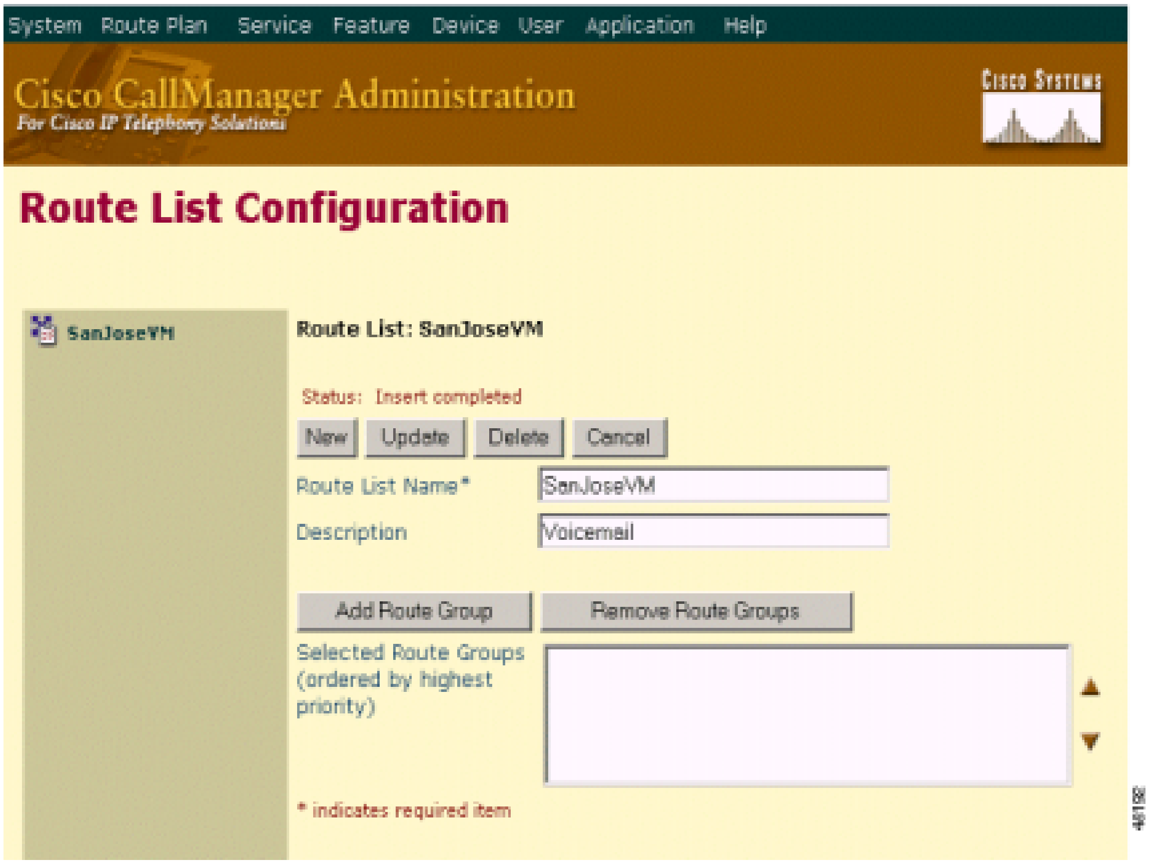
System Route Plan Service Feature Device User Application Help
Cisco CallManager Administration
For Cisco IP Telephony Solutions
Cisco Systems
New
Update
Delete
Cancel
Add Route Group
Remove Route Groups
48192

Route List Configuration

SanJoseVM

Route List: SanJoseVM

Status: Insert completed

Route List Name* SanJoseVM

Description Voicemail

Selected Route Groups (ordered by highest priority)

* indicates required item

c. Click on the **Add Route Group** button.

The following window displays.

Figure 4-53 Adding the Route Group Continued

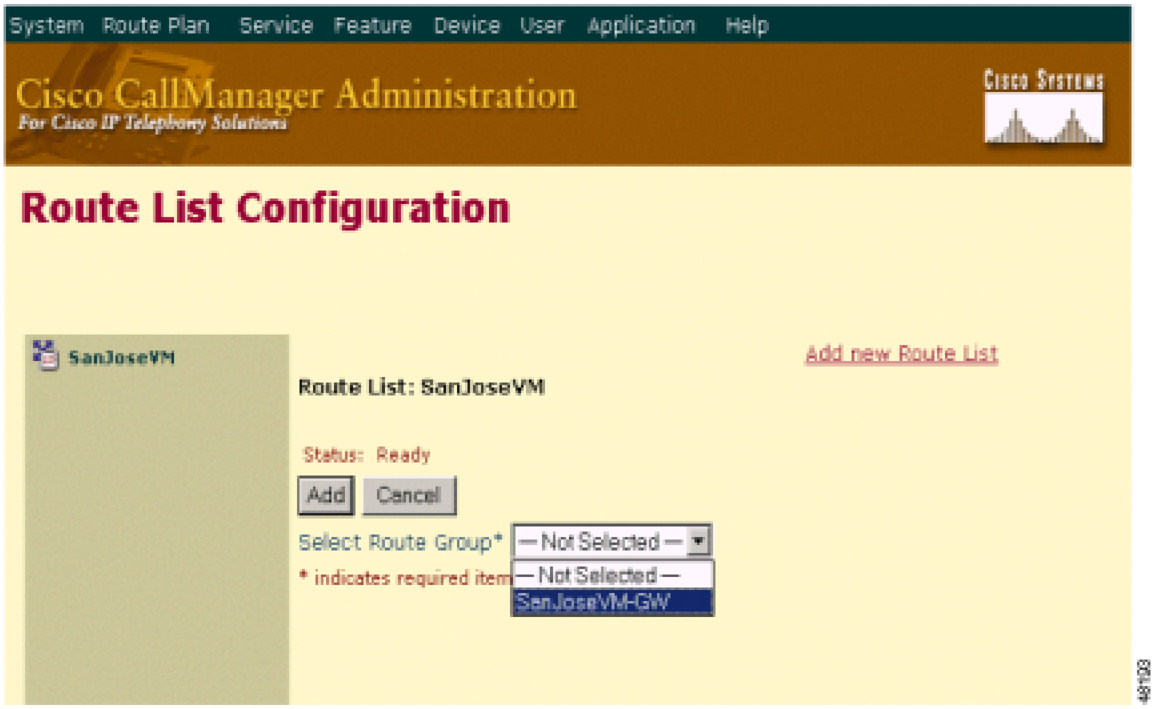

System Route Plan Service Feature Device User Application Help
Cisco CallManager Administration
For Cisco IP Telephony Solutions
Cisco Systems
SanJoseVM
Add new Route List
Add
Cancel
— Not Selected —
— Not Selected —
SanJoseVM-GW
48108

Route List Configuration

Route List: SanJoseVM

Status: Ready

Select Route Group*

* indicates required item

d. Select the appropriate Route Group and click on **Add**.

The Route List is now ready to be inserted.

Figure 4-54 Route Group Selected

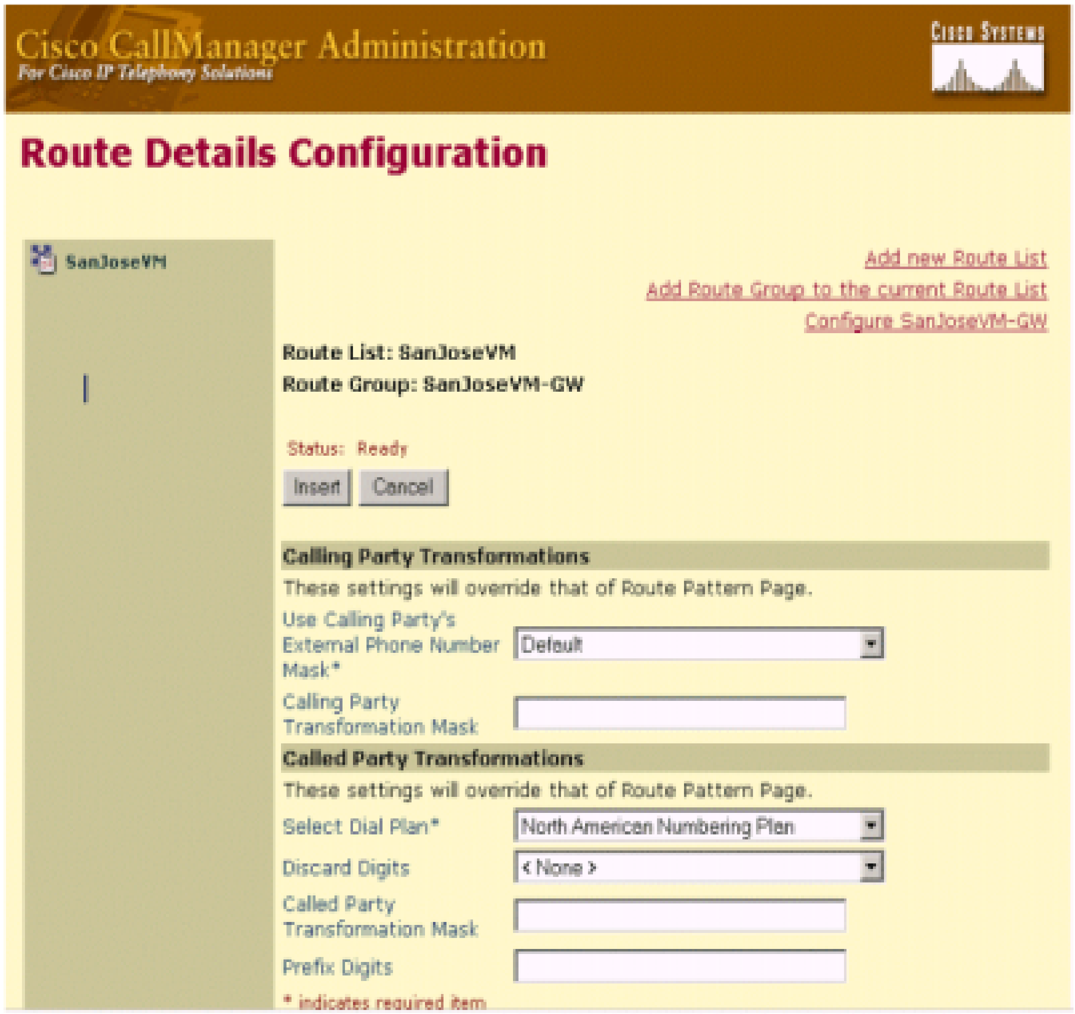

Cisco CallManager Administration
For Cisco IP Telephony Solutions

Route Details Configuration

SanJoseVM

Add new Route List
Add Route Group to the current Route List
Configure SanJoseVM-GW

Route List: SanJoseVM
Route Group: SanJoseVM-GW

Status: Ready

Insert | Cancel

Calling Party Transformations

These settings will override that of Route Pattern Page.

Use Calling Party's External Phone Number Mask*	Default
Calling Party Transformation Mask	

Called Party Transformations

These settings will override that of Route Pattern Page.

Select Dial Plan*	North American Numbering Plan
Discard Digits	< None >
Called Party Transformation Mask	
Prefix Digits	

* indicates required item

e. Click on the **Insert button**.

The Route List is configured and should display like the following screen.

Figure 4-55 Route Details Configuration

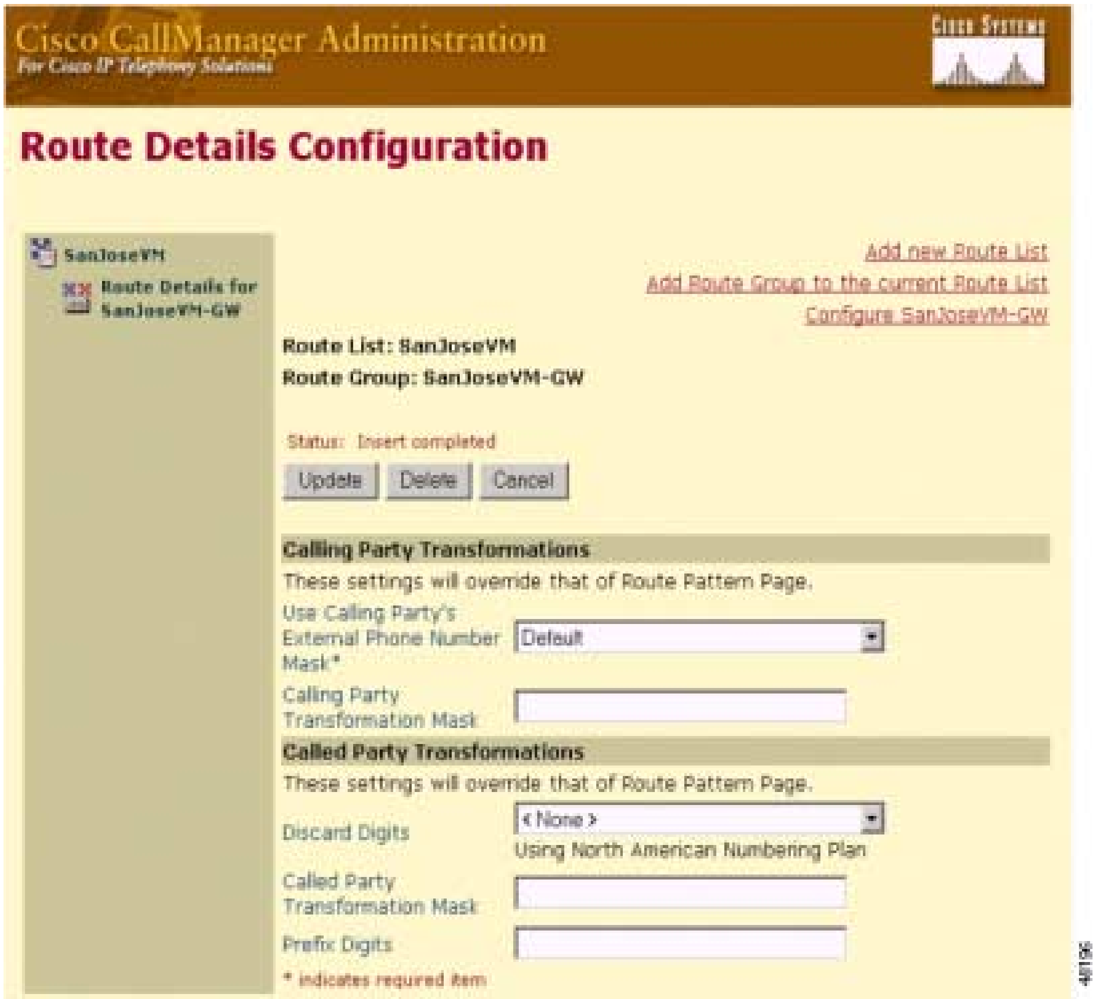
Cisco CallManager Administration
For Cisco IP Telephony Solutions
Cisco Systems
SanJoseVM
Route Details for SanJoseVM-GW
Add new Route List
Add Route Group to the current Route List
Configure SanJoseVM-GW
Update
Delete
Cancel
48196

Route Details Configuration

Route List: SanJoseVM

Route Group: SanJoseVM-GW

Status: Insert completed

Calling Party Transformations

These settings will override that of Route Pattern Page.

Use Calling Party's External Phone Number Mask*: Default

Calling Party Transformation Mask:

Called Party Transformations

These settings will override that of Route Pattern Page.

Discard Digits: < None >
Using North American Numbering Plan

Called Party Transformation Mask:

Prefix Digits:

* indicates required item

Step 5 Create the Route Pattern.

From the window in Figure 4-55, click on **Update**.

In the following window, we assigned a DN of **4000** for our Route Pattern: the access code for voice mail. The **Provide Outside Dial Tone** box is also cleared. If you select this option, subscribers will receive secondary dial-tones as soon as the first digit is entered and some may find this confusing.

Figure 4-56 Completed Route List

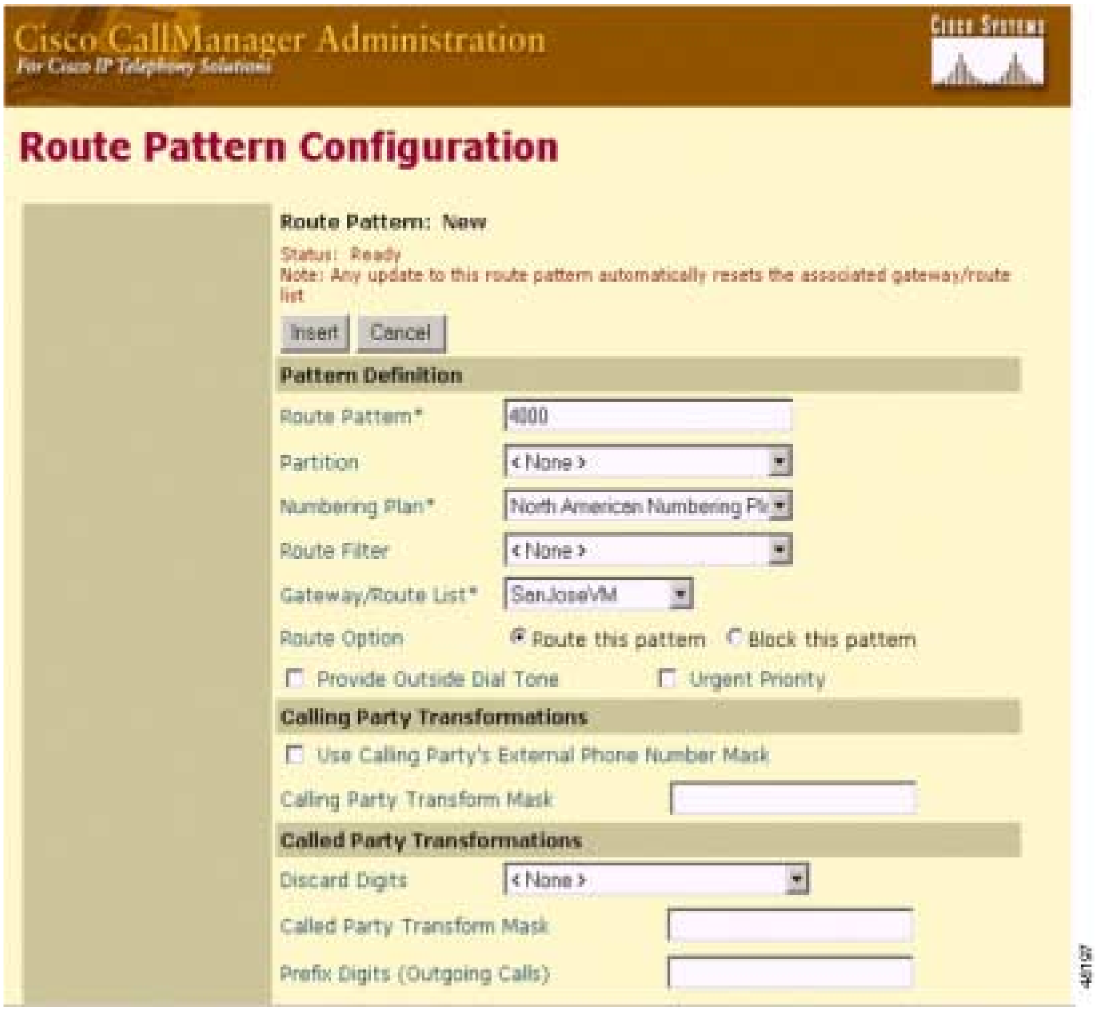

Cisco CallManager Administration
For Cisco IP Telephony Solutions
Cisco Systems

Route Pattern Configuration

Route Pattern: New

Status: Ready
Note: Any update to this route pattern automatically resets the associated gateway/route list

Pattern Definition

Route Pattern* 4000

Partition < None >

Numbering Plan* North American Numbering Pla

Route Filter < None >

Gateway/Route List* SanJoseVM

Route Option ● Route this pattern ○ Block this pattern

□ Provide Outside Dial Tone □ Urgent Priority

Calling Party Transformations

□ Use Calling Party's External Phone Number Mask

Calling Party Transform Mask

Called Party Transformations

Discard Digits < None >

Called Party Transform Mask

Prefix Digits (Outgoing Calls)

Step 6 Configure CMI.

a. From the **Device** menu, select **Cisco Messaging Interface**.

The following window displays.

Figure 4-57 Selecting the Server

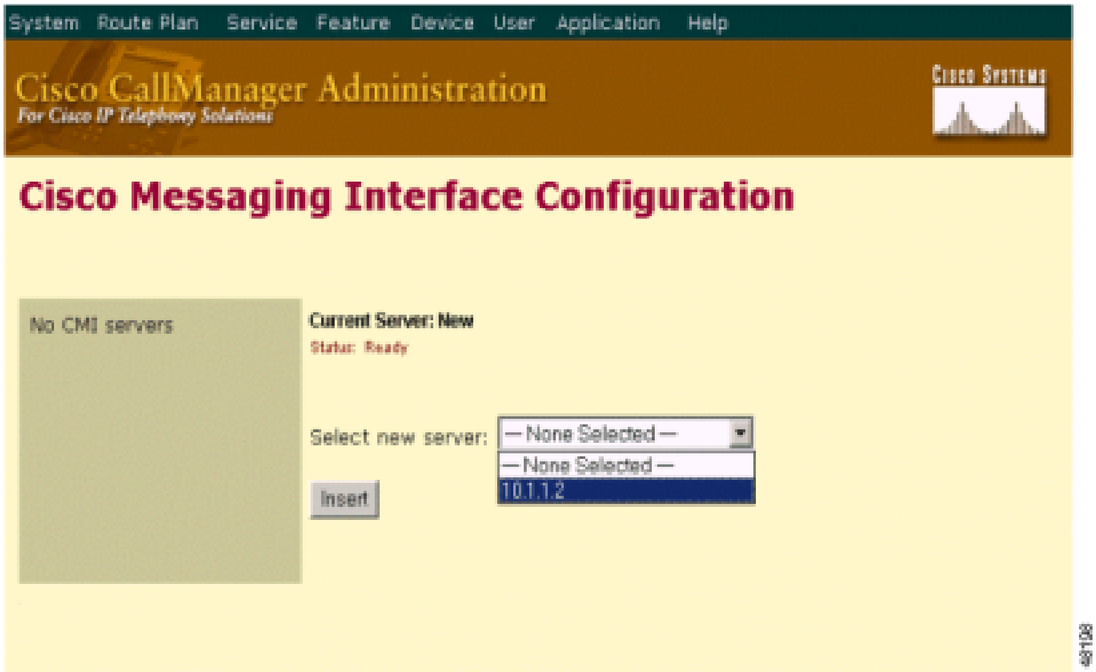

Cisco Messaging Interface Configuration

No CMI servers

Current Server: New

Status: Ready

Select new server: — None Selected —

— None Selected —
10.1.1.2

b. Select the server that will run CMI and click on the **Insert** button.

The following screen displays, which allows you to configure the VoiceMailDN, BaudRate, SerialPort, and other parameters.

Figure 4-58 Setting the CMI Parameters

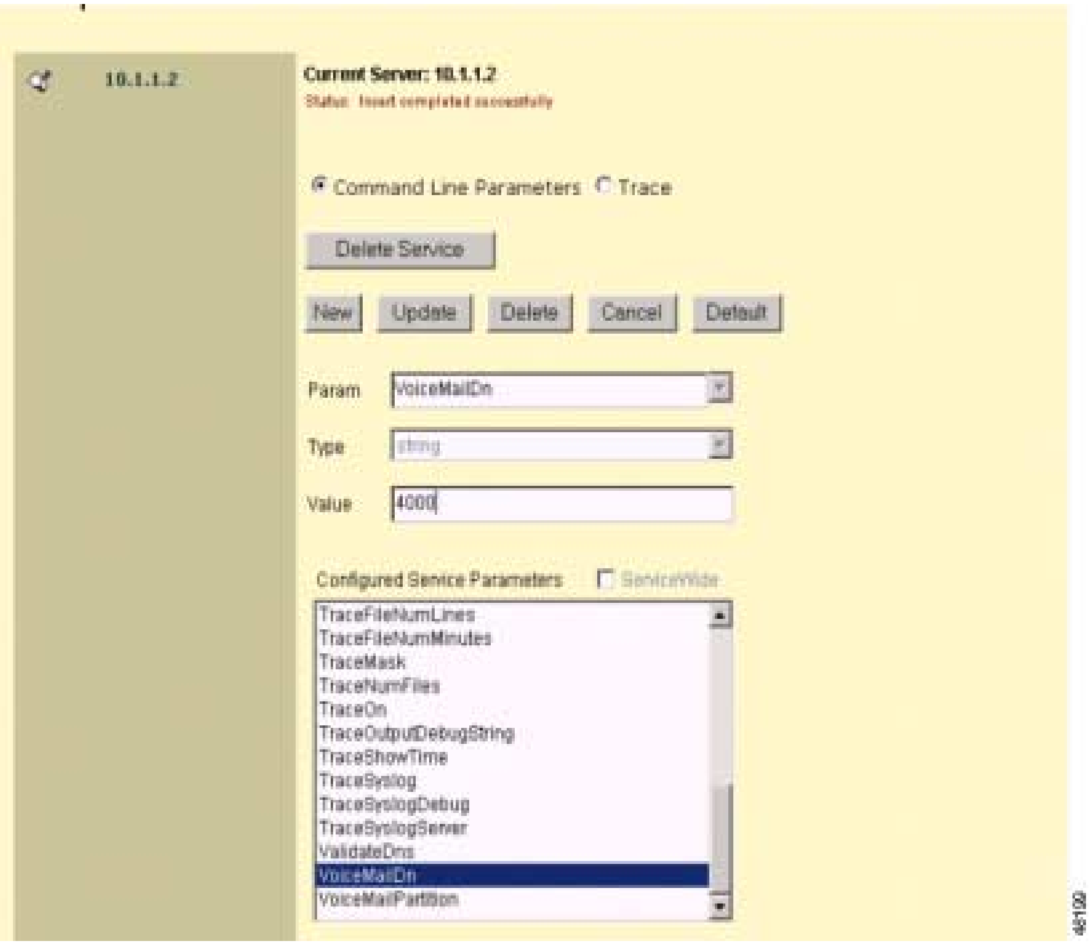

Current Server: 10.1.1.2
Status: Insert completed successfully

Param VoiceMailDn

Type string

Value 4000

Configured Service Parameters □ ServiceWide

TraceFileNumLines
TraceFileNumMinutes
TraceMask
TraceNumFiles
TraceOn
TraceOutputDebugString
TraceShowTime
TraceSyslog
TraceSyslogDebug
TraceSyslogServer
ValidateDns
VoiceMailDn
VoiceMailPartition

c. You need to configure the following parameters in CMI:

Table 4-20 CMI Parameters

Parameter	Explanation
CallManagerName	If you leave this parameter blank, CMI tries to connect to the local host.
BackUpCallManager	In a clustered environment, we recommend that CMI be given a primary (CallManagerName) and a backup (BackUpCallManager). It doesn't cost anything and it ensures that you get voice mail service in the event the primary is down.
VoiceMailDn	You must configure this parameter because it triggers CMI into providing SMDI packets whenever a call is sent to the specified number.
SerialPort	Default is **COM1**. COM port *must* be dedicated to CMI.
BaudRate	Default is **9600**.

Table 4-20 CMI Parameters

Parameter	Explanation
Parity	Default is **Even**.
DataBits	Default is **7**.
StopBits	Default is **1**.
* OutputDnFormat	Default is **%010s**. This parameter allows you to pad the output string with a specified digit. In this case, **%** is a formatting command. The character string **010s** instructs CMI to format the SMDI message as a 10-digit string with leading 0s. For example, in order to have CMI display a 7-digit string, configure **OutputDnFormat** as **%07s**. The **s** simply instructs CMI to display the result as a string.
* OutputExternalFormat	Default is **%010s**. Use this parameter when you have a voice mail system that expects a string length greater than the extension length on CallManager. For example, for a voice mail system expecting a 7-digit string connected to a CallManager configured with 4-digit extensions, you would configure **OutputExternalFormat** as **525%4s**. This would result in a string of **5251234** for extension **1234**.
InputDnSignificantDigits	Default is **10**. In the case of our voice mail system that is sending a 7-digit string for MWI commands to a CallManager system configured for 4-digit extensions, we would configure **InputDnSignificantDigits** as **4**. In other words, we are keeping the 4 least significant digits.

Note * These parameters are optional and in most cases will *not* require any modification.

Also, please note the following:

- In order to configure the subscriber’s **Messages** button, use the **VoiceMail** parameter located under the **Service/Service Parameters/Cisco CallManager** menu. You will need to re-start CallManager in order to activate this change.
- SMDI link messages that CallManager sends to the voice mail system for forwarded calls will always provide **A** as the forwarding reason-code. Forwarding reason-code **A** means it will forward all calls. The correct forwarding reason-codes, **N** (no answer) and **B** for (busy), are not supported at this time. You can configure some voice mail systems, such as Octel products, to provide different greetings based on the forwarding reason code. Cisco CallManager does not support this function.
- SMDI link heart-beat messages – some voice mail systems may send a bogus **OP:MWI** message (e.g., OP:MWI 5551212) expecting a **NAK** response from the far end as a means of determining if the SMDI link is functioning correctly. The **KeepAliveDn** parameter should contain the agreed upon bogus station number for this activity.

- If a message waiting indicator is lit for a phone and the phone is then reset by either cycling power on the device or by the device within the CallManager menus, the message waiting indicator will light up once again after the phone re-registers with the CallManager. Also, if the CallManager is stopped and then re-started, message waiting indicators that were on before the system reset will be turned back on.
- As each parameter within CMI is modified, you *must* perform an **Update**. Otherwise, your modifications will not be written to the database.
- CallManager does *not* currently support configurable Message Desk IDs. The hard-coded ID is **001**. This may present a problem if you are integrating CallManager with a voice mail system that is already integrated with a switch that is using ID **001**. If this is the case, then you must re-configure the switch to another ID to ensure the dual integration functions correctly.
- DB-9 Serial Port Information – CallManager runs on a standard PC, which means it has 9 pins that are actually wired to the serial port. However, CMI only uses three of the RS-232 lines: TD, RD, and GND. There are two outbound control lines: RTS and DTR. These will both be asserted when the port is opened, but CMI doesn't care if the lines are ignored or not. There are four incoming control lines: DCD, CTS, DSR, and RI. All of these lines are ignored. We recommend you connect valid handshaking lines to these lines, but open the port with hardware handshaking disabled so the lines should be ignored.

Troubleshooting SMDI Voice Mail

The only real means of troubleshooting SMDI is to connect a PC with Hyper Terminal to the SMDI port and make test calls.

Tips Don't forget to use a null-modem adapter. This enables you to see exactly what CallManager is sending so you can work with the voice mail technician to resolve any issues.

Remember that the voice mail system will *not* answer a call in the correct manner unless the SMDI packet from CallManager matches the voice mail's port configuration that is receiving the incoming call (e.g., **Message Desk** and **Logical Terminal Number** must match.).

Sample VG200 Configuration Details

Current configuration:

```
!
mgcp
mgcp call-agent 10.1.1.2                        IP Address of CallManager
mgcp dtmf-relay codec all mode out-of-band
mgcp ip-tos precedence 5                        Sets IP “Precedence” of RTP Stream to 5
mgcp default-package hs-package
!
ccm-manager switchback immediate
ccm-manager redundant-host 10.1.1.253           Backup CallManager – if applicable
ccm-manager mgcp
!
interface FastEthernet0/0
ip address 10.1.1.3 255.255.255.0
duplex auto
speed auto
!
ip classless
ip route 0.0.0.0 0.0.0.0 10.1.1.1               Default Route
```

```
no ip http server
!
voice-port 1/0/0
!
voice-port 1/0/1
!
voice-port 1/1/0
!
voice-port 1/1/1
!
dial-peer voice 100 pots
  application MGCPAPP
  port 1/0/0
!
dial-peer voice 101 pots
  application MGCPAPP
  port 1/0/1
!
dial-peer voice 102 pots
 application MGCPAPP
 port 1/1/0
!
dial-peer voice 103 pots
 application MGCPAPP
 port 1/1/1
!
end
```

Example Catalyst 6000 24-Port FXS Gateway Configuration Details

1. From the **Device** menu, select **Gateway**.

   The following window displays.

Figure 4-59 Adding a New Gateway

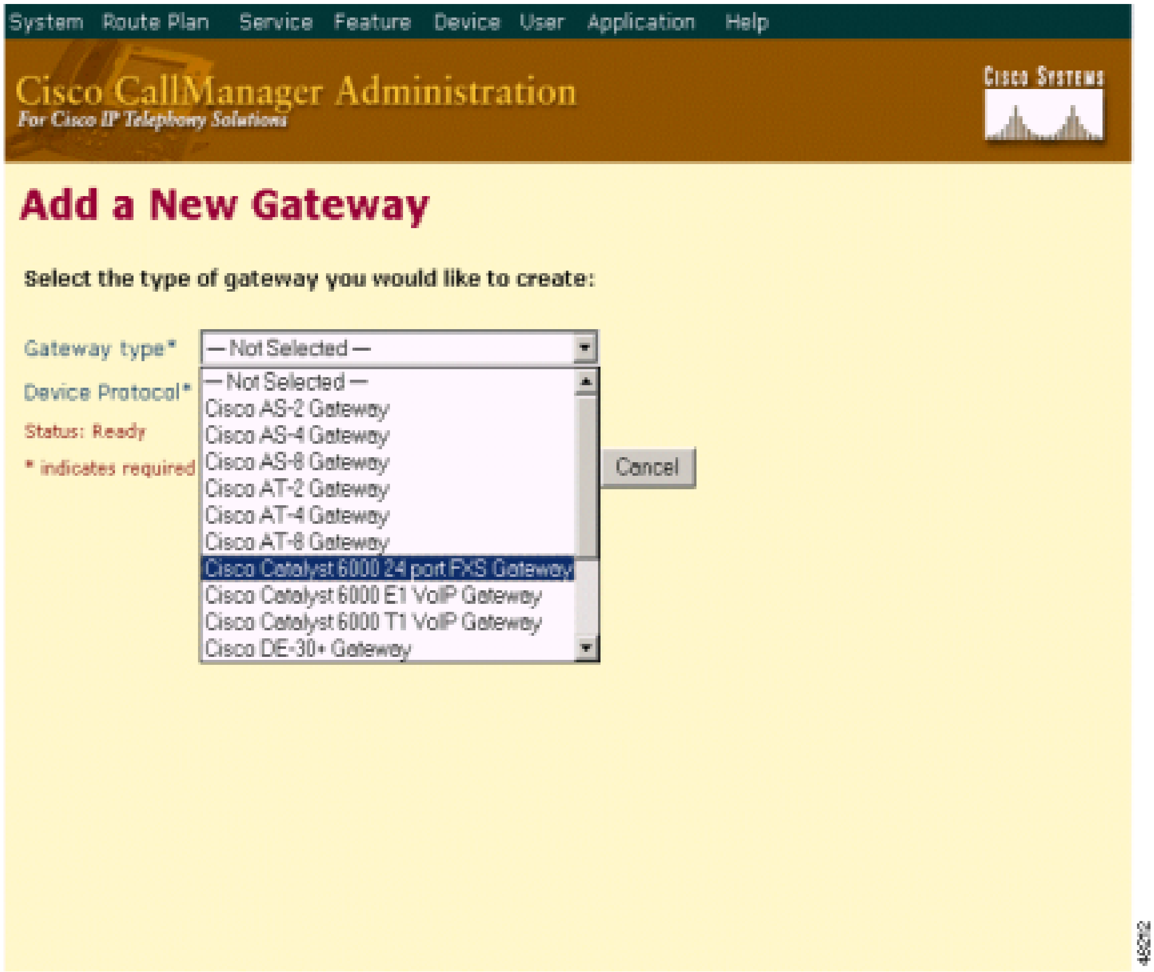

Add a New Gateway

Select the type of gateway you would like to create:

Gateway type* — Not Selected —

Device Protocol*

Status: Ready

* indicates required

2. From the Add a New Gateway window, complete the following fields:
   - **Gateway Type**—select **Cisco Catalyst 6000 24 port FXS Gateway**.
   - **Device Protocol**—select **Analog Access**.
3. Click on **Next**.

   The Gateway Configuration window displays.

Figure 4-60 Gateway Configuration Window

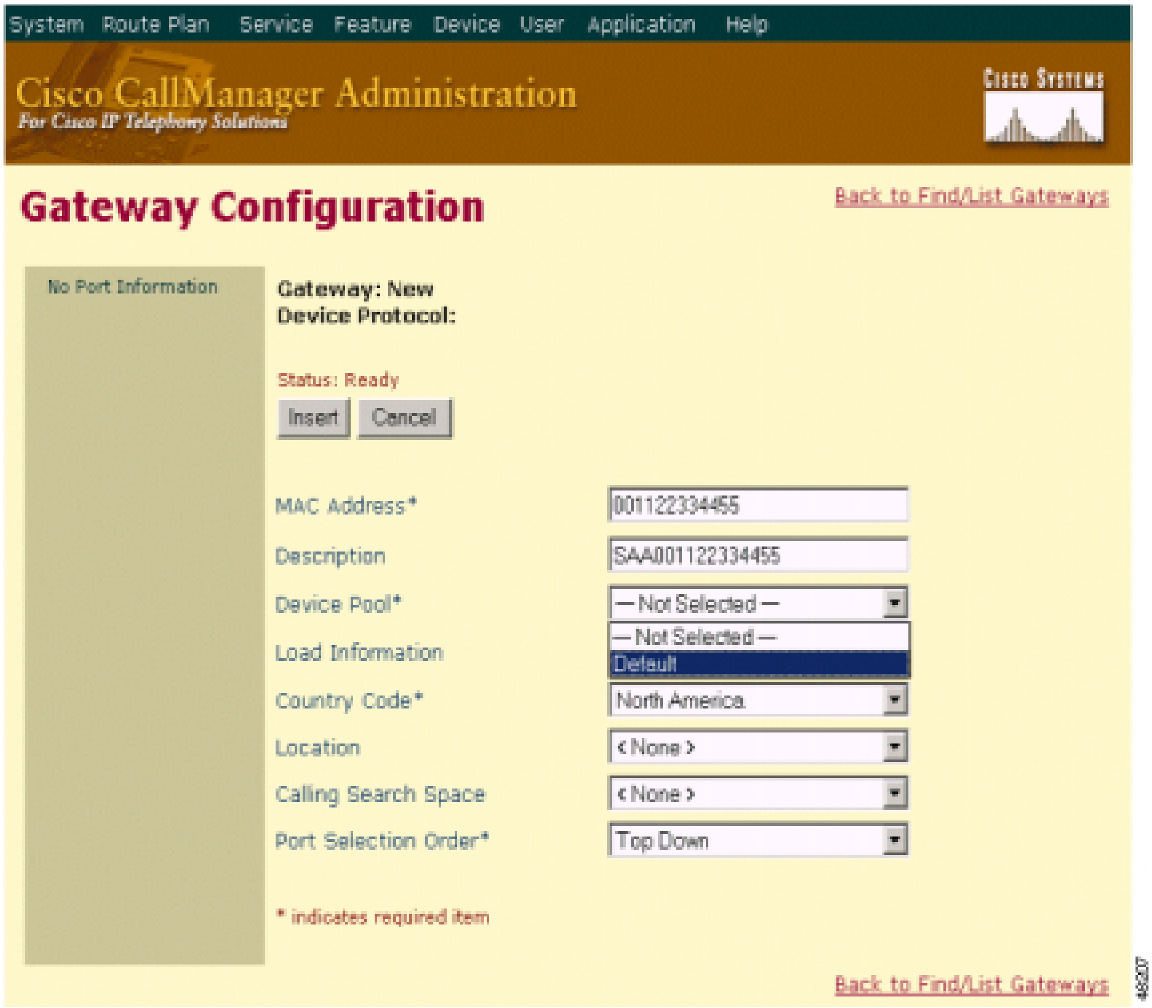
System
Route Plan
Service
Feature
Device
User
Application
Help
Cisco CallManager Administration
For Cisco IP Telephony Solutions
Cisco Systems
Back to Find/List Gateways
No Port Information
Insert
Cancel
Back to Find/List Gateways
48207

Gateway Configuration

Gateway: New
Device Protocol:

Status: Ready

MAC Address* 001122334455

Description SAA001122334455

Device Pool* — Not Selected —

— Not Selected —
Default

Load Information

Country Code* North America

Location < None >

Calling Search Space < None >

Port Selection Order* Top Down

* indicates required item

4. Complete the following fields:
   - **MAC Address**—enter the MAC Address of the Gateway (found from the Catalyst 6000 CLI).
   - **Device Pool**—select **Default**.
   - **Port Selection Order**—select **Top Down** or **Bottom Up**.
5. Click on **Insert**.

   Our Gateway is now inserted.

Figure 4-61 First Gateway Inserted

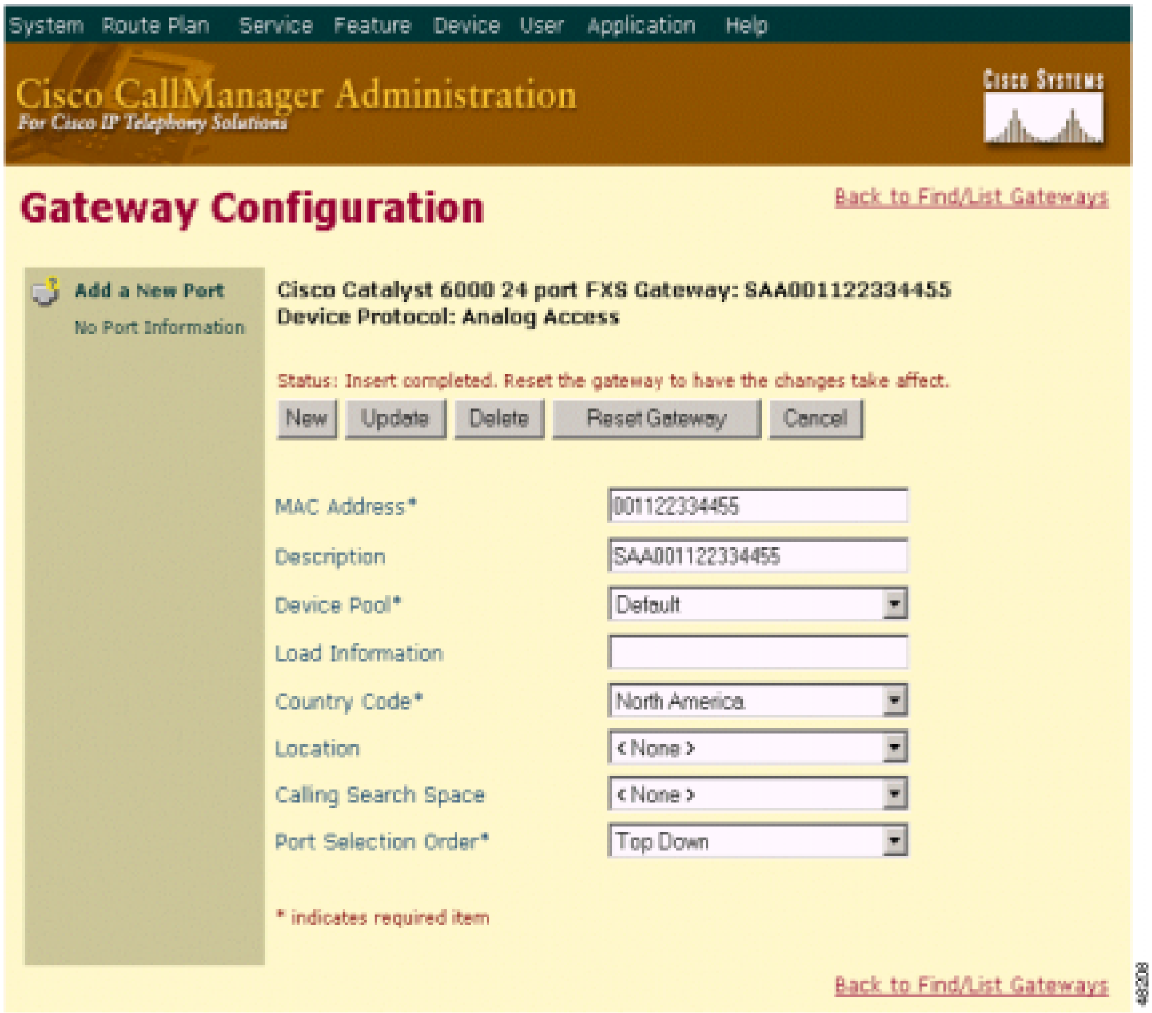
System Route Plan Service Feature Device User Application Help
Cisco CallManager Administration
For Cisco IP Telephony Solutions
Cisco Systems
Back to Find/List Gateways
Add a New Port
New
Update
Delete
Reset Gateway
Cancel
Back to Find/List Gateways
48208

Gateway Configuration

No Port Information

Cisco Catalyst 6000 24 port FXS Gateway: SAA001122334455
Device Protocol: Analog Access

Status: Insert completed. Reset the gateway to have the changes take affect.

MAC Address*	001122334455
Description	SAA001122334455
Device Pool*	Default
Load Information	
Country Code*	North America
Location	< None >
Calling Search Space	< None >
Port Selection Order*	Top Down

* indicates required item

6. Click on **Add New Port**.

   A pop-up window displays.

Figure 4-62 Port Type Pop-up Window

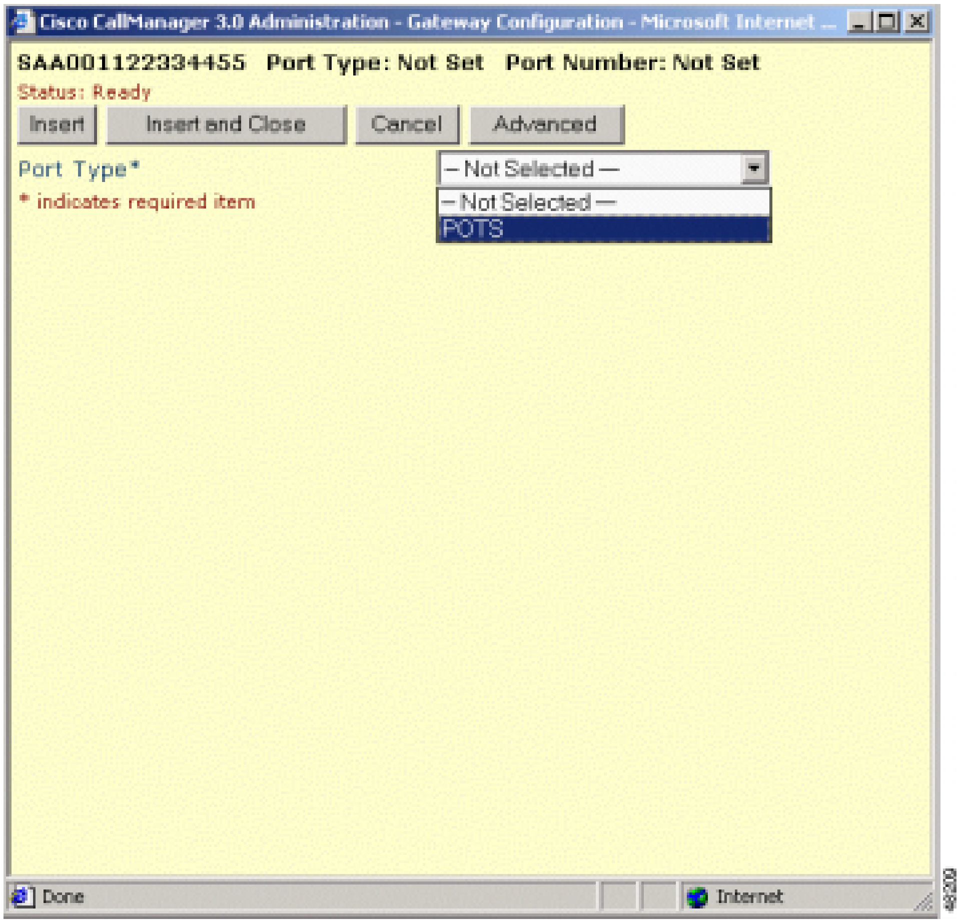

SAA001122334455 Port Type: Not Set Port Number: Not Set

Status: Ready

Port Type* — Not Selected —

* indicates required item

Cisco CallManager 3.0 Administration - Gateway Configuration - Microsoft Internet ...
Insert
Insert and Close
Cancel
Advanced
— Not Selected —
POTS
Done
Internet
48209

7. Complete the following fields:
   - **Port Type**—select **POTS**.
   - **Port Number**—select **Port - 1**. If you select **All Ports**, you would not be able to assign the correct order to the ports. For example, it would always send a Logical Terminal Number (LTN) of 1 within the SMDI packet regardless of which port was actually seized. This results in integration failure.
   - **End Port Number**—select **Port - 1**.
8. Click on **Insert**.

   Our Gateway should now have one FXS Port configured as shown below.

Figure 4-63 Gateway with One FXS Port Configured

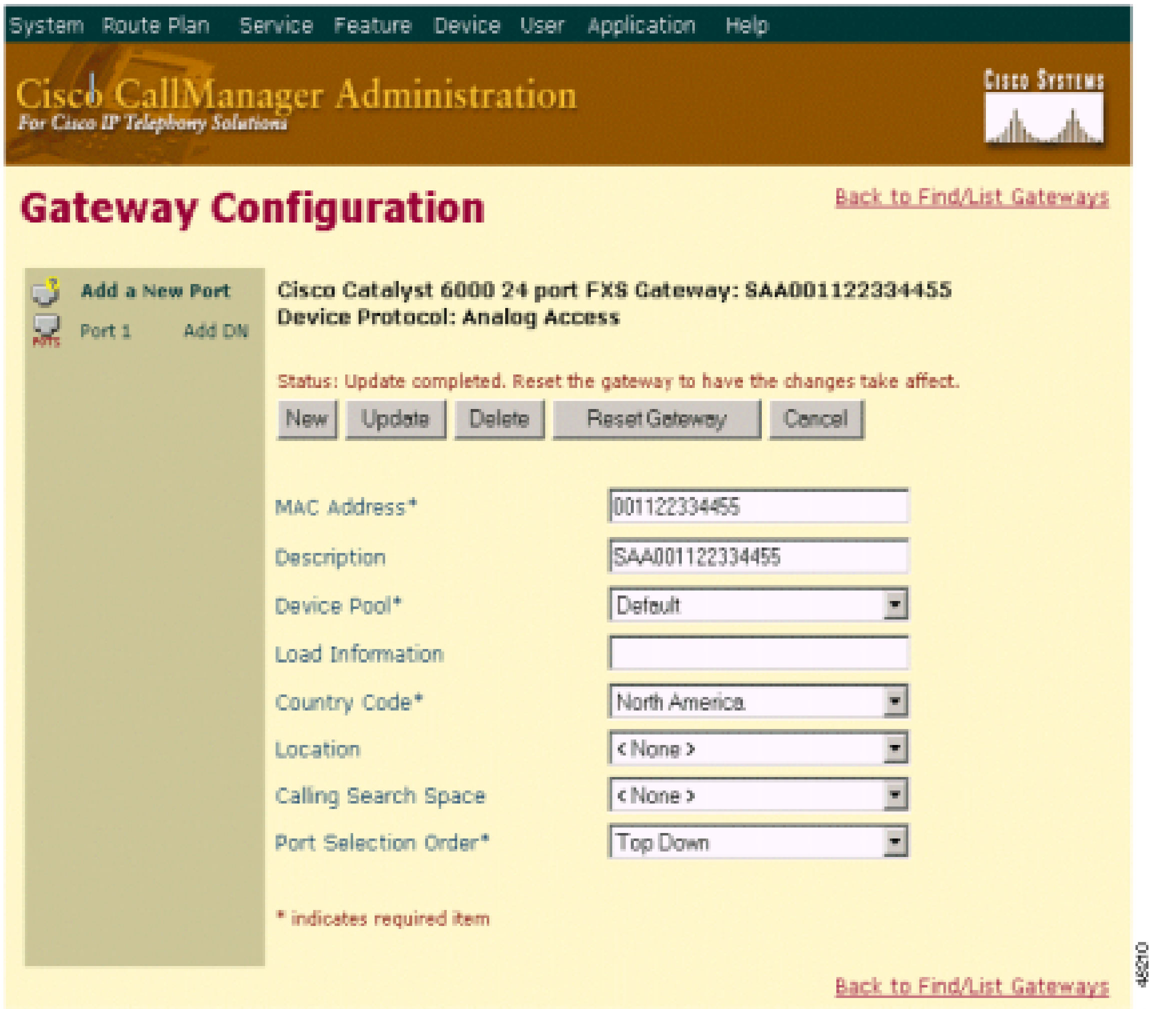
System Route Plan Service Feature Device User Application Help
Cisco CallManager Administration
For Cisco IP Telephony Solutions
Cisco Systems
Back to Find/List Gateways
Add a New Port
Port 1
Add DN
New
Update
Delete
Reset Gateway
Cancel
48210

Gateway Configuration

Cisco Catalyst 6000 24 port FXS Gateway: SAA001122334455
Device Protocol: Analog Access

Status: Update completed. Reset the gateway to have the changes take affect.

MAC Address*	001122334455
Description	SAA001122334455
Device Pool*	Default
Load Information	
Country Code*	North America
Location	< None >
Calling Search Space	< None >
Port Selection Order*	Top Down

* indicates required item

9. Repeat Steps 1 through 7 for each gateway.

Note Ensure you select the appropriate **Port Number** and **End Port Number** in Step 6 for each gateway. For example, for the second gateway, select **Port -2** for these fields; for the third gateway, select **Port-3** and so on.

When finished, the completed configuration should look like the following:

Figure 4-64 Gateway with Twelve FXS Ports Configured

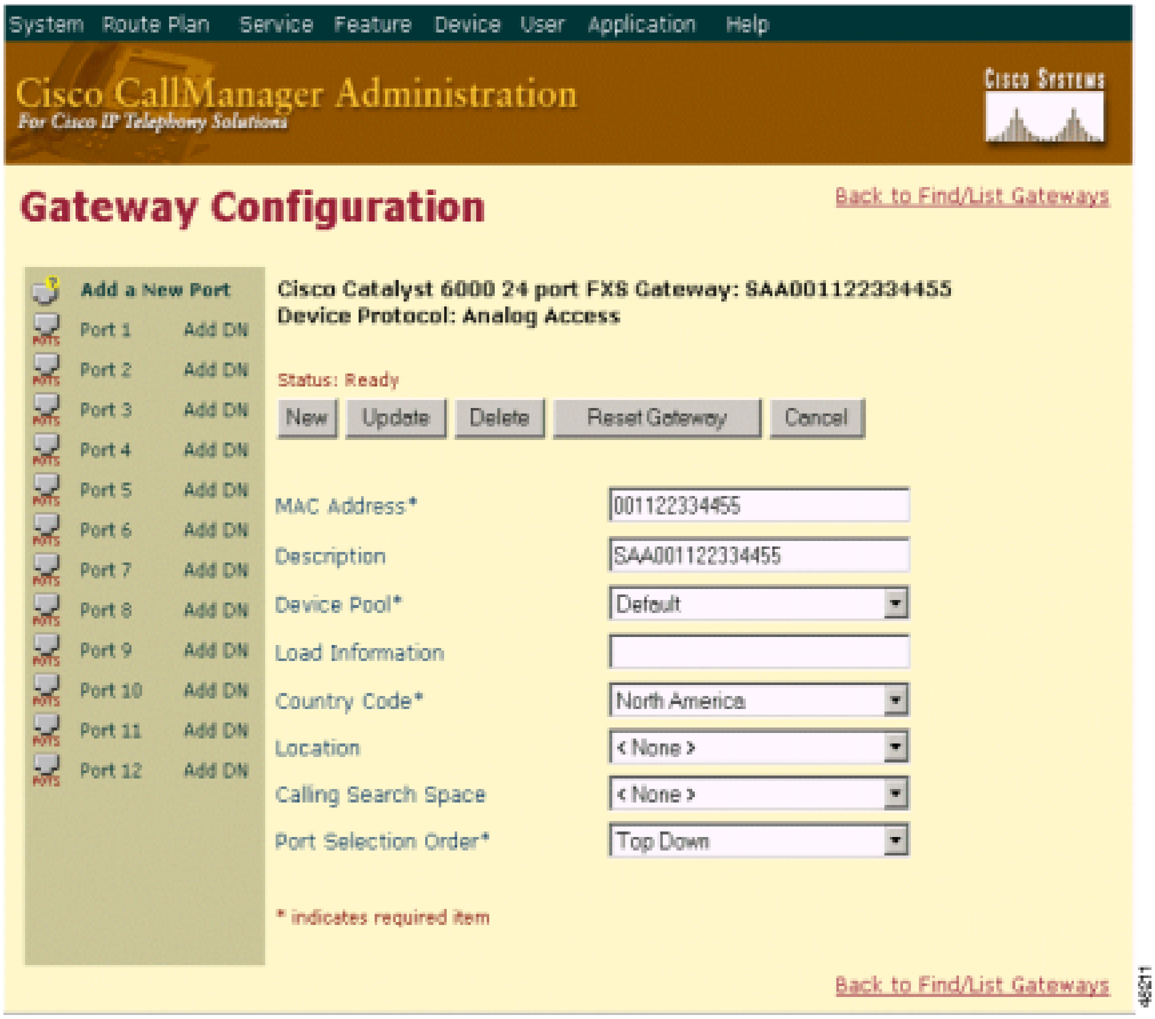

Gateway Configuration

Cisco Catalyst 6000 24 port FXS Gateway: SAA001122334455
Device Protocol: Analog Access

Status: Ready

MAC Address*	001122334455
Description	SAA001122334455
Device Pool*	Default
Load Information	
Country Code*	North America
Location	< None >
Calling Search Space	< None >
Port Selection Order*	Top Down

* indicates required item

Note In order for the FXS Ports to provide dial-tone on disconnect (as opposed to reorder-tone), you must set the **Call Restart Timer** parameter in the Port Configuration window to **1234**. Any other value in that parameter makes the Port provide a reorder-tone on disconnect. See the following figure.

Figure 4-65 Setting the Call Restart Timer Option

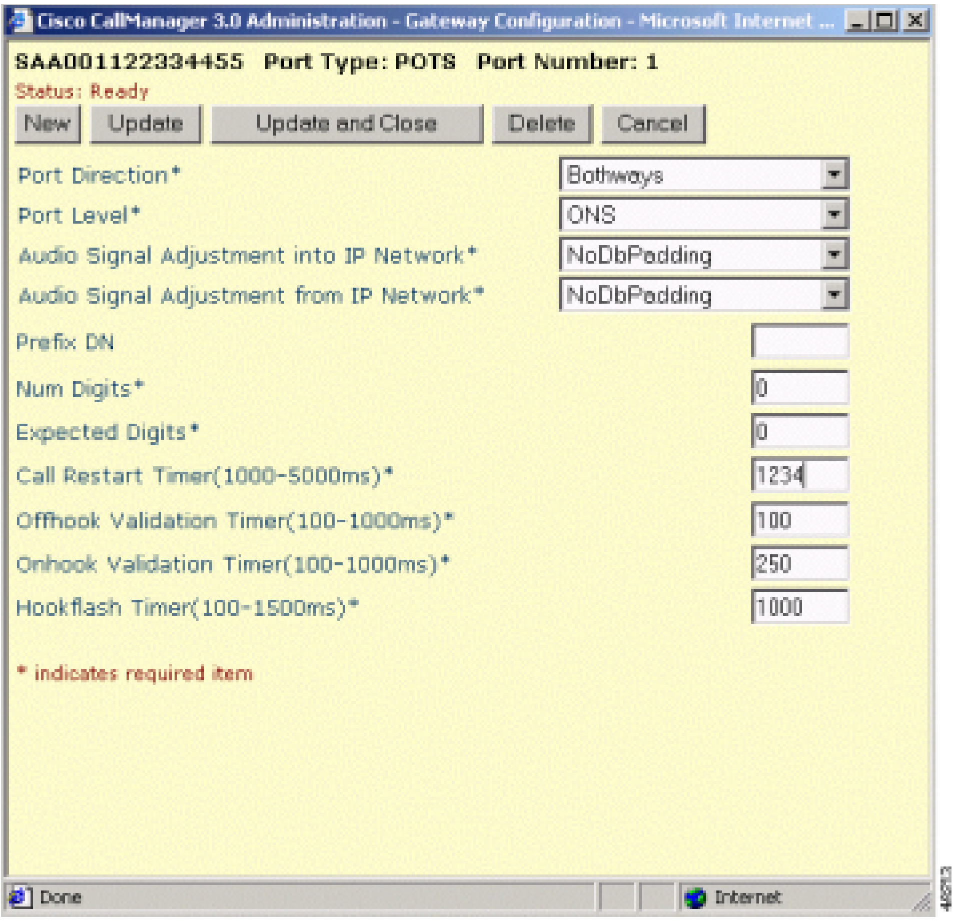
Cisco CallManager 3.0 Administration - Gateway Configuration - Microsoft Internet ...
New
Update
Update and Close
Delete
Cancel
Done
Internet
48253

SAA001122334455 Port Type: POTS Port Number: 1

Status: Ready

Port Direction*	Bothways
Port Level*	ONS
Audio Signal Adjustment into IP Network*	NoDbPadding
Audio Signal Adjustment from IP Network*	NoDbPadding
Prefix DN	
Num Digits*	0
Expected Digits*	0
Call Restart Timer(1000-5000ms)*	1234
Offhook Validation Timer(100-1000ms)*	100
Onhook Validation Timer(100-1000ms)*	250
Hookflash Timer(100-1500ms)*	1000

* indicates required item

10. Repeat this configuration change for *all* Ports.
11. Click on **Reset Gateway** in order to have the changes take affect.

Integrating SMDI Voice Mail Over IP WAN

With SMDI integration over the IP WAN, call information is transmitted over an RS-232 serial link between Cisco CallManager and the asynchronous port of the router located at one end of the IP WAN cloud. Once this call information has reached this router, asynchronous tunneling carries the call information to the router located at the other end of the IP WAN cloud. Once this call information has reached this router, it is transmitted over an RS-232 serial link between the asynchronous port of the router and the voice messaging system.

Voice communications are provided via a separate path created by a hunt group consisting of analog stations (FXS ports derived from either the VG200 gateway(s) or Catalyst 6000 24-port FXS Gateway(s) WS-X6624) between Cisco CallManager and the voice messaging system. The VG200 gateway(s) or Catalyst 6000 24-port FXS Gateway(s) WS-X6624, as well as the voice messaging system, needs to be co-located (not over the IP WAN cloud). When the hunt group receives an incoming call over the IP WAN, it is accompanied by a *digital message* in standard SMDI format over the IP WAN cloud from Cisco CallManager, which contains all necessary call information. The voice messaging system then answers the call on the specified port and plays the appropriate greeting.

RS-232 limitations in traditional SMDI integration cause the voice messaging system and CallManager to be co-located. In several customer scenarios, the IP telephony solution is being installed in small remote sites and the voice messaging system is located at the central site traversing the IP WAN cloud.

By using asynchronous tunneling between the IP WAN, routers located on the remote and central sites allow these RS-232 signals to be tunneled. The RS-232 cable coming out of the SUMI interface of the CallManager located at site A (see Figure 4-66) needs to be connected to the asynchronous port of router A located at site A. You can use any asynchronous port (AUX port is also an asynchronous port) to achieve this goal, depending on the model of router. On the other side of the IP WAN cloud, the RS-232 cable needs to be connected between the asynchronous port (AUX port is also an asynchronous port) of router B located at site B and the SMDI interface of the Octel voice messaging system. The Octel voice messaging system is located at site B.

Before implementing this solution, you should already have configured the appropriate QoS (LLQ) for the number of calls that can roll to voice mail, as well as the capacity required for inter-site calls. Also, you must implement an admission control scheme before implementing this solution.

See Integrating SMDI Voice Mail for more information.

Topology

The Asynchronous Tunneling Configuration is based on the following topology:

Figure 4-66 CallManager SMDI Topology

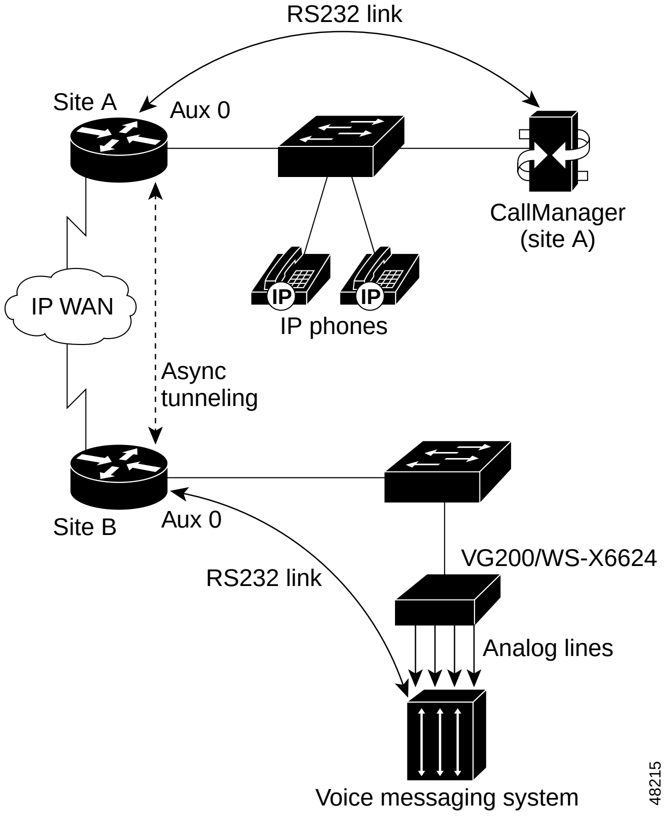

Asynchronous Tunneling Configuration

You need the following asynchronous tunneling configuration to carry the call information over the IP WAN cloud:

Async tunneling configuration on the router A located on site A:

```
! On the router A serial link <10.1.2.3> define an IP hostname to use on the TELNET so we
can use !BUSY-MESSAGE to suppress display message
ip host router_b  4129 10.3.2.1                 ! port 4xxx is raw TCP, 129 is the line
number for aux 0 !on the router B
busy-message router_b  # #
service tcp-keepalives-out                                    ! [2]
!
line 2
no motd-banner
no exec-banner
no vacant-message
autocommand telnet router_b  /stream
autohangup
! The following command means incoming serial data is saved till the TCP
! connection is made.
no flush-at-activation                                            ! not available in all
feature sets
no activation-character                                             ! any character will
create the EXEC                escape-character NONE
! or "escape-char BREAK"
exec                                                                     ! need an EXEC
to do the TELNET
special-character-bits 8
exec-timeout 0 0
session-timeout 0 0
! RS232 config:
no modem inout                                                     ! disable modem
control [1]
no autobaud
speed 9600                                                            ! set the desired
speed
stopbits 1                                                              ! or 2, as
desired
flowcontrol NONE                                                ! or HARDWARE, or
SOFTWARE
transport input NONE                                               ! do not allow
reverse connections
```

Async tunneling configuration on the router B located on site B:

```
! On router_b Ethernet link <10.3.2.1>
no banner incoming
service tcp-keepalives-in                                         ! [2]
line 3
no exec
no exec-banner
no vacant-message
! RS232 config:
modem DTR-active                                                      ! DTR indicates
status of TCP
                                                                                    !
connection
no autobaud
speed 9600                                                       ! as desired - does not
need to match
                                                                        ! the speed on
the called side
```

```
stopbits 1                                                           ! or 2, as desired
flowcontrol NONE                                             ! or HARDWARE, or SOFTWARE
transport input telnet                                              ! allow the incoming
TCP connection
```

Note Use the **no modem inout** command on router A. With modem signaling, if the router sees DSR goes high, it will initiate the autocommand. However, if the router is power cycled and if DSR is high when the router comes up, the autocommand will not be initiated you initiate a **clearline**.

Note Be sure that TCP keepalives are enabled on both sides for the appropriate connection. Otherwise, if site A, or the network path, were to go down, site B (unless it has application data to send) will be unaware that site A's connection has gone away. This causes the new site A connection attempt to fail.

Tips

We used the Octel 250 voice messaging system in our testing. However, this solution works with any voice messaging system that supports SMDI. Follow the tips below when using the Octel voice messaging system:

- Define your available asynchronous port as an Integration port (menu 6.3). You need to restart the system after defining the CPU serial channel as an integration link.
- Configure your Integration link (menu 6.5) as Switch type 3 (1AESS/SMDI, full duplex). Configure it for 9600 bps, 7,E,1.
- Create mailboxes for the IP phones.

On the Cisco CallManager side, make sure CMI is configured and running. Also make sure that the CallManager sends the same number of digits for DNs that the Calista PBXLink devices (if Calista PBXLink is used in the topology) are sending (the DN length usually matches the mailbox length). Finally, make sure that the serial link is set for 9600, 7,E,1.

The Octel Aria 250 supports up to three serial ports for integration. If two of these serial ports are currently in use by the PBX Integration Devices (PIDs), then we have one spare that we can use for CallManager. You must first convert the current Integration to SMDI by hanging out the PIDs and replacing them with Calista PBXLinks. Once the Octel Aria 250 is integrating via SMDI, we can connect to the Cisco CallManager.

You may need a non-DID (direct inward dial) extension on the PBX, which will have a coverage path to the voice messaging system. Voice messages will be delivered to correct voice mailboxes and users will be able to retrieve them. In this scenario, if the WAN goes down, users will still be able to leave and retrieve voice messages. However, they will not get a message-waiting indicator (MWI).

Note You could use an ISDN backup solution if the IP WAN goes down. However, you need to carefully evaluate the limitations associated with this backup solution.

Migrating to an IP Telephony Network

http://www.cisco.com/univercd/cc/td/doc/product/voice/ip_tele/network/dgmigrt.htm

CHAPTER 5

Implementing the IP Telephony Network

In this Chapter

This chapter consists of the following sections:

- Preparing for Implementation
- Conducting the Site Survey
- Determining Site Requirements
- Validating Implementation Readiness
- Implementing the Solution
- Implementing a Migration Strategy
- Solution Implementation Acceptance Testing
- Post-implementation Documentation
- Case Study

Related Information

Visit the following sites for information closely related to this chapter:

- Cisco Traffic Engineering Technical Tip

  http://www.cisco.com/warp/public/788/pkt-voice-general/9.html
- Westbay Engineers Limited Home Page

  http://www.erlang.com/
- Cisco White Paper: Designing High-Performance Campus Intranets with Multilayer Switching

  http://www.cisco.com/warp/public/cc/so/cuso/epso/entdes/highd_wp.htm
- Cisco IP Telephony Network Design Guide

  http://www.cisco.com/univercd/cc/td/doc/product/voice/ip_tele/index.htm
- CallManager 3.0(1) Troubleshooting Guide

  http://www.cisco.com/warp/public/788/AVVID/ts_ccm_301_sec1.htm

- E-911 Fact Sheets
  http://www.fcc.gov/e911
- National Emergency Number Association website
  http://www.nena9-1-1.org/

Preparing for Implementation

While each customer's implementation is unique, properly preparing for each implementation is crucial to the success of every deployment.

General Site Information

This section documents general information about the customer site. The shipping information is always gathered from customer facility management. If union workers operate the site delivery service, make arrangements to coordinate with the facility manager.

Table 5-1 should be completed prior to implementing the solution.

Table 5-1 General Site Information

Shipping and Receiving
Specify any special shipping address, delivery times, and shipping instructions: Street: City and State: Zip Code and Country: Special Instructions:
Specify the hours of operation of the shipping and receiving department for this location:
Specify the secured and insured storage location for the Cisco IP Telephony equipment, telecom equipment, and cabling while it is on site prior to installation.
Is there a loading dock at this location? Yes or No
Is there a pallet jack available at this location? Yes or No
Is there elevator access to all IP Telephony equipment rooms? Yes or No
Are the receiving department and freight elevators available after hours and on weekends? Yes or No

Table 5-1 General Site Information (continued)

If an equipment cabinet (maximum approximate size 7 ft. x 2 ft. x 3 ft.) is required, are there any obstructions such as small stairways, elevators, doorways, and lack of equipment ramps? If so, please describe:
Building/Site Access
Specify the building address and room number where the IP Telephony equipment will be installed: Street Address: Room Number: City and State: Zip Code and Country:
Specify days and hours of normal operation.
Specify any special building access procedures during *normal hours* of operation.
Specify any special building access procedures during *after hours* of operation.
Specify required security procedures and policies that must be followed.
Specify any special Cisco IP Telephony equipment room access procedures such as keys, badges, or escorts.
List the room and building passes, badges, and/or escorts that are required to access the Cisco IP Telephony equipment rooms.

Table 5-1 General Site Information (continued)

Who can arrange for room and building passes, badges, and or escorts, and how can this person be contacted?

Name:

Phone:

Pager:

Email address:

Instructions:

Site Preparation

Who is the site coordinator responsible for ensuring the site preparation, and how can this person be contacted?

Name:

Phone:

Pager:

Email address:

List the telephone numbers nearest the Cisco IP Telephony equipment to be used by the installed if necessary.

CallManager #1:

CallManager #2:

CallManager #3:

CallManager #4:

CallManager #5:

Catalyst Switch #1:

Catalyst Switch #2:

Gateway #1:

Gateway #2:

Gateway #3:

Is union labor required to install, maintain, or support the Cisco IP Telephony equipment? Yes or No

Specify any special dates, times, and locations that work must be performed (quarter end, after hours, weekends).

Table 5-1 General Site Information (continued)

Is there any equipment that needs to be removed from the installation location prior to the Cisco IP Telephony installation? Yes or No If so, indicate the responsible person and date by which the equipment will be removed: Name: Phone: Pager: Email address: Date:
Where can the installation crew dispose of equipment packing materials?
Do all equipment locations have air conditioning and heating equipment that is operational and climate controlled? Yes or No

Team Information

This section provides contact information about the members of the implementation team. This information is necessary because the customer will manage the operation at the end of the implementation cycle. Cisco Systems or the Cisco Partner project team should work closely with the customer's project team for the solution knowledge transfer. The transfer of knowledge increases customer satisfaction and reduces the amount of support needed in the future.

Table 5-2 lists the contact information for the members of the implementation team.

Table 5-2 Implementation Team Contact Information

Cisco Systems/Partner Project Team Member	Customer Project Team Member
Project Manager: Telephone: Email Address:	Project Owner: Telephone: Email Address:
Project Engineer: Telephone: Email Address:	Project Manager: Telephone: Email Address:
Solution Design Engineer: Telephone: Email Address:	Systems Engineer: Telephone: Email Address:
Account Manager: Telephone: Email Address:	
Systems Engineer: Telephone: Email Address:	

Project Management

Cisco project management should establish a single point of contact for the customer during the implementation. The responsibilities of the project manager include:

- Working with the facilities engineers, integration consultants, and other specialists.
- Coordinating the arrival of deliverables from multiple locations with engineers and installers.
- Overseeing communication between all parties to keep the implementation on schedule and according to specifications.

The project manager should provide an implementation task schedule. Table 5-3 is an example of an implementation schedule for one core site and one remote site for an IP Telephony implementation. The number of days required for each task can vary based on the size of the project.

Table 5-3 Implementation Schedule Example

Implementation Task	Days Required	Start Date Schedule	Completion Date Schedule
Implementation Preparation	5		
Site Survey	5		
Pre-implementation Check	10		
Implementation	10		
Acceptance Tests	2		
Post-implementation Documentation	5		
Customer Acceptance	2		

Implementation Considerations

This section provides information for the implementation team to verify that the following implementation considerations have been met.

Assumptions

In order for this implementation to take place, the following assumptions have been made:

- All sites are ready for equipment installation and the customer has ensured that any power, air conditioning, circuit installation, or other work has been completed prior to the installation team's arrival.
- The customer has ensured that all live circuits that are due to connect to the equipment have been fully tested and subsequently proven to be suitable to carry network traffic.
- The customer has taken delivery of all equipment and has installed power rails and circuit breakers in cabinets where required, connected power to them, and tested the power to ensure it is capable of supplying the equipment.
- During implementation, it is assumed that one or more Cisco representatives or implementation engineers will be on site and that a customer representative will be present at all times.

Safety Information

Consult the product literature shipped with each Cisco product for the most current source of safety guidelines. The following links provide safety information regarding the IP Telephony implementations:

- Catalyst 3500

  http://www.cisco.com/univercd/cc/td/doc/product/lan/c2900xl/index.htm
- Catalyst 4000

  http://www.cisco.com/univercd/cc/td/doc/product/lan/cat4000/index.htm
- Catalyst 6000

  http://www.cisco.com/univercd/cc/td/doc/product/lan/cat6000/index.htm
- Cisco CallManager

  http://www.cisco.com/univercd/cc/td/doc/product/voice/c_callmg/index.htm

Implementation Tool Requirements

The Cisco implementation engineer(s) should have the following tools:

- PC with VT100 emulator, 10BASE-T interface, FTP server, TFTP client applications
- Console port cable DB9-RJ45/DB25
- Ethernet transceiver
- 10BASE-T Ethernet cable

Conducting the Site Survey

Site readiness is crucial to rapid deployment since site deficiencies may delay solution implementations. The first step to ensuring readiness is a site assessment that requires a skilled person to perform an on-site survey. Once the survey is complete, the project manager can decide how to implement the solution. A ready site means no surprises at the time of deployment.

This section provides key information, including a gap analysis highlighting the differences between what currently exists at the customer site and what is required to implement the solution.

Site Survey Tables

The first steps to conducting the site survey is to populate the base information tables, then populate the tables that refer to the base information. In practice, you will have to revisit the tables as the site survey progresses.

Table 5-4 lists the tables required to be populated for a successful site survey.

Table 5-4 Site Tables

Base Information Tables	Reference Information Tables	Order-independent Information Tables
Access	Device	Notes
Contacts	Equipment order	IP address range

Table 5-4 Site Tables (continued)

Base Information Tables	Reference Information Tables	Order-independent Information Tables
Features	Intersite circuits and intersite virtual circuits	Documentation
Procedure	Maintenance support	
Service organization	Site	
Room	Site group	
User		

Gathering Site Information

This section provides a brief description of the information to be gathered for the site survey. Each of the site survey tables described below are available at the following location: http://www.cisco.com/warp/public/788/solution_guide/forms/index.html#ss.

- General Site—Record the following general site information during the initial survey.
  - Site Groups—Associate random groups of sites with a site group name. This allows a group of sites to be referred to by one name. These site group names can then be associated (in the tables) to other entities such as contacts persons. Use this table to create your own site groups as needed.
  - Project Contacts—Capture information about people who will need to be contacted during the course of the project for information or assistance.
  - Rooms—Record office, meeting, or equipment space rooms. The phone information is for capturing information on phones that may be replaced with IP phones, in a potential project meeting space, or near equipment to be worked on during the project.
  - Telephones—Capture existing telephone information.
  - IP Addresses—Record the IP address ranges in use at or allocated to the site.
  - Documentation—Record any documentation captured from the customer or during the site survey.
- Individual Equipment—Capture information on communication devices related to the deployment of the Cisco IP Telephony solution. This information should be captured during the site surveys. Be sure to capture information on the following component types: routers, LAN switches, WAN switches, PBXs, voice mail systems, ACDs, IVRs, CSU/DSUs, multiplexers, and so forth.
- Sales Order and Maintenance Support—Record sales order and support maintenance information during the initial survey and site survey.
- Site Access and Procedure Requirements—Include information about special site access and procedural requirements. Examples include requirements to check in with a guard, x-ray gates, security clearances, equipment sign in/out, and so on. This information should be gathered during the initial survey and site survey.
- User Services and Features—List and describe all existing features to maintain under the Cisco IP Telephony solution. This information should be gathered during the initial survey and site survey.
- Inter-site Communications—Capture information on circuits (physical communications path) and links (virtual circuit that is carried over a physical circuit). For leased lines, in most cases, there are no virtual circuits. The majority of this information should be collected during the site survey.

- Service Organizations—Store information concerning outside organizations that provide services to the customer on a site or site group basis. This information should be gathered during the initial survey.
- Site Survey Notes—Enter notes that clarify any ambiguous information in the site survey or record information needed, but is not listed in any of the above mentioned tables.
- Documentation—List all types of documentation, such as floor plans and diagrams, that relate to the site.

Determining Site Requirements

This section provides information about the network infrastructure, telephony infrastructure, IP Telephony, power, and environmental requirements.

LAN Requirements

Table 5-5 lists the types of LAN switches required for the IP Telephony solution.

Table 5-5 LAN Switch Requirements

Role	Required Feature	Reference Platforms
Campus Access Switch	Inline power Multiple queue support 802.1p and 802.1q Fast link convergence	Catalyst 3500 Catalyst 4000 Catalyst 6000
Campus Distribution/Core Switch	Multiple queue support 802.1p and 802.1q Traffic classification Traffic reclassification	Catalyst 6500
Branch Office Switch	Inline power Multiple queue support 802.1p and 802.1q	Catalyst 3500 Catalyst 4000

Table 5-6 lists the capabilities of each type of switch.

Table 5-6 Switch Capabilities

Capabilities	Catalyst 3500	Catalyst 4000	Catalyst 6000/6500
Ability to Trust	Yes	No	Yes
Re-Classify CoS	Yes	Yes	Yes
Re-Classify ToS	No	No	Yes
Congestion Avoidance (WRED)	No	No	Yes
Priority Queue	No	No	Yes
Multiple Queues	Yes	No	Yes

Table 5-6 Switch Capabilities (continued)

Congestion Management (WRR)	No	No	Yes
Policing	No	No	Yes
10/100MB Switch with In-line Power	24 Port	WS-X4148-RJ45V (48-port module)	WS-X6348-RJ45V (48-port module)
Analog Trunk Gateway	—	WS-X4604-GWY [1]	WS-X6624-FXS (24-port module/RJ21)
Digital Trunk Gateway	—	WS-X4604-GWY [1]	WS-X6608-T1 or E1 (8-port module/RJ48)

[1] Cisco IOS based. Functions as router, voice gateway, and DSP Farm. Shares VIC/WIC interfaces with Cisco 2600/3600, voice, and fax support.

WAN Requirements

The IP Telephony installations that adopt the centralized call processing deployment model are limited to a hub-and-spoke topology. This is due to the fact that the location-based call admission control (CAC) mechanism only records the available bandwidth in and out of each location.

WAN Hardware Requirement

Table 5-7 lists the WAN hardware requirements.

Table 5-7 WAN Hardware Requirements

Role	Required Features	Reference Platforms
WAN Aggregation Router	Multiple Queue Support Traffic Shaping LFI Link Efficiency tools Traffic Classification Traffic Re-classification 802.1p and 802.1q	7200 7500
Branch Router	Multiple Queue Support LFI Link Efficiency tools Traffic Classification Traffic Re-classification 802.1p and 802.1q	2600 3600

The WAN connectivity options include:

- Leased lines
- Frame relay
- ATM
- ATM/Frame Relay Service Inter-Working (SIW)

WAN Bandwidth Requirements

Use Table 5-8 to calculate the available bandwidth for the IP Telephony implementation.

Table 5-8 WAN Bandwidth Requirements

Bandwidth	Max. Trunks	Max. Users	Available Bandwidth	Speech Bandwidth	Skinny Bandwidth	Reserved Bandwidth	CallManager Bandwidth
64	1	6	48	22	3	23	24
128	3	12	96	66	6	24	72
192	5	19	144	110	9	25	120
256	7	25	192	154	12	26	168
320	9	32	240	198	16	26	216
384	10	38	288	220	19	49	240
448	12	44	336	264	22	50	288
512	14	51	384	308	25	51	336
576	16	57	432	352	28	52	384
640	18	64	480	396	32	52	432
704	19	70	528	418	35	75	456
768	21	76	576	462	38	76	504
832	23	83	624	506	41	77	552
896	25	89	672	550	44	78	600
960	27	96	720	594	48	78	648
1024	28	102	768	616	51	101	672
1088	30	108	816	660	54	102	720
1152	32	115	864	704	57	103	768
1216	34	121	912	748	60	104	816
1280	36	128	960	792	64	104	864
1344	37	134	1008	814	67	127	888
1408	39	140	1056	858	70	128	936
1472	41	147	1104	902	73	129	984
1536	43	153	1152	946	76	130	1032

Quality of Service Requirements

Quality of service (QoS) mechanisms, such as priority queuing and traffic shaping, are required on the routers that sit at the WAN edges. QoS mechanisms protect the voice traffic from the data traffic across the WAN, where bandwidth is typically scarce. Additionally, a call admission control scheme is required to avoid over-subscribing the WAN links with voice traffic (thus deteriorating the quality of established calls). For the IP Telephony deployment model, this is done using the locations-based call admission control scheme within Cisco CallManager.

Traffic shaping is required for two main reasons:

- To remain within the contracted traffic agreement with the ATM or Frame Relay network to avoid being policed and incurring dropped packets.
- To maintain comparable traffic speeds between sites linked to the Frame Relay or ATM network by different line speeds. For example, the headquarters site may have a DS-3 and the other sites may have a DS-1. Traffic shaping helps prevent buffer overrun within the network, which would result in packet loss.

Remote branches can be provided with PSTN access through a variety of Cisco gateways. When the IP WAN is down, or all of the available bandwidth on the IP WAN has been used, the users at the remote branch can dial the PSTN access code and place their calls through the PSTN.

QoS tools are essential to successfully run voice and data over the slower links that are found in the WAN. The WAN QoS techniques depend on the speeds of the links. At higher speeds (above 768 Kbps), low latency queuing (LLQ) for voice is required to reduce jitter and possible loss during a burst of traffic that might over-subscribe a buffer. This is similar to the LAN infrastructure scenario.

At lower link speeds, other techniques known as link efficiency techniques are needed to minimize the effects of serialization delays.

Table 5-9 summarizes all the features and tools that need to be enabled to support IP Telephony, according to the WAN technology and link speed. The remainder of this section highlights some of the most important techniques to keep in mind when designing a WAN that supports voice and data traffic.

Table 5-9 Feature and Tool Requirements

WAN Technology	Link Speed 56kbps – 768kbps	Link Speed > 768kbps
Leased Lines	LFI (Multilink PPP) LLQ cRTP	LLQ
Frame Relay	Traffic Shaping LFI (FRF.12) LLQ cRTP	Traffic Shaping LLQ
ATM	Traffic Shaping TX-ring buffer changes LFI (Multilink PPP) LLQ	Traffic Shaping TX-ring buffer changes LLQ
Frame Relay/ATM SIW	Traffic Shaping TX-ring buffer changes LFI (Multilink PPP) LLQ	Traffic Shaping TX-ring buffer changes LLQ

Gateway Requirements

In an IP Telephony solution implementation, a gateway interconnects the IP connection between the central office and remote offices. The following tables list the gateway information as organized by gateway protocols, gateway analog interfaces, and gateway digital interfaces.

Table 5-10 Gateway Protocols

Gateway	MGCP	H.323v2	Skinny Station Protocol	Platform
VG-200	Yes	Yes	No	Stand Alone
DT-24+ or 30+	Future	No	Yes	Stand Alone
Cisco 1750	No	Yes	No	IOS—Router
Cisco 3810 V3	Yes	Yes	No	IOS—Router
Cisco 2600	Yes	Yes	No	IOS—Router
Cisco 3600	Yes	Yes	No	IOS—Router
Cisco 7200	No	Yes	No	IOS—Router
Cisco 7500	No	Future	No	IOS—Router
Cisco 5300	No	Yes	No	IOS—Router
Catalyst 4000 WS-X4604-GWY	Future	Yes, for PSTN interfaces	Yes, for conferencing and MTP transcoding	Switch
Catalyst 6000 WS-X6608-T1 or E1	Future	No	Yes	Switch

Table 5-11 Gateway Analog Interfaces

Gateway	FXS	FXO	E&M	Analog DID/CLID
VG-200	Yes	Yes	Yes	Yes
DT-X+	No	No	No	—
Cisco 1750	Yes	Yes	Yes	Future
Cisco 3810 V3	Yes	Yes	Yes	Yes
Cisco 2600	Yes	Yes	Yes	Yes
Cisco 3600	Yes	Yes	Yes	Yes
Cisco 7200	No	No	No	—
Cisco 7500	No	No	No	—
Cisco 5300	No	No	No	—
Catalyst 4000 WS-X4604-GWY	Yes	Yes	Yes	Yes
Catalyst 6000 WS-X6608-x1	Yes	No	No	No/Yes

Table 5-12 Gateway Digital Interfaces

Gateway	T1 CAS	E1/R2	E1 CAS	User Side PRI	Network Side PRI	User Side BRI	Network Side BRI	Digital DID/CLID
VG-200	Yes	Yes	Yes	Yes	Yes	Yes	Yes	—

Table 5-12 Gateway Digital Interfaces (continued)

DT-24+/30+	Future	No	No	Yes	Yes	No	No	Yes
Cisco 1750	No	No	No	No	No	Future	Future	—
Cisco 3810 V3	Yes	No	Yes	No	No	Yes	No	Yes
Cisco 2600	Yes	Yes	Yes	Yes	Yes	Yes	Future	Yes/Yes
Cisco 3600	Yes	Yes	Yes	Yes	Yes	Yes	Future	Yes/Yes
Cisco 7200	Yes	Yes	Yes	Yes	Yes	No	No	Yes/Yes
Cisco 7500	Yes	Yes	Yes	Yes	Yes	No	No	Yes/Yes
Cisco 5300	Yes	Yes	Yes	Yes	Yes	No	No	Yes/Yes
Catalyst 4000 WS-X4604-GWY	Yes	Yes	Yes	Yes	Yes	Future	Future	Yes/Yes
Catalyst 6000 WS-X6608-x1	Future	No	No	Yes	Yes	No	No	Yes/Yes

IP Telephony Requirements

The IP Telephony infrastructure can be built on top of an existing solid network infrastructure. The IP Telephony infrastructure includes the following devices and functionalities:

- Call processing agents
- Call admission control mechanisms for calls that traverse the WAN
- Voice and fax/modem gateways
- Transcoding and ad-hoc conferencing resources
- End-user devices

Table 5-13 lists the Cisco-supported platform(s) for each of the devices/functionalities.

Table 5-13 Cisco Platforms

Device/Functionality	Available Platform
Call processing agents	Software: Cisco CallManager Hardware: MCS-7825-800, MCS-7835-1000, ICS-7750 IOS-based survivable remote: Vespa (IOS 12.1(5)-YD)
Call admission control mechanisms	Centralized call processing model: CallManager Distributed call processing model: Gatekeeper
Voice gateways	H.323: 1750, 2600, 3600, AS5300, 7200, 7500 Skinny: DT-24+, DE-30+, Catalyst T1 and E1 modules MGCP: VG-200, 2600, 3600, IAD2400

Table 5-13 Cisco Platforms (continued)

Fax/Modem gateways	H.323: 2600, 3600, AS5300 MGCP: 2600, 3600, AS5300, IAD2400 Skinny: Catalyst T1/E1 and FXS/FXO modules
Transcoding/Conferencing resources	Catalyst T1/E1 and FXS/FXO modules
End-user devices	Cisco IP Phone 7910 Cisco IP Phone 7940 Cisco IP Phone 7960 Cisco SoftPhone Cisco IP Conference Station 7935

Table 5-14 lists the technical specifications for the Media Convergence Servers.

Table 5-14 Media Convergence Server Technical Specifications

Technical Specifications	MCS 7825-800	MCS 7835-1000
Intel PIII processor	800 MHz	1 GHz
Memory (SDRAM)	512 MB	1 GB
SCSI hard drive (HD)	Single 20GB Fast ATA (7200 RPM)	Dual 18.2GB Ultra3 SCSI (10,000 RPM)
3.5” floppy drive	1.44MB floppy drive	1.44MB floppy drive
CD-ROM drive	24x IDE CD-ROM drive	24x IDE CD-ROM drive
Video controller	ATI Rage XL – 4MB	ATI Rage IIC – 4MB
RAID controller	No	Integrated Smart Array RAID Controller—RAID 0/1 Disk Mirroring
Hot-plug hard drive	No	Yes
Hot-swap redundant PS	No (180W)	Yes (275W)
Rack mount size (1U=1.75”)	1U	3U
Controller module	Integrated Ultra ATA/100 controller	Integrated single-channel wide-ultra SCSI-3 adapter
10/100 PCI UTP controller (embedded)	Dual Port	Single port
Tape drive	—	12/24 GB DAT tape drive (optional)
Software Compatibility		
CallManager 3.0.x	Yes	Yes
IP IVR/Auto Attendant	Yes	Yes
Unity voice mail	No	Yes
IPCC/ICM server	No	Yes
IPT Personal Assistant server	No	Yes
Flexibility		
IP IVR/Auto attendants ports	30	48

Table 5-14 Media Convergence Server Technical Specifications (continued)

CCM 3.0 phone support	500 lines	2500 lines
IP IVR/Auto attendants ports	30	48

Table 5-15 lists the typical hardware and software requirements for IP Telephony solution devices.

Table 5-15 Hardware and Software Requirements

Product	Recommended Release
Catalyst Switch (4000, 6x00)	Catalyst OS 5.5 or higher
Catalyst Switch (35xxXL)	Cisco IOS 12.0(5)XU or higher
Cisco Router (17xx, 26xx, 36xx, 72xx, 75xx)	Cisco IOS 12.1(5)T or higher
Media Convergence Server	CallManager 3.0(5) or higher

Power and Environmental Requirements

Table 5-16 lists the power consumption and heat dissipation requirements for Cisco IP Telephony devices.

Table 5-16 Power Consumption and Heat Dissipation Requirements

Item	Input VA Rating	Heat (BTU)
Cisco 7204	370	
Catalyst 5509	1000	
Cisco 7206	370	1262
MCS-7835	432	1475
Catalyst 6006	1800	6140

Validating Implementation Readiness

Prior to configuring the implementation, a highly skilled and specialized technical Cisco representative performs a design review to confirm that high- and low-level designs meet the customer's requirements. The design review may result in recommendations for improvement.

Document the results of the design review, along with the customer readiness check information. Complete the pre-implementation checklist, Table 5-17, before implementing the solution.

Table 5-17 Pre-implementation Checklist

Component	Readiness	Note
Network topology analysis		
Voice network analysis		

Table 5-17 Pre-implementation Checklist

Data network analysis: LAN requirement		
Data network analysis: WAN requirement WAN bandwidth requirement		
IP Telephony requirement analysis		
Solution implementation templates		
Customer ordered equipment		
Customer premise equipment (CPE) interfaces		
Customer site readiness		

Solution Design Review

The design review verifies that the design and the configuration meets the customer's expectations. If any modification to the design is necessary, it should be done before it impacts the network rollout.

The solution design review includes:

- Network topology analysis
- Voice network analysis
- Data network analysis
  - LAN requirement analysis
  - WAN requirement analysis
  - WAN bandwidth requirement analysis
  - IP Telephony requirement analysis

Network Topology Analysis

The first step of the solution design review is to gather the necessary information for the voice and data network diagrams. Include the following components (at each site) for the voice network diagram:

- Termination point for phones (PBX/Centrex/Key system)
- Number of trunks between termination point and PSTN network
- Number of trunks (private or VPN) to other sites
- Number of trunks to local voice mail system, if applicable
- Number of sets (users)

Include the following components (at each site) for the data network diagram:

- Routers and links (with bandwidth) to distant sites
- Ethernet switches at the site
- Ethernet hubs at the site
- Number of attached devices

Voice Network Analysis

Voice trunks are used for digital and analog gateway products. The voice network analysis determines the optimal number of voice trunks required to support a specific function. Trunk types include two-way trunks, Direct Inward Dial (DID) trunks for inbound calls only, and Direct Outward Dial (DOD) trunks for outward dial trunks. The assessment determines the capacity for three trunk types in three different applications as follows:

- PSTN trunks—Used in gateways to the PSTN
- Voice mail trunks—Used for connectivity to voice mail servers
- Site-to-site trunks—Represent WAN connectivity and may initially use a different gateway and then use WAN VoIP at a later date

Table 5-18 shows a sample table to complete for each site to account for the current trunk usage. It is common for some fields in the table to contain a zero value.

Table 5-18 Current Trunk Usage

Trunk Type	Two-way Calling	DID Trunks	DOD Trunks
Voice mail trunks			
PSTN trunks			
Trunks to site x			
Trunks to site y			
Trunks to site z			

To determine trunking requirements, complete a traffic study. Three applicable methods are listed below in the order of desirability:

1. Conduct a traffic analysis that shows the busy hour traffic (BHT) value for each trunk type, location and application. Most telephone companies and PBX vendors can readily produce such a report for a trunk group over a specified period. If, for example, two-way voice mail trunks at site A show two BHT, this indicates that the two-way voice mail trunks experience 120 minutes of call volume during the busy hour. The traffic analysis should ideally be done over a peak period of telephone usage, but can be done for any random interval if it is believed that the telephony traffic volume does not vary widely on a seasonal basis. This method is the preferred method for data collection.
2. Collect call data records (CDRs), which you can examine to determine approximate BHT values. This method varies based on the level of detail available in the call record and can take a significant amount of time. Since it is less precise and less straightforward, this technique is usually not recommended except as a last resort. Failed call attempts and call overhead (ringing, busy, dialing) estimates need to be considered.
3. Look at the existing trunking and ask the customer about their level of satisfaction. If the customer is currently satisfied with the trunk quantity, Cisco can simply allocate a similar amount of gateway trunks for the Cisco IP Telephony solution. More trunks can be added as desired. When estimating

the load generated by an existing trunk group where no traffic study has been done, take the number of trunks and a blocking factor of .01 to arrive at a load estimate in Erlangs (using Erlang-B). This number can be used to provide a "best guess" estimate for the load that will be introduced into the data network if a gateway were servicing those trunks.

When you obtain the BHT value, the assessor can use the Erlang-B calculator at http://www.erlang.com or through any number of commercial packages or Erlang-B tables to understand trunking requirements. Cisco generally recommends a blocking factor of .01 or 1 percent; however, the customer may wish to alter this for the assessment. Find out what blocking factor the customer requires, given that Cisco Systems recommends 1 percent. The recommended value and the customer desired value could be included in the assessment tables for trunk analysis.

In order to populate the tables in the report, the customer-provided load factor, plus the desired blocking factor (.01 recommended), can be inserted into the Erlang-B formula to produce the number of trunks required. This number can be compared to the customer's current trunking.

Data Network Analysis

Data network analysis includes:

- LAN Requirement Analysis
- WAN Requirement Analysis
- WAN Bandwidth Requirement Analysis

LAN Requirement Analysis

- Phones and a maximum of one daisy-chained PC must be on one individual switched port. This minimizes the collision domain of the media and prevents performance degrading delay.
- Servers and gateways must be on individual switched ports to minimize the collision domain of the media and prevent performance degrading delay.
- Full-duplex should be used wherever possible.
- Normal best practices for LANs apply, including size of broadcast domains, traffic, utilization on Ethernet, collision rate, low-speed device connectivity, and so on. A good reference for campus LAN design is the *Designing High Performance Campus Intranets with Multilayer Switching* white paper. This document can be found at the following location: http://www.cisco.com/warp/public/cc/so/cuso/epso/entdes/highd_wp.htm.
- The links in the LAN topology should get faster as they move toward the core (no bottlenecks).
- Links connecting LAN switches should not be on shared media (hubs).
- Firewalls cannot be present in any proposed voice path over the data network, unless they have specific provisions to handle the voice protocols, which opens up a complex set of issues beyond the scope of this audit.
- Neither network address translation (NAT) nor port address translation (PAT) systems can be present in any proposed voice path over the data network, unless they have specific provisions to handle the voice protocols.
- The total aggregate offered load at each switch port with a phone should be less than 80 percent of the media capacity at that port over five-second intervals. This approximation is a guideline to prevent voice degradation.

- The total aggregate offered multicast load at each switch port should be less than 20 percent of the media capacity. This type of load is common only in networks with multiple video multicast streams and can be mitigated on individual switch ports by the use of multicast control schemes such as:
  - Cisco Group Management Protocol (CGMP)
  - GARP Multicast Registration Protocol (GMRP)
  - Internet Group Management Protocol (IGMP) snooping

WAN Requirement Analysis

Call admission control is needed to allocate bandwidth for a WAN deployment. In rare cases, a WAN may be used when admission control has not yet been implemented. In these cases, features such as queuing, traffic shaping, and fragmentation may be needed for optimization. In addition, the following capabilities may be necessary:

- IP RTP priority—12.0(5)T
- IP Precedence—11.2
- Differentiated services—12.0(6)T

Prioritization

- IP RTP priority—12.0(5)T
- WFQ—11.0
- CBWFQ—12.0(6)T

Traffic Shaping

- Frame relay traffic shaping—11.2
- ATM Traffic shaping—11.2

Fragmentation

- FRF.12-Frame relay—12.0(4)T
- MLPPP – Serial lines—12.0(6)T Fast switched

Efficiency

- Low speed Codecs—varies by Codec
- Compressed RTP—12.(1)T Fast switched

WAN Bandwidth Requirement Analysis

For estimating purposes, the bandwidth consumed by a call is as follows:

- G.711—80 kbps
- G.729—30 kbps
- G.723—18 kbps

The consumed bandwidth number can be reduced by approximately 10 to 40 percent if voice activity detection (VAD) is used. VAD can be turned on or off through a CallManager configuration.

The bandwidth estimates can vary based on encapsulation types of intermediate links (such as ATM versus PPP). The use of Real Time Protocol (RTP) header compression can further reduce the bandwidth needs on slow links at the cost of CPU time and delay on the devices doing the compression. The specific pros and cons of various bandwidth reduction schema are beyond the scope of this analysis since these configurations will need to be made on intermediate routers on the WAN, not on the CallManager or IP phones. Further, it may be necessary to update software, memory, or hardware on these intermediate routers in order to get adequate QoS capability to reliably carry voice.

Note If the QoS capabilities of the WAN routers are not known, an audit of those capabilities may be necessary; however, such an audit is not part of this analysis. If a network audit is required, please contact your Cisco project manager or account manager.

The formula for recommended bandwidth for data networks should be: The number of trunks that would have been required to carry the call load (see the "Voice Network Analysis" section on page 5-18) times the expected bandwidth required per call.

TRUNKS x EXPECTED BANDWIDTH REQUIRED PER CALL = RECOMMENDED BANDWIDTH

For example, if it is determined that six trunks are required between sites, and G.729 is the encoding technique, the bandwidth required would be approximately 30 kbits / second X 6 = 180 kbps to handle the busy hour load.

IP Telephony Requirement Analysis

The IP Telephony solution has the following characteristics:

- Compressed voice across the WAN
- Low bit-rate fax-relay requirement
- Use of advanced quality of service tools
- Compressed RTP for voice media streams

The Media Convergence Server (MCS) is the computing platform that runs the telephony applications supporting the IP phones. The applications include:

- Cisco CallManager
- Conference bridge
- Media transfer point
- Voice/unified messaging

The extent to which these applications can all coexist on a single MCS depends on the load presented to the combination of applications. An MCS can run all of these applications or a subset of applications.

The engineering guidelines and considerations for MCS are:

- Maximum number of IP phones per CallManager—2500
- Maximum number of CallManagers per Cluster—3
- Failover limitation—If two Cisco CallManagers are at the site, the failover would be approximately 90 seconds and could require manual replication of the programming on the backup Cisco CallManager.
- Power to the Cisco IP Phone 7960—Derived from either of two sources, the inline power 10/100 BASE-T Ethernet switching module, HYDRA, or from an external transformer.
- IP phone network connections—IP phones connect to a 10 Mbps switched or 10/100 Mbps switched Ethernet connection only.

- DHCP addressing for daisy-chained PCs—A PC or workstation that is connected to the IP Phone obtains its IP address from the same address pool (subnet) as the phone if DHCP is used for addressing.
- Possible voice degradation for daisy-chained IP phones—When a PC or workstation is connected to an IP phone so that they share a network connection, it is possible that the PC can cause some voice degradation in voice quality if both devices are being used at the same time with high traffic loads or IP Multicast applications such as IPTV. In practice, however, it is difficult to achieve this degradation with normal PC applications.
- No firewalls between IP Telephony devices—There can be no firewalls/NAT between any of the IP phone, gateway, or CallManager components.
- The network must be designed for average or higher use—The LAN/MAN network should be appropriately designed to handle expected data traffic bandwidth, deliver acceptable voice quality of service, and account for normally expected over-subscription. This guideline may dictate the LAN/MAN hardware/Cisco IOS upgrade.

The CallManager and the MCS may not enforce the guideline limits listed above. However, exceeding the limits of the recommended numbers will result in gradual degradation of voice quality for calls passing through the MCS. Therefore, it is good practice to keep the number under these limits in order to assure the service quality.

The rack mount requirements for the Media Convergence Servers and Catalyst switches are:

- The width of the rack between the two front mounting strips or rails must be 17.75 inches (45.09 cm).
- The depth of the rack between the front and rear mounting strips must be at least 19.25 inches (48.9 cm) but not more than 32 inches (81.3 cm).

The rack must have sufficient vertical clearance to insert the chassis of various systems. For example, the chassis height for the 6509 switch is 25.5 inches (64.8 cm), and the height for the MCS-7835 is 5.1 inches (13 cm).

Solution Implementation Templates

The project engineer should fill in the solution implementation templates based on the design prior to the implementation. This section provides the solution implementation templates for:

- WAN connection
- Address plan
- CallManager configuration
- Dial plan

WAN Connection Templates

The following templates are intended for the WAN configuration preparation for the implementation.

Remote Site Information Template

No.	DID	Remote Site	Route	Site Name
1				
2				
3				
4				
5				
6				
7				
8				
9				
10				

Remote Site Carrier Circuits Template

Region	Site Code	Circuit #	FRY/ATM	Speed	VPI	VCI	DLCI

Router Global Configurations Template

Name	Function	Equipment Type	IP Address	Gateway

Router Ethernet Configurations Template

Interface	IP Address	Subnet Mask	Connected To	Bandwidth	VLAN	Port

Router ATM/FR Configurations Template

Interface	IP Address	Subnet	Connected To	Bandwidth	DLCI	VBR	VPI	VCI

Address Plan Templates

The following templates are intended for address plan preparation.

Core Site IP Address and VLAN Assignment Template

Node Name	IP VLAN 1	IP VLAN 2	IP VLAN 3	IP VLAN 4	Subnet Mask	Device Type

Remote Site IP Address Assignment Template

Region	Site Code	Start IP	End IP	LAN	Subnet Mask	WAN	LOOP

CallManager Configuration Templates

The following templates are intended for CallManager configuration preparation.

CallManager Configuration Attribute Template

Parameter	Value
Primary CallManager host name	
Secondary CallManager host name	
Primary CallManager IP address	
Secondary CallManager IP address	
Site Win NT workgroup name	
NTP server address	
SNMP read community	
SNMP write community	

CallManager Port Template

CallManager	Description	IP Address	Ethernet Port	Digital Access	Analog Access

CallManager Group Template

CallManager Group	CallManager	Priority

CallManager Location Template

DID	Remote Site	Site Name	Bandwidth (Kbits)	Route ID

CallManager Device Pool Template

Device Pool	CallManager Group	Date/Time	Region

CallManager Partition Template

Partition	Description

CallManager Route Plan Template

DN/Route Pattern	Discard Digits	Prefix	Partition	Route List

CallManager Route Group Template

Gateway Name	Route Group	Selection Order

CallManager Route List Template

Route List	Route Group

Route List	Route Group

CallManager Gateway Template

Gateway	Device Protocol	Description	Device Pool	Calling Search Space

CallManager Pickup Group Template

Directory Number	Partition

CallManager Call Park Range Template

Call Park Number/Range	Partition	Cisco CallManager

IP Phone Template

Template Name	Number of Buttons	User Modifiable

IP Phone Button Template

Template Name	Button Number	Feature	Index	Label

Dial Plan Templates

The following templates are intended for the CallManager dial plan configuration.

IP Phone Extension Range Template

Cluster	Area Code	Extension Range

DID Number Assignment Template

Site	Phone	Organization	DID Block	DID Range

Customer Ordered Equipment

It is important to provide part numbers and quantities of the Cisco-supplied equipment. This can be done using Table 5-19 and Table 5-20, or by obtaining the information directly from the customer order details.

Table 5-19 Cisco Supplied Equipment

Part Number	Description	Quantity

Table 5-20 Customer Supplied Equipment

Description	Quantity

Customer Premises Equipment (CPE) Interface

Use Table 5-21 to document site-specific cable information.

Table 5-21 CPE Interface Checklist

From	To	Interface	Cable Type/s	Supplied By	Installed By

Customer Site Readiness

After performing a site survey, the site readiness check fully prepares each customer site to receive the installation. This customer site readiness is the summary result from the site survey.

Implementing the Solution

The Cisco project manager manages the on-site service installation and provides the field engineers with the information. The information is required for the field engineers to perform the physical installation.

Solution staging involves loading software and powering up all hardware at a single location prior to connecting the customer network. All devices should be connected and fully configured, tested, and commissioned. The customer should have technical staff participate at the network staging so they can become familiar with the IP Telephony solution. The result is a quick, seamless installation and cutover.

The implementation engineer should perform the following tasks in the following order:

1. Unpack equipment.
2. Verify installation of cabinet power feeds, rails, and earthing.
3. Physically install equipment in cabinet including cables between new network devices.
4. Record equipment serial numbers and verify against delivery documentation.
5. Verify equipment slot allocations.
6. Install intra-cabinet power and earthing cables for equipment
7. Install intra- and inter-cabinet communications cables.
8. Verify circuit termination in customer patch panel.
9. Power up Cisco equipment.
10. Verify and load system software and firmware.
11. Configure equipment.
12. Implement dial plan.
13. Configure E-911.
14. Conduct installation tests.
15. Add equipment to customer network.
16. Conduct solution acceptance tests.

Details for each activity are provided in the following subsections.

Unpacking the Equipment

The customer is responsible for ensuring that equipment is delivered to the installation location. The implementation team verifies that packaging has not been damaged in transit (check tip and shock indicators). The implementation team verifies that the equipment is in good condition when removed from the packaging. Cisco personnel assembles the equipment adjacent to the installed position.

Verifying Cabinet Power Feeds, Rails, and Earthing

The customer is responsible for providing the correct power supply. For AC supplies, appropriate sockets are to be provided adjacent to the rack position. For DC supplies, the customer must provide cabling to equipment with an appropriate crimp connector. Obtain an appropriate crimp tool from RS Components at http://rswww.com. (The product code 445-611.) It is the responsibility of Cisco personnel to connect to the equipment with the customer present to verify that the power is isolated. All power leads must be labeled.

The customer is responsible for providing an isolated earth cable to the equipment position with an appropriate crimp connector and documenting the details of earthing requirements. It is the responsibility of Cisco personnel to connect to the rack and equipment.

The implementation team verifies the customer's installation of cabinet power rails, power feeds, and earthing, and ensures that all supplies are isolated. It is the responsibility of Cisco personnel to label the power supply cables.

Physically Installing Equipment in Cabinet

Cisco personnel clearly defines how the equipment should be physically positioned within the cabinet, including cables between new network devices. Include information from the site survey detailing rack/equipment positioning and specific implementation instructions. Refer to the installation documentation that is supplied with each piece of equipment and provide the web address of specific Cisco product documentation. Highlight any points that relate to the specific implementation and include details where the standard installation document does not provide sufficient information.

Recording Equipment Serial Numbers

Serial numbers are required to track the location of equipment throughout the implementation. The Cisco project manager should ensure that the serial numbers recorded in this document at the time of implementation are supplied to the Cisco customer service team. The Cisco customer service team will ensure that the records used for support purposes are updated as required. Only serial numbers of field-replaceable items are recorded.

Verifying Equipment Slot Allocations

Cisco personnel clearly defines how the cards are to be positioned in the Cisco equipment. This is important for the routers, switches, and gateways. Use Table 5-22 as a matrix for recording card slot allocations.

Table 5-22 Equipment Slot Allocation

Equipment Number	Slot Number	Card Name	Card Function

Installing Intra-Cabinet Power Cables

Cisco personnel installs customer-prepared power cables between cabinet power rails and equipment power entry modules. Cisco personnel connects earth cable to cabinet earth point. The customer must present prepared cable to the cabinet position. Cisco personnel neatly ties and labels cables with key fob style holders.

Installing Intra- and Inter-Cabinet Communications Cables

Identify the intra-cabinet power cables to be installed by the Cisco engineer. The Cisco-provided intra-net cables are included in the order. The customer provides the inter-cabinet cables. Cisco personnel installs all provided cables between the equipment used in the same rack. Cisco neatly ties and labels cables with key fob style holders.

Verifying Circuit Termination in Customer Patch Panel

Confirm the responsibility of the customer in providing carrier circuits and cabling/patching to the Cisco cabinet. The customer confirms that the circuit designations between the patch panel and IP Telephony equipment are correct. The customer confirms that all cabling between the IP Telephony equipment and the patch panel is correct and has been tested. The customer confirms that all circuits have been successfully tested.

Powering Up Cisco Equipment

Cisco personnel provides details of the power-up procedure. Refer to the Site Survey report to identify any power-up restrictions for the equipment.

Cisco personnel is responsible for the following tasks:

- Turning on switches on all equipment power supplies (refer to any local requirement to power up in stages).
- Confirm that all equipment begins the power-up cycle.
- Confirm that all equipment provides a user prompt when a VT100 compatible laptop is connected to the console, or equivalent, port.
- Log in to each device in turns and verify there are no outstanding/unexplained alarms or unexplained equipment failures following powerup.

Verifying and Loading System Software and Firmware

Cisco personnel provides software/firmware revision requirements *and* an upgrade procedure in the event that the on-site engineer is required to perform an upgrade. A Cisco engineer connects to each Cisco WAN device using a VT100 compatible terminal and verifies switch software, boot code, and firmware versions. Cisco personnel connects to each Cisco router and verifies Cisco IOS software versions. A Cisco engineer corrects any variation to the defined releases.

Configuring the Equipment

A Cisco engineer configures the IP Telephony equipment based on the solution templates. This includes configuring CallManager servers, gateways, routers, switches, and other related devices, as well as the following:

- Call Admission Control
- CallManager clusters
- DSP resource provisioning (transcoding)

- Voice messaging system

Implementing the Dial Plan

A Cisco engineer implements the dial plan based on the design.

Dial Plan Architecture

The following list defines the functions in the dial plan.

- Route patterns—The matching of an E.164 address range or specific address points to a single route list. The CallManager software matches most specific wild card pattern and ranges as follows:

Wild Card	Range
X	Single Digit (0-9)
N	Single Digit (2-9)
@	North American Numbering Plan
!	One or more digits (0-9)
[x-y]	Generic Range Notation
[^x-y]	Exclusion Range Notation
.	Terminates Access Code
#	Terminates Inter-digit timeout

- Route lists—A list of route groups prioritized in “order” to reach a destination (private network or PSTN)
- Route groups—A prioritized trunk group made up of a list of devices/gateways to reach a particular destination (private network or PSTN)
- Devices/Gateways—Devices and gateways that interface with a remote private network or PSTN. Devices include:
  - H.323- or MGCP-based routers
  - Skinny-based Catalyst switches
  - Stand-alone Skinny devices such as DT24-T1 or E1

Dial Plan Configuration

The dial plan configuration varies depending on how voice calls are being routed and the number of paths the calls are being sent to in order to reach a specific destination. Calls can be routed on-net, which means inter-site or intra-site calls, and off-net, which means calls being sent outside the private network to the PSTN.

Calls are routed within and outside the IP Telephony network through a route pattern configured in the CallManager. The following subsections outline the basic steps required to successfully configure the CallManager to route calls.

Configuring CallManager Route Patterns

Perform the following steps to configure a CallManager route pattern.

Step 1 Add devices/gateways into the CallManager. Typical device types are Skinny-based, MGCP-based or H.323-based.

Step 2 Create a route group in CallManager and add devices available to CallManager to route calls.

Step 3 Create a route list to prioritize the routing of calls.

Step 4 Create or add route patterns to match an E.164 address range or specific address points to a single route list.

Configuring Dial Plan Groups and Calling Restrictions

Configuring dial plan groups and calling restrictions enables users to be grouped into communities of interest on the same CallManager that have the same calling restrictions and the same dial plans. Different communities of interest can share the same gateways and have overlapping dial plans. These capabilities are achieved in CallManager 3.0 with the use of call partitions and calling search spaces.

- Call partition—A group of devices with similar reachability characteristics. Devices such as IP phones, directory numbers (DN), and route patterns.
- Calling search space—An ordered list of partitions that a user/phone subscriber may look in before being allowed to place a call. Calling search spaces are assigned to devices that may initiate calls. These devices include IP phones, IP soft phones, and gateways.

Refer to the *Cisco CallManager Administration Guide* for detailed instructions on how to configure call partitions and calling search spaces. This document is available at the following location: http://www.cisco.com/univercd/cc/td/doc/product/voice/c_callmg/3_0/admin_gd/.

Configuring Digit Translation Tables

A digit translation table translates a dialed number to another number or changes the number of digits before forwarding to the destination. This can be achieved for internal and external calls, whether inbound or outbound, to the PSTN. The translation table may be configured so that when a user dials a telephone number (for example, "0" for the operator) the call gets *translated* to a directory number that corresponds to a user directory number.

Translation patterns are similar to route patterns in their use of wild cards and transformations. The primary use of translation patterns is for extension mapping.

For more information about digit translation, refer to the *CallManager Administration Guide*. This document is available at the following location: http://www.cisco.com/univercd/cc/td/doc/product/voice/c_callmg/3_0/admin_gd/.

Configuring E-911

A systems architect should provide an explicit dedicated routing plan as to how the system will handle emergency calls. When a user dials 911 or 9911, this call should be treated differently than any other call coming into the system.

The general responsibility of a telephony system is to:

- route the call to the appropriate point (either on-net or off-net).
- deliver calling party identification, typically in the form of the caller's phone number (either as an extension number, a PSTN number, or a combination of the two).

IP Telephony VoIP voice solutions can be configured to yield the following *essential* functionality:

- The appropriation of on-net and off-net call routing—On-net routing is the responsibility of the CallManager.

  The off-net selective routing always relies on an E.164 number being presented to the E-911 router as the call is initiated. The IP Telephony solution must provide the public network with a number for routing the call to the appropriate public safety answering point (PSAP).
- The forwarding of callback number information to the PSAP—The callback number provided to the PSAP is the E.164 automatic number identification (ANI) of the phone.

  There will be a sharp distinction on the level of 911 service available to DID and non-DID extensions (that is, whether a phone is directly reachable from outside of the IP Telephony system).
- The forwarding of location identification to the PSAP.

  This function is enabled at the PSAP when it receives a valid ANI for which there is an associated valid entry in the automatic location identification (ALI) database.

An IP Telephony architecture may be designed according to one of four general deployment models:

- Individual campus deployments
- Isolated multisite deployments
- Multisite IP WAN deployments with distributed call processing (Cisco CallManager at each site)
- Multisite IP WAN deployments with centralized call processing (one Cisco CallManager at the central site)

The four deployment models listed above have many similarities in the way they can each be configured to handle 911 calls. Therefore, this section examines a generic approach and itemizes the details and caveats for each of the four models.

This section examines the implementation of lifeline call processing capabilities whereby users dialing the 911 (or 9911) string will be routed off-net through a gateway to the PSTN which, in turn, switches or routes the call through an E-911 network to a public safety agency. This type of implementation is specific to the North American market.

All of this functionality can be used to route emergency calls to an on-net destination (for example, campus security) by using appropriate route patterns (911 or other security number such as 5111). Also, the decision to route emergency calls on- or off-net can be made independently for each branch office.

The processing of 911 calls can be broken down into the following components:

- Dial Plan
- Gateway Selection
- Gateway Interfaces

Dial Plan

The identification of the dialed 911 string can be made through the use of the “**@”** wildcard, which stands for the North American Numbering Plan. A special service labeled “911” can be used to filter out all calls but 911 calls.

The following patterns cover a user dialing either 911 or 9911. The **@** wildcard represents patterns that describe the North American Numbering Plan (E.164 numbers). The “WHERE” statements precede the route filter criteria. The first two patterns combined are equivalent to the bottom two patterns combined, and cover the two ways in which users may dial 911.

Pattern
911
9.911
@ WHERE SERVICE=911
9.@ WHERE SERVICE=911

Gateway Selection

The North American E-911 system is a collection of separate regional networks. Since they are typically not connected to allow rerouting of calls between jurisdictions, a 911 call must enter the E-911 network in the geographical jurisdiction of the appropriate PSAP.

In order to support a WAN deployment, there are geographical considerations to evaluate when configuring the system. Figure 5-1 shows how users in Branch X need to be connected to the E-911 network in their community through a gateway located at the same geographical location (Gateway X). Users at Branch Y need to use Gateway Y.

Figure 5-1 Gateway Selection

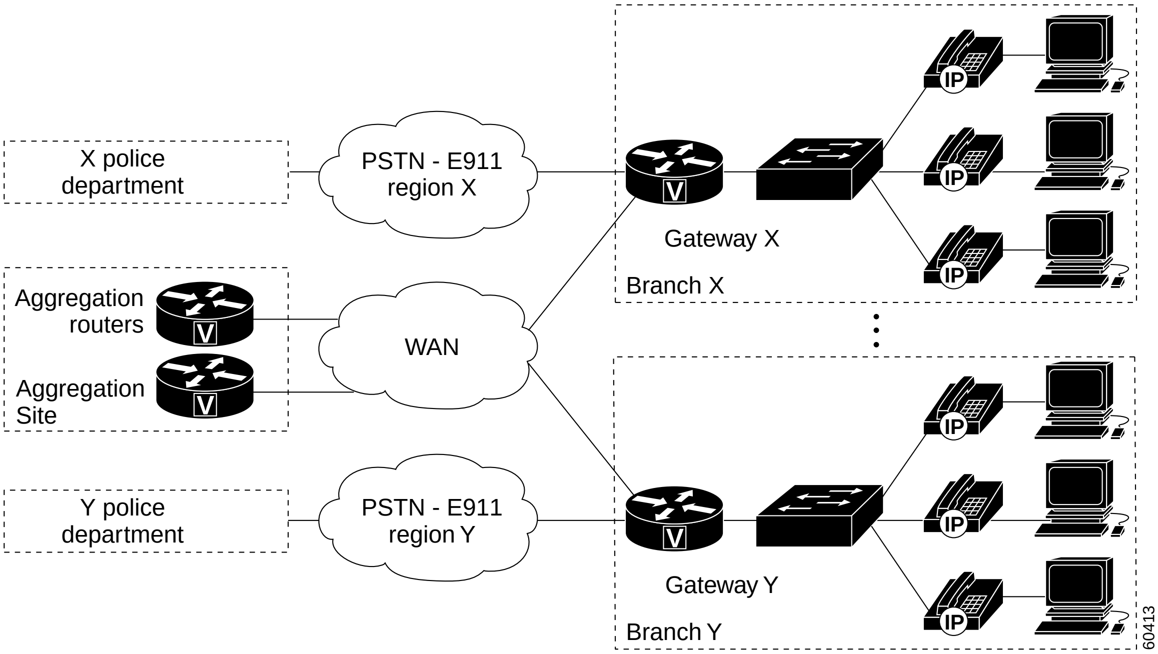

For 911 calls, users must be associated with a particular gateway according to geography since E-911 networks have only a local significance. In Figure 5-1, the clouds labeled "PSTN - E-911 region" represent the regional characteristic of the PSTN for E-911. For example, a 911 call cannot be sent to the Y Police Department using Gateway X.

Also, it is important to associate route patterns with partitions tailored to the geographical location (E-911 jurisdiction) of the users who have access to it.

Defining Partitions and Calling Search Spaces

The following example, beginning with Figure 5-2, defines how Partition Users_X_911 are dedicated to handle 911 calls for users in Branch X.

Figure 5-2 Partition Configuration Window

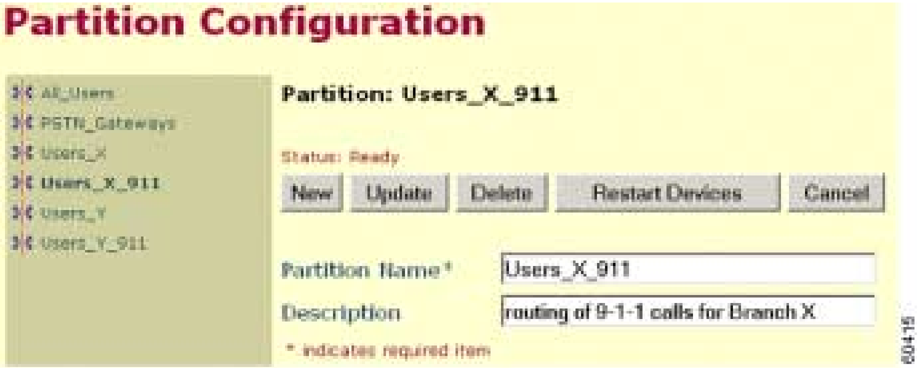

Partition Configuration

Partition: Users_X_911

Status: Ready

Partition Name* Users_X_911

Description routing of 9-1-1 calls for Branch X

* indicates required item

Figure 5-3 shows how the partition should then be included in the calling search space for users located in Branch X.

Figure 5-3 Calling Search Space Configuration Window

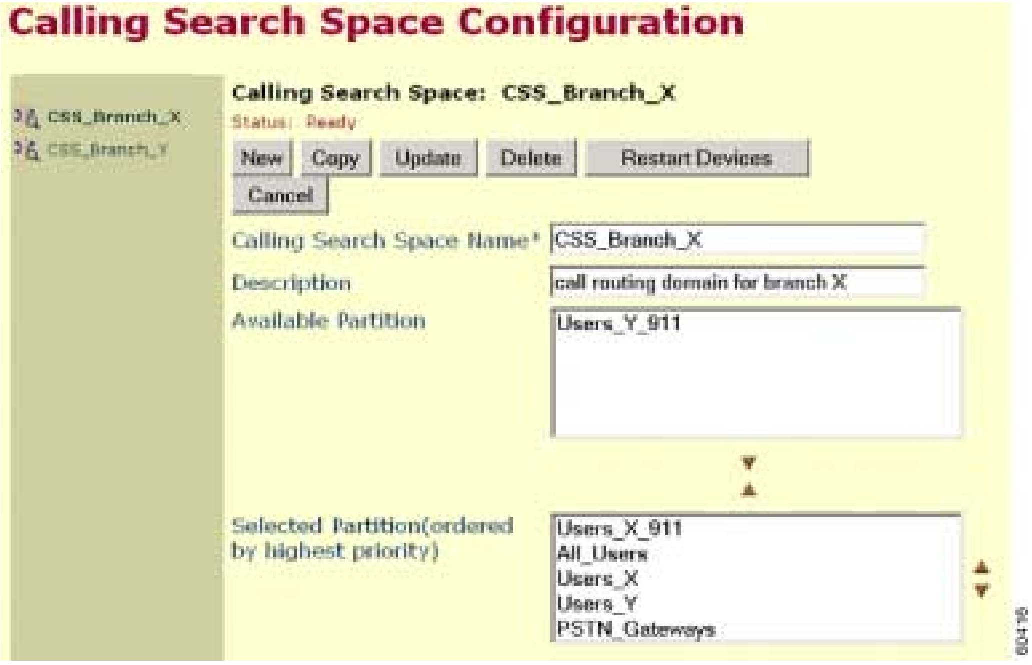

Calling Search Space Configuration

Calling Search Space: CSS_Branch_X
Status: Ready

Calling Search Space Name* CSS_Branch_X
Description call routing domain for branch X
Available Partition Users_Y_911
Selected Partition(ordered by highest priority) Users_X_911, All_Users, Users_X, Users_Y, PSTN_Gateways

Route Patterns

The route patterns need to be defined as belonging in partition Users_X_911, as shown in Figure 5-4.

Figure 5-4 Route Pattern Configuration Window

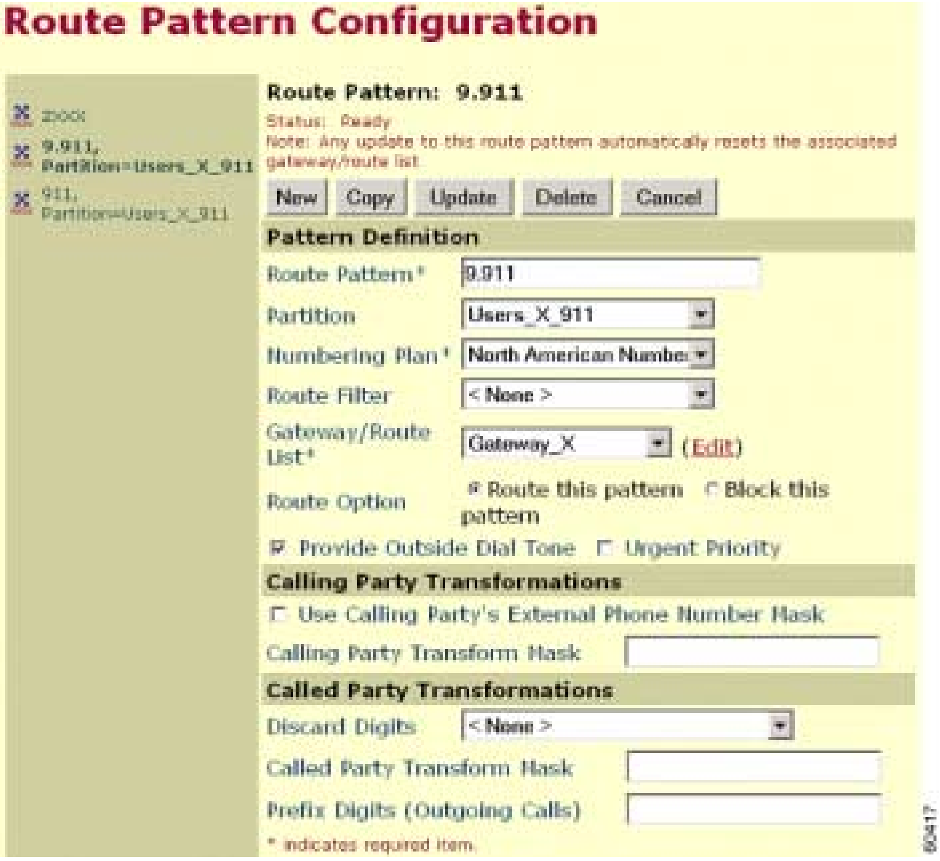

Route Pattern Configuration

zxxx
9.911, Partition=Users_X_911
911, Partition=Users_X_911

Route Pattern: 9.911

Status: Ready
Note: Any update to this route pattern automatically resets the associated gateway/route list

Pattern Definition

Route Pattern* 9.911
Partition Users_X_911
Numbering Plan* North American Numbe
Route Filter < None >
Gateway/Route List* Gateway_X (Edit)
Route Option ● Route this pattern ○ Block this pattern
☑ Provide Outside Dial Tone □ Urgent Priority

Calling Party Transformations

□ Use Calling Party's External Phone Number Mask
Calling Party Transform Mask

Called Party Transformations

Discard Digits < None >
Called Party Transform Mask
Prefix Digits (Outgoing Calls)
* indicates required item.

The value in the Gateway/Route List field should reflect a gateway that has access to the E-911 network for Branch X. In the above example, the Gateway_X value was selected because the Branch X gateway is connected to a PSTN switch that can reach the X police department: This is an acceptable point to which 911 calls from Branch X could be sent. The same type of configuration can be set for Branch Y.

Note For a given cluster, all Route Pattern/Partition combinations are unique. In the Figure 5-4 example, one instance exists for a 911 call and one instance exists for a 9.911 call. In some configurations, there may be two instances of a route pattern but they reside in different partitions. For example, an entry may exist for the 9.911 route pattern for partition X and a route pattern may exist for route pattern 9.911 for partition Y.

Figure 5-5 highlights the CallManager decision flow for two instances of 911 calls: a user in Branch X dialing 911, and a user in Branch Y dialing 9911.

Figure 5-5 CallManager Decision Flow

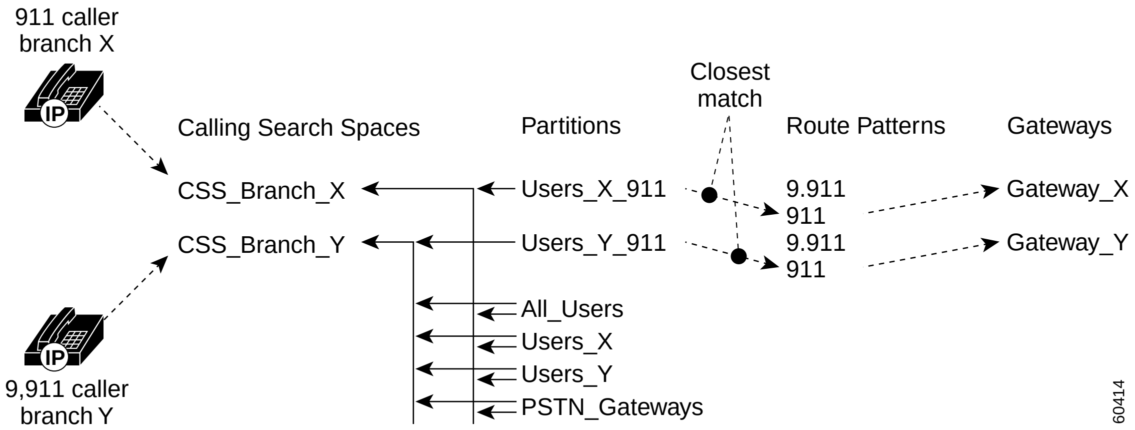

The gateway selection process is a function of the partitions to which a given phone has access. The partitions are controlled by the calling search space assigned to the phone. The closest-match routing process then selects a route pattern, which then points to a gateway (either directly or through a route list).

Route Filters

A generated route filter permits or restricts access through a route list using route patterns. Route filters are required if the route pattern is based on the **@** wildcard. The @ wildcard represents the collection of patterns that makes up the North American Numbering Plan.

Figure 5-6 shows a route filter configured to eliminate matching of all patterns represented by **@**, except for when SERVICE= 911. You can apply the same route filter to the 9.@ *and* @ route patterns.

Figure 5-6 Route Filter Configuration Window

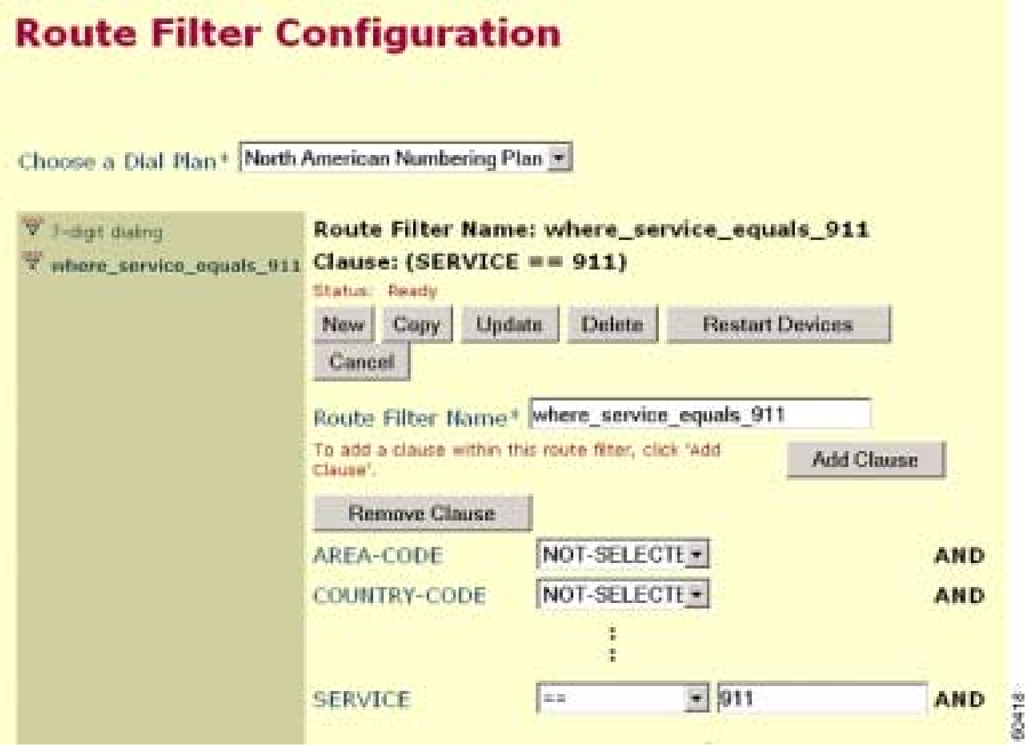

Route Filter Configuration

Choose a Dial Plan* North American Numbering Plan

Route Filter Name: where_service_equals_911
Clause: (SERVICE == 911)
Status: Ready

Route Filter Name* where_service_equals_911

To add a clause within this route filter, click 'Add Clause'.

AREA-CODE NOT-SELECTE AND
COUNTRY-CODE NOT-SELECTE AND
SERVICE == 911 AND

Calling Party Transformation Masks

The entire E-911 functionality relies on the calling party having a valid, fully qualified E.164 number (ANI) presented to the E-911 network. Based on the ANI, a 911 call can be enhanced to allow:

- routing to the appropriate PSAP
- delivery of a number that can be dialed, which public safety authorities may use to call the user back
- automatic retrieval of location information, which may be used to pinpoint the location of a silent caller

For 911 functionality, there are fundamental differences between DID and non-DID phones. These differences are *not* unique to the Cisco CallManager/IP Telephony architecture.

DID Phones

When a phone is assigned a DID number, it is best to forward that number as the Calling Party Number (CPN) to be used by the E-911 network as the ANI. In such a case, the calling party transformation mask should be used to translate the internal abbreviated dialing number of the caller into a fully qualified E.164 number. For example, if the number contains the area code, office code, and four-digit number; for extension 2003, the transformation would be to 408.555.2003.

The CallManager software offers two main locations where the calling party number can be transformed: the device level and the route pattern level.

The route pattern level can be used to transform the CPN as it goes out to the E-911 network without imposing any changes on the way internal calls are presented to the called phone or gateway. This means that when extension 2003 calls extension 2004, the called party sees an incoming call from 2003. If extension 2003 dials 911, the CPN extended to the local exchange carrier (LEC) through the gateway will be transformed, providing the called party with an ANI of 408.555.2003 (as an example only).

Non-DID Phones

When a phone is not assigned a DID number, there are several choices:

- Extend the call through the gateway with no CPN. This disallows use of any of the enhanced features of an E-911 network and may be illegal in some states. This forces the E-911 network to, by default, route the call based on either the listed directory number (LDN) or trunk group of the PSTN ingress point. This is definitely the least desirable option.
- Extend the call using a fixed CPN. This fixed CPN may be that of a campus security force, a DID phone near the location of the calling party, or the LDN of a trunk group (for example, the main number for a PSTN trunk). This fixed number is best used if it has entries in the E-911 selective routing and ANI/ALI databases of the local LEC. This option is more desirable since it at least allows the routing of the 911 call according to some administrator-controlled factor. Also, knowing the general origin of the call is better than knowing nothing.
- Use a third party calling line identification-automatic number identification (CLID-ANI) translator box. This may be required in some states.

The limitations on E-911 functionality imposed by the use of non-DID phones apply to all IP-based or non IP-based enterprise telephony systems.

Routing Overflow to an Alternate Number

If a preferred 911 gateway is not available, an alternative is to reroute the call through a regular PSTN gateway and transform the called party number to a fully qualified E.164 number. For instance, if the preferred gateway is located in San Francisco and is not available, it may be acceptable to overflow the call to a gateway in Oakland, and use (415) 553-8090 as the called number to be dialed on the PSTN. This number is the published seven-digit emergency number for the city and county of San Francisco. Admittedly, a call routed using this method does not utilize any of the advantages of a 911 call, but it is a preferable option to blocking the call altogether.

Gateway Interfaces

The level of 911 functionality available to an IP Telephony system depends not only on the internal configuration, but also on the type of interface used to extend the calls to the LEC.

PRI interface

In the current IP Telephony gateway portfolio, PRI interfaces offer the only native choice for *dynamic* 911 interfacing (for example, an interface that can present a different CPN for each call). This approach allows CallManager to dial 911 on one of the PRI bearer channels and present the calling party's phone number as the CPN ID. This assumes that CallManager has a *fully qualified E.164 number to send to the PSTN*. See "Calling Party Transformation Masks" on page 41.

Some LECs may not allow the CPN presented during call setup to be used as ANI. Instead, it uses the LDN of the trunk. This seems to be more a matter of LEC policy than technical limitation of the LEC's Class 5 switch.

POTS Interface

One of the simplest and most widely supported PSTN interfaces is the plain old telephone service (POTS) line. Any POTS line comes ready with its serving Class 5 switch connected to an E-911 tandem with its own E.164 phone number, emergency service number (ESN) assignment, and ALI database entry. The existing E-911 infrastructure supports 911 calls from POTS lines very well; it is the original technology around which E-911 was designed. Most of the current efforts around wireless and enterprise telephony support of E-911 are to bring these newer technologies to par with POTS 911 functionality.

For example:

- The foreign exchange office voice interface card (FXO VIC) is already supported by the IP Telephony solution.
- POTS lines may be dedicated for 911 calls only.
- The capital costs associated with POTS line interfaces may be low if the rest of the system's configuration already requires an FXO-capable chassis. If an entire chassis is required for a single POTS line, this approach may prove expensive.
- The OPEX (operating expenses) costs associated with a POTS line are low.
- The POTS line could serve as a backup line in the event of a power failure.
- The POTS line's number can be used as the callback number to be populated into the ALI database.
- POTS lines would represent the lowest cost 911 support for locations where user density does not justify local PRI access into the PSTN.
- POTS lines are ubiquitous PSTN installations.
- Any international lifeline scheme is supported by this approach.

Of course, this approach does not differentiate between DID and non-DID phones. All outgoing 911 calls are treated the same by the E-911 network, and the calling party transformation masks are irrelevant.

Centralized Automated Message Accounting Interface

Legislative efforts are in progress that will require enterprise telephony systems to interface directly into the E-911 networks. All 911 networks use Centralized Automated Message Accounting (CAMA) trunks, either in analog or digital (CAS over T1) form. There are two possibilities:

- CLID-CAMA converter—Cisco does not have a native CAMA solution. If you must connect an IP Telephony solution directly into the E-911 network, you must use a third party solution.
- CAMA VIC—Internal projects are in progress to add CAMA functionality to the IP Telephony solution.

Critical E-911 Considerations for All IP Telephony Deployment Models

This section contains important information to consider when implementing E-911 functionality for all deployment models.

User Mobility

User mobility is one of the main advantages of IP-based telephony systems. It allows automated relocation of a physical phone and a phone number to follow a physical device, even if the relocation crosses organizational, geographical, or technological boundaries. A user could have a phone number at work and then telecommute from home through a digital subscriber line and be seen as “in the office” by the other users. Dialing 911 from home would effectively connect the user to a PSAP that may be hundreds of miles away.

The current E-911 functionality described in this document relies on an entirely manual administrative process. User mobility poses a challenge since once a phone is configured as pointing into a certain calling search space, that configuration is maintained even if the user relocates to another physical location. The E-911 functionality within the IP Telephony solution currently does not adapt to a user’s mobility.

User Data Administration

The maintenance of the user names, location, and number information in a Public Safety Automatic Location Identification database is an entirely manual process. This model is obviously not scalable, and even on small IP Telephony deployments, keeping track of Add/Drop/Moves is a tedious process.

Always-on

One of the most critical attributes of emergency telephony devices is the ability to survive loss of commercial power. The typical LEC POTS line survives power outages with no loss of functionality. This can be accomplished by using in-line power of IP phones and UPS/generator backup of common infrastructure, gateway, and call processing equipment.

Also, you may want to configure a few LEC loop-start lines in the system. This would allow for general communications with the “outside world” during a power outage and could be used to provide fax and modem switched access at any time. The lines may also be connected to dedicated POTS phones for emergency use.

If such an approach is employed, the E.164 numbers of such lines may be used as CPN for 911 calls coming from non-DID phones. See “Non-DID Phones” on page 42 for more details.

Local Call Processing Resiliency

If a remote site is severed from its CallManager, the phones are not usable. A current project is underway to address this situation.

TDD Functionality

TDD (Telecommunications Device for the Deaf) is a modem-based communications device used to exchange characters-based conversations over regular phone lines. The 911 community, under the Americans with Disabilities Act, has to treat TDD calls with the same level of service as regular phone calls. As of December 2000, Cisco VoIP systems have been tested to work with TDD devices.

TDD machines use one of two modulation schemes:

- Baudot (for the vast majority)—The only technology formally mandated for compliance with the Americans with Disabilities Act (ADA)
- ASCII (a misnomer for frequency shift keying (FSK), carrier-based modulation)—Rarely used

Two types of TDD machines are available:

- Dedicated—Connects directly to the line through a modular connector
- Acoustically coupled—Uses the handset of a phone to the TDD machine

For both schemes, G.711 provides better results. Voice activity detection (VAD) should be turned off to facilitate communications.

For more information about telecommunications for the deaf, go to the following location: http://www.zak.co.il/deaf-info/old/tty_faq.html.

LEC Interfacing

See "Centralized Automated Message Accounting Interface" on page 43 for a description of the required LEC telephony interfaces currently under development.

911 Call Notification and Accounting

It is a requirement of many campus installations that a notification system be in place to signal that a user is calling 911. The obvious example is that of an educational institution whose campus security force would monitor 911 calls and be provided with the extension number of the caller. A way to actually monitor the audio conversation and potentially "barge-in" may be required. The IP Telephony architecture does not currently offer a native way of performing these functions.

Call detail records (CDRs) dedicated to reporting 911 calls are also often stated in customer requirements. From an IP Telephony perspective, there are no special accommodations for 911 calls compared to any other type of call. Third party development of a CDR package that would allow the parsing of records based on called party number may be developed, but is not readily available.

Critical E-911 Considerations for Single Site Deployment Models

This section contains important information to consider when implementing E-911 functionality for all single site deployment models.

Individual Campus Deployments

Configuring a small-scale deployment is fairly straightforward because users are concentrated within a single building. All users may share the same gateway association for 911 calls. For very small implementations, LEC POTS lines may be the only type of interfacing required to satisfy 911 call requirements.

For isolated multisite deployments, each of the sites is a separate entity that must be separately configured for the processing of 911 calls.

Multisite Deployment Models

Multisite IP WAN deployments (with either distributed or centralized call processing) share the same challenge of correctly associating phones with gateways. This association must be based on the actual geographical location of the phone *and* the E-911 domain into which the gateway has access.

Multisite deployments can rely on the functionality described in the previous sections; the system designer has to recognize the jurisdictional boundaries of the E-911 network and provision the system so that an E-911 network ingress point is provided for all locations where phones are physically installed. This implies that wherever a single phone is located, an appropriate local gateway should be provisioned to offer E-911 functionality. Alternatives may be considered if called party number translations are allowed (for example, if it is acceptable to route a call originally dialed as "911" as if a corresponding 7-digit emergency number had been dialed). However, this process may not be considered legal since it does not offer all of the enhanced functions of E-911.

Conducting Installation Tests

Cisco personnel provides a table of installation tests to be carried out. Each test is to be witnessed by the customer. The individual test sheets must be completed. The individual test sheets are located in the "Solution Implementation Acceptance Testing" section on page 5-48. In addition, any failures must be indicated on the System Handover Certificate.

Adding Equipment to Customer Network

Cisco personnel provides clear instructions as to how the equipment is to be added to the customer network. Reference to the customer change control process may be required.

Conducting Solution Acceptance Tests

Provide a table of commissioning tests to be carried out. Each test is to be witnessed by the customer. The individual test sheets must be completed. The tests are located in the "Solution Implementation Acceptance Testing" section on page 5-48.

Fallback Procedures

This section describes the fallback procedures for the IP Telephony implementation.

Falling Back from IP Telephony to TDM

To fall back from the IP Telephony solution to the customer's previous time division multiplex (TDM) configuration, perform the following steps:

Step 1 Power down the CallManager 3.0 server.

Step 2 Restart the TDM PBX switch.

The original configuration settings should be preserved so no further modifications to the TDM configuration should be necessary.

Step 3 Perform regression tests to verify the pre-migration configuration.

Implementing a Migration Strategy

This section provides information to ensure an overall smooth migration process to deploy and upgrade Cisco CallManager 3.0 and to equip Cisco Professional Services personnel with a checklist before they install or upgrade a new Cisco CallManager 3.0 solution.

Cisco Systems recommends that you plan the upgrade in two phases: dry run and actual run. The dry run phase allows you to roll back to an earlier software release in the event that further preparation or modifications are required.

1. Dry Run

   During the dry run, follow the planned upgrade procedures. If issues arise during the dry run within the designated time frame, you can roll back to the pre-upgrade configuration. Record all issues that arise from the dry run. These issues should be investigated, isolated, and fixed prior to performing the actual migration. Procedural changes should be modified in the migration plan as necessary.

2. Actual Run

   The actual upgrade should occur after all of the identified issues from the dry run have been addressed. During this phase, follow the procedures from the dry run. If issues arose during the dry run, use the modified procedures. Upon successful completion of the actual run, the site should be upgraded to the configuration defined in the system architecture for the Cisco IP Telephony solution.

Each upgrade phase should be scheduled during off-business hours. If the first dry run is successful, accept the dry run as the actual run and eliminate the second step to complete the migration.

Migrating from a TDM Network to Cisco IP Telephony Solution

To upgrade a customer site from a TDM voice network environment to the Cisco IP Telephony solution, a series of actions should be planned and executed prior to and during the transition. A Cisco representative should adhere to the following steps to ensure a smooth transition:

1. Communicate the site requirements to the customer.
2. Conduct a site survey to collect the information pertaining to the customer site.
3. Prepare a detailed customer data network architecture diagram to use as a starting point for planning the adequate Cisco IP Telephony network backbone.
4. Document a detailed design and engineering plan specific to the customer site and communicate these plans to the customer.
5. With feedback from the customer, finalize and execute a thorough deployment plan.

A separate document for customer solution verification and an acceptance plan may be presented to the customer as part of the contract. Cisco Professional Services and technical assistants will participate with the Cisco IP Telephony solution deployment to ensure the sanity of the Cisco IP Telephony solution, satisfactory network performance, and network management.

Upgrading Cisco CallManager

For complete upgrade instructions, refer to the following document:
http://www.cisco.com/univercd/cc/td/doc/product/voice/c_callmg/3_0/install/305upgra.htm

Migration Phases

Depending on the size, multiple locations, and complexity of the customer data network, a phased migration may be necessary. In this case, Cisco personnel should plan guidelines with regard to the various phases of the migration. These guidelines should be negotiated with the customer. Once the number of phases and extent of each phase for the migration is agreed upon by the customer, a detailed plan containing procedures, entry criteria and exit criteria for each phase should be documented before migration starts, and be updated as the migration proceeds.

Solution Implementation Acceptance Testing

This section describes the processes of Solution Implementation Acceptance (SIA) tests and the criteria for the SIA testing. All solution implementation tests are the responsibility of Cisco Systems Professional Services engineers or Cisco Professional Services Partners, and should be carried out in the presence of a customer representative. Any issues arising from the SIA test program will be entered in the project Issues Log. The expected duration of the SIA test program is seven working days.

Verification Process

After the IP Telephony solution has been implemented, and before being turned over to the customer for operation, SIA tests should be conducted to verify the correctness and completeness of the implementation. Table 5-23 lists the SIA designed to verify the IP Telephony solution.

Note The Solution Implementation Acceptance tests are available at the following location: http://www.cisco.com/warp/public/788/solution_guide/forms/index.html#iat.

Table 5-23 Implementation Solution Acceptance Test Master List

Test Number	Test Title
CallManager Checkup Tests	
1001	CallManager Service Status
1002	CallManager Configuration
1003	CallManager Device Default Configuration
1004	CallManager Gateway
1005	CallManager Voice Mail Port (unified message port)
General IP Phone Tests	
1101	Basic IP phone
1102	Call On-hold and Retrieve
1103	Call Park
1104	Group Pick Up
1105	Call Waiting
1106	Shared Line Appearance
1107	Call Transfer to an Off-net Phone
1108	Call Forward to an Off-net Phone
Campus Call Processing	
1201	Phone-to-phone Dial-up on LAN

Table 5-23 Implementation Solution Acceptance Test Master List (continued)

Test Number	Test Title
1202	Call Transfer to an On-net IP Phone
1203	Call Forward to an On-net IP Phone
1204	Ad-hoc Conference
1205	Meet-me Conference
Centralized Call Processing	
1301	Phone-to-phone Dial-up to a Remote Site
1302	Call Transfer to an IP Phone in a Remote Site
1303	Call Forward to an IP Phone in a Remote Site
1304	Ad-hoc Conference over WAN Link
1305	Meet-me Conference over WAN Link
Distributed Call Processing	
1401	Phone-to-phone Dial-up in Different Clusters
1402	Call Transfer to an IP Phone in Another Cluster
1403	Call Forward to an IP Phone in Another Cluster
1404	Ad-hoc Conference Between Two Clusters
1405	Meet-me Conference Between Two Clusters
Advanced IP Telephony	
1501	(optional) Four-party Conference
1502	(optional) Six-party Ad-hoc Conference
1503	(optional) Modem-to-modem
1504	(optional) Fax-to-fax
1505	(optional) CallManager Failover

Acceptance Criteria

The solution implementation acceptance test results will be judged to have been successfully completed if the network is deemed fit to carry live customer traffic and to be supported by the Cisco TAC (assuming that a suitable service contract is in place). If, for example, as a result of this test procedure being carried out, a software or hardware defect is discovered that does not severely impact the operation of the network, the investigation and rectification of this defect will continue to be managed under the Cisco support process after the Professional Services engagement has completed. Cisco Professional Services will manage the transition of any outstanding technical issues as part of the disengagement process.

Post-implementation Documentation

After the implementation phase of the IP Telephony solution is complete, the project team documents all equipment for asset tracking purposes and to verify ownership of the equipment. All Cisco-provided equipment should be removed from the customer premises. It is important that the project team records the serial numbers of all equipment, especially if the customer purchased SmartNet support.

Asset Tag and Cable Labeling

Asset Tag and Cable Labeling is an optional service but is very important for operations management and troubleshooting. The benefits of this service include:

- Enables customer to conveniently manage new equipment arrivals and enter information into existing databases
- Improves ongoing network equipment tracking
- Reduces overall installation cost and on-site time

Customer Acceptance Certification

The customer completes and signs the Customer Acceptance Certification, and returns it to the Cisco Systems project manager. An example of a Customer Acceptance Certificate is shown below.

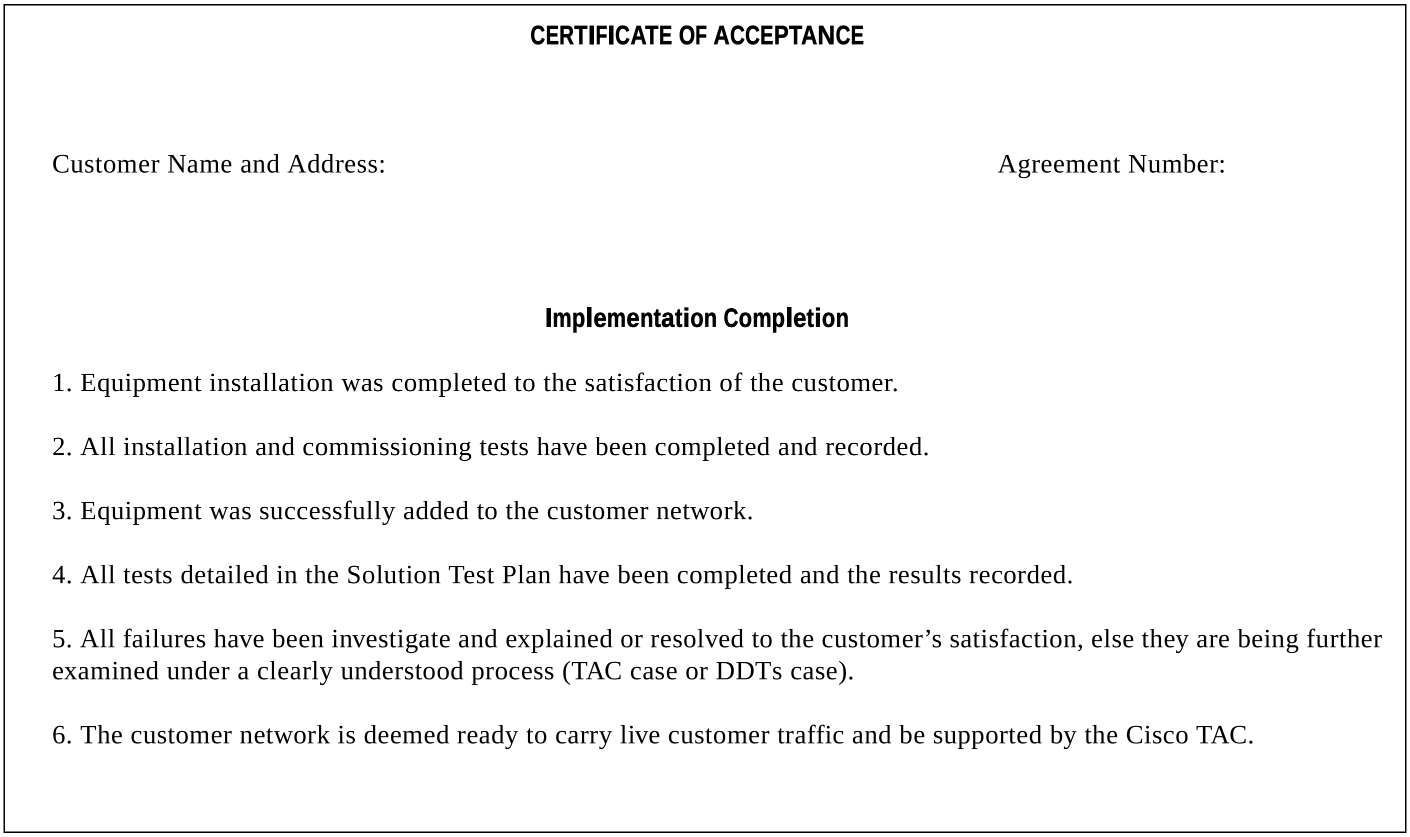

CERTIFICATE OF ACCEPTANCE

Customer Name and Address:

Agreement Number:

Implementation Completion

1. Equipment installation was completed to the satisfaction of the customer.

2. All installation and commissioning tests have been completed and recorded.

3. Equipment was successfully added to the customer network.

4. All tests detailed in the Solution Test Plan have been completed and the results recorded.

5. All failures have been investigate and explained or resolved to the customer’s satisfaction, else they are being further examined under a clearly understood process (TAC case or DDTs case).

6. The customer network is deemed ready to carry live customer traffic and be supported by the Cisco TAC.

CERTIFICATE OF ACCEPTANCE

THE UNDERSIGNED ("CUSTOMER") IS A CUSTOMER UNDER THE INSTALLATION AGREEMENT WITH CISCO SYSTEMS, INC. CUSTOMER REPRESENTS AND CERTIFIES THAT THE IMPLEMENTATION OF THE IP TELEPHONY SOLUTION HAS BEEN ACCEPTED BY CUSTOMER ON THE ACCEPTANCE DATE INDICATED BELOW AND LABELS, IF SUPPLIED, HAVE BEEN AFFIXED TO EACH ACCEPTED ITEM OF EQUIPMENT.

CUSTOMER CERTIFIES THAT THE ITEMS LISTED ABOVE HAVE BEEN ACCEPTED BY CUSTOMER ON THE ACCEPTANCE DATE INDICATED BELOW

DELIVERY OF AN EXECUTED COPY OF THIS CERTIFICATE OF ACCEPTANCE BY FACSIMILE OR ANY OTHER RELIABLE MEANS SHALL BE DEEMED TO BE AS EFFECTIVE FOR ALL PURPOSES AS DELIVERY OF A MANUALLY EXECUTED COPY. CUSTOMER UNDERSTANDS THAT CISCO SYSTEMS MAY MAINTAIN A COPY OF THIS CERTIFICATE IN ELECTRONIC FORM AND AGREES THAT A COPY PRODUCED FROM SUCH ELECTRONIC FORM OR BY ANY OTHER RELIABLE MEANS (FOR EXAMPLE, PHOTOCOPY, IMAGE OR FACSIMILE) SHALL IN ALL RESPECTS BE CONSIDERED EQUIVALENT TO AN ORIGINAL.

Accepted by: ______________________________ Accepted by: ______________________________

(Authorized Cisco Signature) (Authorized Customer Signature)

______________________________ ______________________________

(Name - Type or Print) (Name - Type or Print)

Acceptance date: ___________________________ ______________________________

(Must Be Completed) (Must Be Completed)

Please return to the Cisco Systems Project Manager

Completing the Implementation Reports

At the end of the solution implementation, the project manager needs to complete the implementation reports for the account team. This can be easily accomplished by summarizing all of the implementation documents and the solution acceptance test results.

The date of the completion and customer acceptance certification should be documented for official record.

Case Study

An IP Telephony case study is available at the following location: http://www.cisco.com/warp/public/788/AVVID/hdppcasestudy.html.

While the names in this case study are fictional, it is based on a real-world IP Telephony implementation.

CHAPTER 6

Operating the IP Telephony Network

This chapter includes the following sections:

- Operations Support and Planning
- Network Management
- Security
- Troubleshooting

Related Information

You can also refer to the following websites for related information:

- Cisco Network Monitoring and Event Correlation Guidelines

  http://www.cisco.com/warp/partner/synchronicd/cc/pd/wr2k/tech/cnm_rg.htm
- Cisco MIB Files Part I

  http://www.cisco.com/public/mibs/v1/
- Cisco MIB Files Part II

  http://www.cisco.com/public/mibs/v2/
- Cisco IP Telephony Troubleshooting Guide for Cisco CallManager Release 3.0(1)

  http://www.cisco.com/warp/public/788/AVVID/ts_ccm_301_sec1.htm
- CiscoWorks Enterprise Network Management Solutions

  http://www.cisco.com/univercd/cc/td/doc/pcat/index.htm#CFHBHEAF

Operations Support and Planning

Telecom networks, data networks, and servers are currently managed somewhat independently. Administrators can differentiate service issues fairly easily. Operational processes such as performance, capacity, provisioning, fault management and inventory management are typically managed by separate groups without much interaction. Each group may also have its own independent support plan with its own unique goals or service requirements that meet existing requirements.

The IP Telephony solution typically requires a new support model:

- Groups or individuals that have historically not interacted with one another may now work more closely together.
- New support processes that meet specific IP Telephony requirements may be required.
- New roles and responsibilities to ensure support at each level for all areas of the solution may also be required.
- Service requirements to improve availability for voice traffic may need to be enhanced.

To achieve a consistent reliable solution over time, Cisco recommends that the operations support organization review operational support requirements. These requirements may include role and responsibility changes, new process additions, process modifications, and service definition changes. This review is best done by first reviewing the technical constraints of the IP Telephony solution, determining the roles and responsibilities, and then defining and approving service elements.

Table 6-1 lists many of the operations requirements for IP Telephony operations.

Table 6-1 Technical Support Matrix

	Configuration	Fault	Performance	Security	Accounting
IP Telephony application level	Change management MAC process Dial plan	Monitor CallManager application and availability	CallManager utilization calls Voice quality	User IDs Passwords	Call detail records Billing fraud detection
Platform operating system	CallManager backup and recovery	Monitor device, SNMP monitoring	Disk space RAM CPU	User IDs Passwords	System event logs
LAN/WAN network	DCHP, DNS, TFTP, and IP addresses	Catalyst switches and routers	LAN traffic device resources	User IDs Passwords	SysLogs NATkit
Hardware	Server, switches, routers, and cabling	24 x7 x4 breakfix on-site support	MTBF (mean time between failure) MTTR (mean time to repair)		

This chapter also identifies unique IP Telephony operational service requirements and provides service definition examples recommended to help ensure consistent, reliable availability and performance.

Defining Technical Goals and Constraints

Analyzing technical goals and constraints of the Cisco IP Telephony solution helps the support organization understand service level requirements for IP Telephony services and whether the service goals are achievable. This process can also help set expectations regarding achievable service levels to users or customers.

The organization should first work to identify all technical goals, constraints, and service risks involved in achieving the desired service level for the IP Telephony solution. The following categories can be used to help identify constraints:

- Network or IP Telephony technology, resiliency, and configuration
- Life cycle practices, including planning, design, implementation, and operation

- IP Telephony application behavior and/or requirements

The constraints should then be prioritized in terms of the greatest risk or impact on service level goals. This helps the organization prioritize support initiatives and to understand how easily the constraints can be addressed.

Network or IP Telephony Technology, Resiliency, and Configuration

Network technology, resiliency, and configuration constraints for IP Telephony can be defined as any limitation or risk associated with the current technology, hardware, links, design, or configuration.

Technology limitations cover any constraint posed by the technology. For example, no current technology allows for sub-second convergence times in redundant network environments, which are critical for sustaining IP Telephony voice quality across the network. Another limitation is the raw speed that data can traverse terrestrial links, which is approximately 100 miles per millisecond. This is an important consideration in WAN deployment models.

Network hardware resiliency risk investigations should concentrate on hardware topology, hierarchy, modularity, redundancy and mean time between failure (MTBF) along defined paths in the network. Network link constraints should focus on network links and carrier connectivity for enterprise organizations. Link constraints may include link redundancy and diversity, media limitations, wiring infrastructures, local loop connectivity, and long distance connectivity.

Design constraints relate to the physical or logical design of the network and include everything from available space for equipment to scalability of the routing protocol implementation. All protocol and media designs should be taken into consideration in relation to configuration, availability, scalability, performance, and capacity.

In addition, the following network service constraints should be taken into consideration:

- CallManager functionality and reliability
- Gateway functionality and reliability
- DHCP and DNS reliability

Life Cycle Practices

The network life cycle refers to the cycle of planning, design, implementation, and operations. Life cycle practices define the processes and management of the network used to:

- Successfully deploy solutions
- Quickly detect and repair problems
- Prevent capacity or performance problems
- Configure the network for consistency and modularity to reduce undesirable complexity and unanticipated behavior.

The impact of life cycle practices is important because expertise and process are typically the largest contributors to non-availability.

IP Telephony Application Behavior and Requirements

IP Telephony traffic has unique performance constraints unlike data traffic. Most IP application traffic uses TCP, which allows for re-transmission and efficient use of bandwidth. IP Telephony uses the RTP (real time protocol), which does not accommodate re-transmission. Voice is also sensitive to delay and jitter. A reasonable constraint for end-to-end delay is 250 milliseconds or 1/4 of a second. Jitter, or the amount of time between voice packets, should be less than 25 milliseconds. In many cases, the current data network will not accommodate this consistent performance problem because of queuing delay or sporadic network congestion. In WAN deployments, delay is almost always a factor because of the delay inherent in all WAN links. Data can generally travel about 100 miles per millisecond. In coast-to-coast connections, this may be somewhere between 25 and 50 milliseconds, depending on the routing of the circuit. Expect longer delay for many international connections.

A worksheet similar to the one in Table 6-2 can help identify IP Telephony solution constraints. The organization should also consider their current support capabilities in addition to anticipated technology constraints in their environment. It is recommended that the architecture, engineering, and operations teams meet to discuss and identify all potential constraints.

Table 6-2 Technical Goals and Constraints

Risk or Constraint	Type of Constraint	Potential Impact
VoIP delay requirement of consistent 250ms or below of delay across the network	IP Telephony or VoIP application	high
VoIP jitter requiring consistent 25ms jitter or less	IP Telephony or VoIP application	high
IP Telephony technology stability and solution age	IP Telephony or VoIP application	medium
No performance tools to manage end-to-end performance jitter and delay is implemented	Life Cycle Process/Technology	medium/high
Single point of failure in some gateways	Technology	medium, but could be redundant
Single point of failure in some CallManager clusters	Technology	medium, but could be redundant
No UPS systems in wiring closets and no backup power for phones that exists with today’s telecom infrastructure	Technology	medium, but solution exists
No IP management of phones	Technology	medium to low
Non-redundant DHCP services	Technology	medium
Non-redundant DNS services	Technology	medium
No tool or process to identify down ISDN channels on voice gateways	Life Cycle process/Technology	medium
Additional voice traffic impact on capacity	Technology	medium low
No QOS implemented in LAN environment	Technology but can upgrade	medium
Must refine process for alerts and SYSLOG review for identifying CallManager problems	Life Cycle practices	medium
No team identified to achieve 4 hour CallManager or Gateway replacement with hardware failure	Life Cycle practices	medium
Ease of resolving network versus CallManager issues	Life Cycle practices	medium
No bandwidth redundancy on redundant WAN links to critical sites	Technology but can upgrade	medium

Service Level Goals

Availability and performance standards set the service expectations for the support organization and help to define service and support requirements. Service goals are typically set for availability and performance. Performance may include factors such as delay, jitter, maximum throughput, and bandwidth commitments. Availability for IP Telephony will include IP Telephony and gateway availability, in addition to overall network availability.

Performance goals based on IP Telephony or VoIP requirements should be defined. Availability goals based on business requirements should be created; however, technical constraints and cost will always be a major factor. Analyze the following areas to determine the potential availability of the IP Telephony solution:

- Hardware path MTBF (mean time between failure) and MTTR (mean time to repair)
- Software reliability
- Power/Environment availability, including disaster preparedness
- Carrier or link availability
- Network design, including redundancy and convergence capabilities
- User error or process considerations, including the time it takes to isolate and resolve technical problems

Once the service areas and service parameters have been defined, use information from previous steps to build a matrix of service standards. Availability may be based on expected availability that is defined by investigating each of the above areas and the expected support capabilities. It is recommended to also define areas that may be confusing to users and IT groups. For instance, maximum response time will be very different for a round-trip ping than what users may experience on a voice call. Table 6-3 lists the service standards that should be acceptable to a majority of organizations.

Note Ping does not always accurately measure the response time for RTP or voice traffic due to QOS configurations for RTP and ping process priority for many platforms. Instead, measure RTP traffic performance with a performance tool such as Cisco Internet Performance Monitor (IPM), which is bundled with CiscoWorks2000.

Table 6-3 Service Goals for IP Telephony

Network Area	Availability Target	Measurement Method	Average Network Response Time Target	Maximum Response Time Accepted	Maximum Jitter Accepted
LAN	99.99%	Ping monitoring (tool to be determined)	under 50 ms (round trip ping)	250 ms	20 ms (jitter measurement to be determined)
WAN	99.9%	Ping monitoring (tool to be determined)	under 100 ms (round trip ping)	250 ms	20 ms (jitter measurement to be determined)

Determining the Relevant Parties

IP Telephony support may require service level agreements, or at least service level definitions, between a variety of parties, including the telecom organization, the data networking organization, the NT server administration group, and users or organization leaders.

If the organization expects to craft service level agreements for IP Telephony service, then representatives and managers from each of these groups should be involved in understanding their roles and potential responsibilities along with the user point of view. In many cases, service definitions for telecom support and data support will need to be combined or re-crafted to form a new basis that synchronizes with the current support capabilities of each organization.

See the following section for potential service requirements for IP Telephony. The parties involved should examine these support requirements to further understand the roles and responsibilities from each team. In addition, the organization should define additional staffing and expertise required to fully meet service requirement goals.

Defining Service Elements

Service element requirements for IP Telephony may already dovetail into many existing defined service elements within the organization. In other cases, the organization may define new service elements in conjunction with the IP Telephony solution. However, all service definitions should be reviewed to ensure that roles and responsibilities are properly defined for different aspects of IP Telephony and that the defined service level meets IP Telephony and business requirements.

Service elements are grouped into two different categories: reactive service elements and proactive service elements. Reactive service elements define the processes required to report and repair IP Telephony problems. Proactive service elements are also needed to help ensure consistent performance, voice quality, provisioning, and change success and security. Service elements currently defined for IP Telephony include the following:

- Reactive Service elements
  - Service repair expectation and problem priorities
  - Problem reporting and escalation paths
  - Fault monitoring
  - Device file backups and recovery
  - Hardware replacement/on-site assistance
  - Escalation paths and requirements
  - Disaster recovery processes
  - Metrics and reporting
- Proactive service elements
  - Change management
  - MAC (move, add, change) process
  - Telephone number and dial plan management
  - IP/DHCP management
  - Proactive network management processes

- Fraud detection
- Billing

The following sections provide collateral and example material for service level definitions related to the above service level elements.

Reactive Service Elements

Reactive service elements define the processes needed to consistently and quickly isolate and repair IP Telephony faults. Each service area includes a discussion of the required parameters of the service elements and in some cases, examples of service level elements that can be used as a template when creating the IP Telephony operations support plan within the organization.

Service Expectations and Problem Priorities

The first set of service level definitions normally requires document service expectations and problem priorities for identified problems. These service level areas are typically measured using help desk database statistics and periodic auditing. The organization will normally want to review their existing service levels in this area and define service expectations and problem priorities specific to the IP Telephony application.

Table 6-4 identifies the support severity levels associated with problems that may arise in an IP Telephony network. Some organizations may also have a severity 5 request for new service if handled through the same support process.

Table 6-4 Support Severity Levels

Severity 1	Severity 2	Severity 3	Severity 4
Severe business impact	**High business impact through loss or degradation, possible workaround exists**	**Some specific network functionality is lost or degraded such as loss of redundancy**	**A functional query or fault that has no business impact for the organization**
• LAN user or server segment down • Critical WAN site down • 25 percent or more phones at site down, or functionality severely degraded	• Campus LAN down, 5 to 99 users affected • Domestic WAN site down • International WAN site down • Critical performance impact • 15 percent or more phones at site down, or functionality somewhat degraded	• Campus LAN performance impacted • LAN redundancy lost • Single phone outage or service-affecting problem	

Once a problem severity has been defined, the organization may need to define or investigate the support process to create service response definitions. In general, service response definitions require a tiered support structure coupled with a help desk software support system to track problems via trouble tickets. Metrics should also be available on response time and resolution time for each priority, number of calls by priority, and response/resolution quality. To define the support process, it helps to define the goals of each support tier in the organization and their roles and responsibilities. This helps the organization understand resource requirements and levels of expertise for each support level. Table 6-5 shows an example of a tiered support organization with problem resolution.

Table 6-5 Tiered Support Example

Support Tier	Responsibility	Goal
Tier 1	• Full-time help desk support • Answer support calls, place trouble tickets, work on problem up to 15 minutes • Document ticket and escalate to appropriate tier 2 support	Resolve 40 percent of incoming calls
Tier 2	• Queue monitoring, network management station monitoring • Place trouble tickets for software problems • Implementation • Take calls from tier 1, vendor and tier 3 escalation • Retain ownership of call until resolved	Resolve 100 percent of calls
Tier 3	• Provide immediate support to tier 2 for all priority 1 issues • Agree to help with all problems unresolved by tier 2 within SLA resolution period	No direct problem ownership

The next step is to create the matrix for the service response and service resolution service definition. This sets goals for how quickly problems, including hardware replacement, are resolved. It is important to set goals in this area because service response time and recovery time directly impact network availability.

Problem resolution times should also be aligned with the availability expectations. If high numbers of high severity problems are not accounted for in the availability goal, the organization can then work to understand the source of these problems and a potential remedy.

Table 6-6 shows an example of service response and resolution goals.

Table 6-6 Service Response and Resolution Goals

Severity	Help Desk Response	Tier 2 Response	On-site Tier 2	Hardware Replacement	Resolution
1	Immediate escalation to tier 2, network operations manager	5 minutes	2 hours	2 hours	4 hours
2	Immediate escalation to tier 2, network operations manager	5 minutes	4 hours	4 hours	8 hours
3	15 minutes	2 hours	12 hours	24 hours	36 hours
4	15 minutes	4 hours	3 days	3days	6 days

Reporting and Escalating Problems

A defined, documented, and well-published process should be in place for reporting IP Telephony issues. Different support organizations may field problems depending on geographic boundaries. In addition to service response and service resolution, an escalation matrix should be created. The escalation matrix helps ensure that available resources are focused on more severe service problems. In general, when analysts are focused on fixing problems, they rarely focus on bringing additional resources in on the problem. Defining when additional resources should be notified helps to promote problem issue awareness in management and can generally help lead to future preventative measures.

Table 6-7 shows an example of a severity level escalation matrix.

Table 6-7 Severity Level Escalation Matrix

Elapsed Time	Severity 1	Severity 2	Severity 3	Severity 4
5 minutes	Network operations manager, tier 3 support, director of networking			
1 hour	Update to network operations manager, tier 3 support, director of networking	Update to network operations manager, tier 3 support, directory of networking		
2 hours	Escalate to VP, update to director, operations manager			
4 hours	Root cause analysis to VP, director, and operations manager, tier 3 support, director of networking Unresolved requires CEO notification	Escalate to VP Update to director and operations manager		
24 hours			Network operations manager	
5 days				Network operations manager

The following is an example problem reporting and escalation plan:

1. Tier 1and Tier 2 Support

   For all international sites, the first and single point of contact for all reported IP Telephony problems is the International Response Center (IRC). For all corporate, remote, and field sales sites in the Americas, the first and single point of contact for all reported IP Telephony problems is the Telecom Help Desk.

   Within international sites, the IRC will log information for the problem and attempt to resolve the reported issue. The IRC is also responsible for alerting managers and technicians responsible for problem resolution through paging services when appropriate. Technicians providing support will be contacted through on-duty pagers. Managers will be contacted through group pager lists. If the issue cannot be immediately resolved by the IRC, the problem trouble ticket will be escalated and transferred to Network and Telecom Operations for tier 2 support.

Within the Americas, The Telecom Help Desk will log information for the problem and attempt to resolve the issue for both tiers. The Telecom Help Desk also contacts the TRC to provide the appropriate paging and notification services, as appropriate. If the problem is not resolved immediately, then tier 2 support processes will be invoked and applied by the Telecom Help Desk.

2. Tier 2 Escalation and Support

   Once a reported problem has been escalated to tier 2, Telecom Operations (within the Americas) or Network and Telecom Operations (international) will continue to be responsible for the trouble ticket until it is resolved. If the issue can be isolated to a LAN or WAN-specific issue, the appropriate LAN and WAN Operations groups will be contacted for problem isolation and resolution at tier 2 level. Similarly, if the issue can be isolated to an NT Operations-specific issue, NT Operations will be contacted for problem isolation and resolution. All groups will coordinate with the Telecom Operations or Network and Telecom Operations for resolution of the problem.

3. Tier 3 Escalation and Support

   Once Telecom Operations (within the Americas) or Network and Telecom Operations (international) have determined that they can not resolve the issue in collaboration with LAN, WAN, and NT Operations, the problem will be escalated to tier 3 support and a problem case will be opened with the World Wide Technical Assistance Center (WW TAC).

   It is expected that the WWTAC will work with Telecom Operations to resolve the issue as quickly as possible to include escalation to the vendor’s other support team and On-Site Support (OSS) team if required.

   In the event that the issue appears to be NT platform-related or the WWTAC dispatches On Site Support (OSS) to replace a server running the NT Operating System, NT Operations will be contacted for coordination and collaboration in resolving the problem. Telecom Operations or Network and Telecom Operations will have the responsibility for contacting NT Operations.

Figure 6-1 shows a graphical representation of the escalation process.

Figure 6-1 Escalation Process

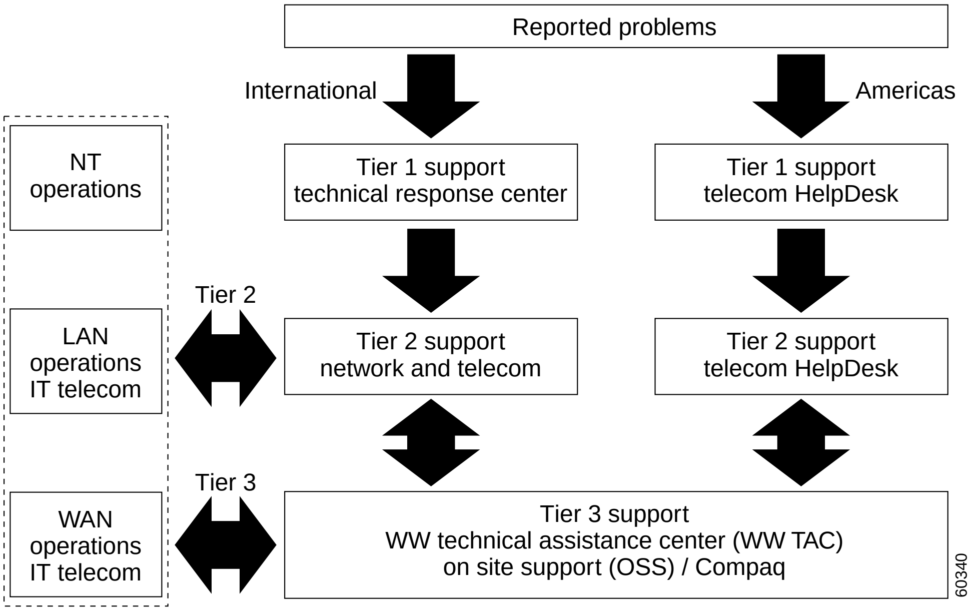
Reported problems
International
Americas
NT operations
Tier 1 support technical response center
Tier 1 support telecom HelpDesk
Tier 2
LAN operations IT telecom
Tier 2 support network and telecom
Tier 2 support telecom HelpDesk
Tier 3
WAN operations IT telecom
Tier 3 support WW technical assistance center (WW TAC) on site support (OSS) / Compaq
60340

Limitations of Support and Responsibilities of Telecom Operations, Local Site, Network and Telecom Operations

P1 and P2 problem resolution, including break-fix support, shall be provided in four hours or less. This resolution is also contingent upon having on-site access provided by local personnel at the site. If this can not be provided at the local site, then the problem case will be re-classified as a P3 problem case by the WW TAC and Cisco IT and a one-day problem resolution interval objective will apply.

Dual CallManagers at every Site for P1 and P2 Support

NT Operations will provide P1 and P2 support to only those sites that will use dual primary and backup CallManagers.

NT Operations will provide best effort support for those cases where LAN or WAN problems prevent NT Operations from accessing the site. For example, if the LAN is down and not reported as a P1, and these problems prevent NT Operations from accessing a CallManager for resolution and repair, NT Operations will downgrade the problem case to the status of the reported priority of the networking problem case. In all cases, NT Operations will provide support equal to the lowest problem case priority of causing restricting remote network access to the CallManager. If NT Operations discovers that an equipment unit part or server requires replacement, and that corporation IT department failed to ensure that 24x7x4 on-site support contracts were purchased for the equipment, NT Operations will not be responsible for 24x7x4 break-fix support.

Security and Access Privileges to CallManagers Servers

Providing P1 and P2 support will also depend on NT Operations having full and complete access to all NT servers running IP Telephony CallManager.

It is the responsibility of IT-Telecom to provide, maintain, notify, domain names, URLs, user logons, and passwords to all IP Telephony CallManager servers. If dial-in remote access is provided at a site through the use of dial-up remote control software, it is the responsibility of IT-Telecom to provide, maintain, notify, and communicate dial-in numbers to the remote access console servers, including user login and passwords. Supplying appropriate remote access client and server software on the IP Telephony CallManagers and to NT Operations is the responsibility of IT-Telecom.

Fault Monitoring

Fault monitoring processes for IP Telephony networks should include device-down monitoring, link down monitoring, and network error monitoring. Defining the processes in relation to IP Telephony components helps to clarify the support and network operation center (NOC) responsibilities and will result in a more efficient and reliable operational environment.

Table 6-8 is an example of fault monitoring service level definitions for device-down monitoring that provides a clear understanding of alerts and how they are handled. The organization may still need additional efforts as defined above to ensure success. In some cases, the organization may have different alert, priority, and escalation paths based on the time of day. For instance, some organizations may not have a 7x24 NOC, but do have 24-hour pager notification to support personnel.

Table 6-8 Device-Down Detection Monitoring

Device Type	Problem Identification Method	Polling or Detection Intervals	Action Taken
CallManager	SNMP polling	5 minutes	Priority 1 ticket created in NOC (priority 3 with CM redundancy)
PSTN gateway	SNMP polling	5 minutes	Priority 1 ticket created in NOC (priority 3 with CM redundancy)
Voice mail gateway	SNMP polling	5 minutes	Priority 2 ticket created in NOC (priority 3 with CM redundancy)
Core switch	SNMP polling and trap notification	5 minutes	Priority 1 ticket created in NOC (priority 3 with CM redundancy)
Edge switch	SNMP polling and trap notification	5 minutes	Priority 1 ticket created
Core router	SNMP polling and trap notification	5 minutes	Priority 1 ticket created (priority 3 with CM redundancy)
Desktop	none	N/A	N/A
IP phone	none	N/A	N/A
Data center switch links	SNMP polling and trap notification	5 minutes	Priority 1 ticket created in NOC or priority 3 with link redundancy)
Core switch links	SNMP polling and trap notification	5 minutes	Priority 1 ticket created in NOC or priority 3 with link redundancy
Core router links	SNMP polling and trap notification	5 minutes	Priority 1 ticket created in NOC or priority 3 with link redundancy
WAN Switch Links	SNMP polling and trap notification	5 minutes	Priority 1 ticket created in NOC or priority 3 with link redundancy
Edge Switch trunks	SNMP polling and trap notification	5 minutes	Priority 1 ticket created (no redundancy)
Edge switch user ports	none	N/A	User reports problem
Desktop	none	N/A	User reports problem
IP Phone	none	N/A	User reports problem

Table 6-9 is an example of service level definitions for network errors that provide a clear understanding of error alerts and who is responsible, how the problem will be identified, and what will happen when the problem occurs. The organization may still need additional efforts as defined above to ensure success.

Table 6-9 Network Errors

Error Category	Detection Method	Threshold	Action Taken
Software errors (software-forced crashes)	Daily review of syslog messages using SYSLOG viewer. Done by tier 2 support.	Any occurrence for priority 0, 1, and 2. Over 100 occurrences of level 3 or above.	Review problem, create trouble-ticket and dispatch if new occurrence or if problem requires attention.
Hardware errors (hardware-forced crashes)	Daily review of syslog messages using SYSLOG viewer. Done by tier 2 support.	Any occurrence for priority 0,1 and 2. Over 100 occurrences of level 3 or above.	Review problem, create trouble-ticket and dispatch if new occurrence or if problem requires attention.
Protocol errors (IP routing protocols only)	Daily review of SYSLOG messages using SYSLOG viewer. Done by tier 2 support.	10 messages/day of priority 0, 1 and 2. Over 100 occurrences of level 3 or above.	Review problem, create trouble-ticket and dispatch if new occurrence or if problem requires attention.
Media control errors (FDDI, POS and Fast Ethernet only)	Daily review of syslog messages using SYSLOG viewer. Done by tier 2 support.	10 messages/day of priority 0, 1 and 2. Over 100 occurrences of level 3 or above.	Review problem, create trouble-ticket and dispatch if new occurrence or if problem requires attention.
Environmental messages (power and temp)	Daily review of syslog messages using SYSLOG viewer. Done by tier 2 support.	Any message.	Create trouble ticket and dispatch for new problems.
Accuracy errors (link input errors)	SNMP polling at 5-minute intervals. Threshold events received by NOC.	Input or output errors. 1 error in any 5 minute interval on any link.	Create trouble ticket for new problems and dispatch to tier 2 support.

Device File Backups and Recovery

The potential always exists for device file corruption or loss due to hardware problems. The organization should prepare for this by having a defined process for backing up network devices and CallManager systems. Most network devices including IOS gateways and MGCP gateway devices support TFTP for configuration file backups. DT-24 gateways keep their configuration on the CallManager so if a new one is required, a new MAC address must simply be configured on the CallManager. The CallManager system may require a system software load, as well as a set of configuration files for recovery, so these should be on hand in case of a needed recovery. CallManager backups can be done using a supported tape drive backup or network backup to another system.

The organization should define when backups occur, who performs the backup, where the backup tape or directory can be found, and who is responsible for recovery. The organization should have a service plan for device backups and recovery.

Table 6-10 shows the file backup and recovery plan that can be used within the organization.

Table 6-10 Network File Backup and Recovery Service Plan

Device	Backup Method	Backup Responsibility	Backup Period	Recovery Responsibility
CallManager	Network—CallManager utility to back up server XX	Tier 2 NT operations (no remote CallManagers backed up)	Full backup daily at 06:00	Tier 2 NT operations
IOS gateway	Network TFTP	Data network tier 2 operations	After configuration changes	Tier 2 data network operations
IP phone	None—information stored on CallManager	N/A	N/A	N/A
DT-24 gateway	None—information stored on CallManager	N/A	N/A	N/A
Other network devices	Network TFTP	Data network tier 2 operations	After configuration changes	Tier 2 data network operations

In addition to CallManager backups, the organization should be clear on who manages CallManager configuration and change. As the CallManager is an NT device, NT server administration groups may be responsible for this activity. Responsibilities for CallManager configuration may include the following:

- Tracking, managing, and archiving all CallManager change control logs
- Maintain CallManager configuration consistency
- Maintain CallManager software consistency, including versions and patches
- Backup schedule
- Backup recovery procedure

Hardware Replacement and On-site Assistance

The organization should review the hardware replacement process and on-site assistance support capabilities to help ensure it meets new IP Telephony availability goals and business requirements. This is especially true when you have new IP Telephony availability targets and a larger number of device types in the network.

Two factors contribute to hardware availability:

- MTBF

  Mean time between failure is the failure rate of the device. Cisco maintains theoretical availability based the Telcordia (formerly Bellcore) “Parts Count Method”. Individual device failure rates are taken from the latest issue of the Telcordia specification TR-332 tables. The final predicted MTBF value is two times this calculated result, based upon Cisco’s historical experience.

  Once a product has been released to the market, field replacement data is collected to compute a demonstrated MTBR value that is updated each month on a rolling 3-month basis. This MTBR number is based on the running install base and the number of Returned Material Authorizations (RMAs) reported by Cisco Metrics database. MTBR is calculated by dividing the total install base operating hours for 3 months by the number of returned units for the past 3 months. This information is available from the Cisco account team upon request.

- MTTR

  Mean time to repair is the organization’s responsibility to implement support contracts or to stock spares so that overall service and support goals are met. Table 6-11 shows expected availability with two theoretical devices given different MTTR values.

Table 6-11 Availability and MTTR Values

Device	MTBF	Ongoing Hardware Availability with 4-hour MTTR	Ongoing Hardware Availability Percentage with 24-hour MTTR	Availability with 4-hour MTTR and Device Redundancy
Device 1	50,000 hours	99.992%	99.952%	99.99999%
Device 2	10,000 hours	99.996%	99.976%	99.99999%

The organization has several choices for hardware replacement. The first and most common is a Cisco SmartNET support contract that provides approximately one business day replacement of returned hardware or components. This is often combined with a sparing program to improve the MTTR. Another method may be a four hour support contract where on-site assistance and hardware replacement is four hours.

A sparing program should also have several components to ensure success. Some organizations have even set up groups called Materials Management to handle the sparing service. The following sparing program components are helpful to ensure overall success:

- Maintain spares inventory and location
- Periodic review to ensure hardware components meet current software requirements
- RMA tracking to ensure broken components are returned and new spares are replaced in spares inventory
- Hardware handling procedures to ensure proper grounding and module insertion

On-site assistance and hours of coverage should also be a concern in a hardware sparing program. Cisco support can provide four hour 7x24 service contracts to help in this area. If the organization does not currently have 7X24 support or a 7X24 NOC, they may need to consider how problems will be identified and repaired outside of business hours.

Disaster Recovery

A disaster recovery plan covers both the hardware and software required to run critical business applications and the associated processes to transition smoothly in the case of a disaster. The IP Telephony application may then create new implications for disaster recovery as the voice capability is closely tied to the data network.

Management awareness is the first step in Disaster preparedness. To obtain the necessary resources and time required from each area of the organization, senior management has to understand and support the business impacts and risks. Like many large projects, disaster recovery requires strong commitment from senior management to support the implementation of a disaster recovery plan and have the required processes and resources in place. There are several key tasks required to achieve management awareness:

- Identify and make a list of the top ten potential disaster types and impact on the business—for example fire, earthquake, and storm
- Review the list with senior management and build management awareness
- Obtain signoff from management for disaster recovery planing process and funding

Key elements of disaster recovery planing include the following:

- Establish a planing group
- Perform risk assessments and audits
- Establish priorities for network and applications
- Develop recovery strategies
- Up-to-date inventory and documentation of the plan
- Develop verification criteria and procedures
- Implementation

Resiliency and backup services form a key part of disaster recovery and require diligent reviews to meet the criteria for disaster recovery. A high availability design is very often the foundation for disaster recovery and can be sufficient to handle some minor or local disaster. Key tasks for resiliency planning and backup services include the following:

- Resiliency assessment—identification of gaps and risks
- Review of backup services
- Implementation of network resiliency and backup services

Vendor Support Services add a strong value to disaster recovery planing. For example, specific managed "hot standby sites" or on-site services with rapid response times. Many organizations use specific disaster recovery services from organizations such as Sunguard to provide adequate backup data services. Key questions on the topic of vendor support include:

- Are support contracts in place?
- Has the disaster recovery plan been reviewed by the vendors and are the vendors included in the escalation processes?
- Does the vendor have sufficient resources to support the disaster recovery?

Metrics and Reporting

Any defined service level goal should always be measurable. This allows the organization to measure service levels, identify root-cause service issues, and make improvements that are aimed at specific targets. Overall, metrics are simply a tool that allows network managers to manage service level consistency and to make improvements according to business requirements.

Many organizations do not currently collect availability, performance, and network metrics. The primary reasons are the inability to provide complete accuracy, cost, network overhead, and available resources. Fortunately, many organizations have been able to create low cost, low overhead metrics that may not provide complete accuracy, but do satisfy the primary goal of service level management.

Measuring availability and performance is one area typically neglected in service level metrics. Organizations that are successful with these metrics use two fairly simple methods. One method is to send ICMP ping packets from a core location in the network to edges. Performance can also be obtained using this method. Organizations that are successful with this method also group like devices into "availability groups", such as LAN devices or domestic field offices. This is also attractive because organizations usually have different service level goals for different geographic or business critical areas of the network. This allows the metrics group to average all devices with the availability group to obtain a reasonable result.

The other successful method of calculating availability is done using trouble-tickets and a measurement called IUMs (impacted user minutes). This method tabulates the number of users that have been affected by an outage and multiplies it by the number of minutes of the outage. When expressed as a percentage of total minutes in the time period, this can be easily converted to availability. In either case, it can also be helpful to identify and measure the root cause of downtime so that improvement can be more easily targeted. Root cause categories may include hardware problems, software problems, link or carrier problems, power or environment problems, change failures, and user error.

Measuring performance and jitter across a large network infrastructure on a consistent basis is still a difficult task. Several tools, including Internet Performance Monitor (bundled with CiscoWorks2000), can help measure key paths across the network. Other tools such as Ganymede's Chariot can help measure performance and jitter with VoIP traffic in real time to specified servers or workstations.

Other measurable reactive support goals include reactive service response time by call priority, problem resolution goals, or MTTR and problem escalation time. Reactive support goals are normally measured by generating reports from help desk databases. This requires database fields for the time a call was initially reported, (or entered into the database), the time the call was accepted by an individual working on the problem, the time the problem was escalated, and the time the problem was closed. These metrics may require management influence to enter problems consistently in the database and update problems in real time. In some cases, organizations are able to automatically generate trouble-tickets for network events or E-mail requests. This helps provide accuracy for identifying the start time of a problem. Reports generated from this kind of metrics will normally sort problems by priority, work group, and even individual users to help determine potential issues.

Another measure of service level management success is the service level management review. This should be done whether or not service level agreements are in place. Service level management review should be done by holding a monthly meeting with individuals responsible for measuring and providing defined service levels. User groups may also be present when service level agreements are involved. The purpose of the meeting is to then review performance of the measured service level definitions and to make improvements.

Each meeting should have a defined agenda that includes review of measured service levels for the given period, review of improvement initiatives defined for individual areas, current service level metrics, and then a discussion of what improvements are needed based on the current set of metrics. Over time, the organization may also trend service level compliance to determine the effectiveness of the group. This process is not unlike a quality circle or quality improvement process. The meeting helps target individual problems and determine solutions based on root cause.

Proactive Service Elements

Proactive service elements define unique processes that are part of a required business practice or are recommended for overall network manageability, availability and performance. One area of critical success in IP Telephony networks is also maintaining configuration, version and module consistency across the infrastructure to help reduce complexity and facilitate rapid resolution of problems when they occur. Each of the proactive service level elements discussed in this section includes a discussion of the recommended service level parameters and, in some cases, examples of service level elements that can be used as a template when creating an IP Telephony operations plan within the organization.

Change Management

Change management is a process typically implemented within enterprise and service provider organizations to help ensure network consistency and change success. Managing change properly improves availability by helping to ensure the following:

- Network change is in the best interest of the organization
- The change will be successful

- The change will create the least amount of user or customer impact
- The effects of the change are captured and understood prior to the change

Change management is one of the most important processes related to network availability. Without change management, organizations can expect to experience much higher unscheduled downtime, less network consistency, and a much higher rate of change failure. Network inconsistency also leads to increased network failure and extended repair time. For this reason, change management is often considered a primary targeted practice for new and growing networks.

For more detailed information on change management, go to the following location: http://www.cisco.com/warp/public/126/chmgmt.html.

Move, Add, Change Process

Experience with larger IP Telephony implementation shows that the organization must prepare a move, add, change, (MAC) process for IP phone provisioning and removal. The process will help ensure that consistent implementation quality is being used, helps ensure that proper IP Telephony resources are being configured, and helps track current provisioning and voice IP phone service levels. If the organization has multiple sites, this process is critical in providing timely service for phone MAC changes. The organization may also have different implementation types if analog phones are being connected to VoIP gateways.

The organization must first identify the support group responsible for phone MAC changes. The business requirements may also help determine the expected service times and turnaround time for phone MACs. The identified organization must then determine staffing levels and the procedure used to perform phone MACs.

The organization needs to collect and maintain information on all phones in a database or spreadsheet. This information typically includes phone, phone number, and user information.

Telephone Number and Dial Plan Management

Most organizations will need to manage phone number information and dial plans. Typically phone information is archived into a telecom database or possibly a spreadsheet. This information can be used for phone directories, track phone numbers and ranges, CallManager scaling, billing, phone locations, phone re-configuration, or replacement and inventory management. Some of this information may be the default for all phones, but may be added so that the entire phone configuration is covered. The information that should be tracked by user is shown in Table 6-12.

Table 6-12 Telephone Number and Dial Plan Management

Database Record Information	Purpose
User ID./User name	Identify phone user
Phone Number	Identify phone number
Phone location	Identify phone location
CallManager	Identify configured CallManager associated w/phone
Cluster	Identify configured CallManager cluster
Phone Model	Identify phone model (may want serial # as well)
Phone button template	Standard or configured IP phone button template used by user for phone

Table 6-12 Telephone Number and Dial Plan Management (continued)

Database Record Information	Purpose
Device Pool	Device pools allow the organization to scale and simplify the distribution of Cisco CallManager redundancy groups. A device pool can include region, date/time group, and Cisco CallManager redundancy group.
Calling search space	The list of partitions the user can search when placing call. Dialing a number outside the search space results in a busy signal.
Partition	User/phone reachability characteristics including phones, route-patterns and directory numbers. An example might be users in Building D with partition name Bldg-D.

The organization may also require some standards for phone and dial plan management. The following are some standards that may be used:

- Standard formula for deriving directory number from telephony number management (“5-digit dial prefix + extension” or “7 + 2620")
- Standardize on company name “organization” as default global partition name for all company IP phones
- Standardize on names for calling search space across clusters
- Use default phone button template that comes with product
- Standardize on “Nickname Lastname” for line display
- Standardize phone description based on Cisco product ID

Ongoing dial plan management may also be a consideration for the organization that is scaling phone numbers. Many organizations now have dial plans that include support for abbreviated dialing, such as four or five-digit extensions. The organization has to track the digit identifiers for each region or phone group and periodically determine scaling issues. In many cases, the organization may need to move from 4-digit dialing to 5-digit dialing, and maybe even 6 digit dialing, to accommodate the quantity of phones within the organization.

IP/DHCP/DNS/TFTP Management

The organization will also need a support plan for IP address space, DHCP servers, and TFTP servers. By now the organization has defined how IP addressing will work with phones. This is typically done using one of three methodologies as follows:

- Assigning IP addresses using the same subnet as data devices
- Modify the current IP addressing plan
- Create a separate IP subnet for IP phones

Once the methodology has been defined, you can allocate the IP addresses for IP phones. Many organizations also choose to have standard IP address ranges and DNS names associated with the IP addresses. The information must then be configured into the primary and possibly backup DHCP server. The organization will also need TFTP server capability for backup storage. The following responsibilities should be considered for IP/DHCP/DNS/TFTP management.

IP Address Management

IP address management is typically done by the data networking group. These activities should already be accomplished for data network management.

- Identify Phone IP addressing methodology
- Make network configuration changes necessary for IP phones
- Manage/allocate/track IP address ranges for IP phone and data device subnets
- Create IP phone DNS standard for IP phones allocated via DHCP

DHCP Management

DHCP management normally includes NT server administration requirements and DHCP application support requirements. The DHCP service should support option 150, which allows the IP phone to download the IP address from the DHCP server and forward a request to the CallManager to also download the configuration. Without this option, you must manually configure the phones. Identify the following processes and owners as part of DHCP management.

- DHCP server design, including DHCP file redundancy
- DHCP server administrator and support for server issues
- DHCP database backup and restore
- DHCP subnet configuration
- DHCP subnet range capacity review
- DHCP allocation configuration
- DHCP problem support

DNS Management

DNS management is typically already part of the data networking processes. DNS servers are used primarily to perform name-IP binding, but are also used for IP address place-holding. This helps identify whether an IP address has been allocated or not. DNS servers should be investigated during the design phase to ensure that the appropriate redundancy is in place to support VoIP as a high availability application. DNS processes that are important under IP Telephony include the following:

- DNS server administration
- DNS server backups and restore
- DNS updates for device MACs or DHCP IP address ranges
- DNS resolution problem support

TFTP Management

TFTP servers are typically used to archive data networking device configuration files including IOS gateways and IP phones. TFTP services can also be configured as part of CallManager cluster. Cisco recommends that IP Telephony TFTP services be done in the cluster configuration however other TFTP servers may be used for general device configuration backups. The following responsibilities are typical of TFTP server administration:

- File cleanup
- TFTP server backups and restore
- TFTP server problem support

Proactive Network Management Processes

Proactive network management processes identified for IP Telephony network management help to ensure continued voice quality following organization growth and change. Specific goals for proactive network management include the following:

- Maintain or reduce network complexity to ensure ongoing network manageability and supportability
- Maintain or improve IP Telephony performance by ensuring continued scalability and performance of the network and IP Telephony solution over time
- Quality improvement to improve availability and possibly total cost of ownership over time

The goals can be broken down into the following three management areas:

- Capacity/performance management
- Configuration management
- Availability management

Capacity/Performance Management

Table 6-13 defines the proactive processes that may be needed along with a recommended review period.

Table 6-13 Capacity/Performance Management Processes

Process	Goal	Review Period
IP Telephony scalability	Ensure the IP Telephony scalability requirements are being met, including phones per CallManager and CallManager application scalability requirements.	Monthly or quarterly
Network performance baselining	Baseline network on a regular basis to ensure that consistent network performance goals are being met and that adequate capacity exists within potential upgrade cycle time.	Monthly or quarterly
Application profiling	Benchmark new applications and understand network behavior to ensure no impact on IP Telephony.	Before new application deployment
What-if analysis	Ensure network changes will not impact IP Telephony performance and voice quality.	Before major risk network changes occur
Performance exception analysis	• Monitor and report network performance exceptions that impact voice quality such as slow performance, high CPU, high output queue quantities, and high link utilization. • Identify problem and resolve ASAP	Daily

Configuration Management

Configuration management processes help to ensure consistent configurations, software versions, hardware modules and network topology. Consistent configuration, versions and topology help to ensure a less-complex environment over time, which is ultimately easier to support and manage.

A less-complex environment is also expected to have higher availability due to fewer bugs, issues and repair times. Table 6-14 identifies some processes that have been identified for proactive network management in relation to configuration management.

Table 6-14 Configuration Management Processes

Process	Goal	Review Period
Software version control	Ensure consistent software versions and upgrade/change as needed	Monthly
Hardware/topology review	Ensure consistent hardware rev levels, modules and platforms	Monthly or quarterly
Configuration standards	Maintain, review and push out standard device configuration files, related to device access/management, security, device features/behavior	Ongoing
Configuration Audit	Audit configuration changes recorded with TACACS or change management process to ensure consistency	Ongoing

Quality Improvement

The organizations that have the highest availability and consistent performance typically have a quality improvement process. General quality improvement processes such as “6-sigma” or “Total Quality Control” created for manufacturing processes can often be directly applied to network quality. The key to these processes is identifying targets, measuring results associated with targets and making improvements needed to achieve the desired goals. Targets associated with IP Telephony deployments may be availability, voice quality, trouble-ticket reduction, change management success and rapid problem resolution. To be successful with improvements the goal must be measurable and the organization must be committed to setting obtainable goals and achieving results.

Call Data Record Reporting and Billing

Many organizations will want to create call data record reporting capability from the CallManager. Call data records can be used to:

- Determine load balancing across load sharing CallManagers systems
- Troubleshooting (cause codes)
- Internal billing
- Fraud detection (normally requires a third-party application)
- PSTN costing (normally requires a third-party application)

Call data record reporting must be enabled on the CallManager (off by default) to accumulate call data records. The records are stored in the CallManager SQL database. Administrators can also enable a trace function to get real-time call data information. The files also accumulate indefinitely so some regular archival or deletion is needed. The files can be copied via TFTP from the CallManager and imported into spreadsheet applications. Commas delimit all fields.

Billing is normally done using third-party billing for chargebacks. The organization may also purchase a costing engine to estimate PSTN costs for billing and chargebacks. Software is also available for fraud detection where thresholds can be set to determine unusual call volume, calls outside business hours over a defined amount, or unusually high call costs.

Basic Call data records include the following information:

- Call Record ID
- Call origination time stamp and date
- Originating phone number
- Originating IP address
- Destination phone number
- Destination IP address
- Call disconnect time stamp and date
- Raw connect time
- Disconnect code
- Call quality information, including jitter and delay

For more information on configuring call data record functionality, a complete list of data record fields, cause codes, and call data record types, go to the following location: http://www.cisco.com/warp/public/788/AVVID/ts_ccm_301_sec2.htm.

Staffing and Support Model

You may need to revise the staffing and support model to support IP Telephony. Data network administrators, system administrators, and telecom operations may be deployed in the overall support of the IP Telephony solution. Overall, staffing requirements are a factor of the expertise levels, support contracts, and tool automation that the organization has implemented. The organization should however, look at all the requirements for operating the IP Telephony solution and determine resource counts based on individual and administration group capabilities. The organization may also wish to outsource portions of IP Telephony support, for instance remote support or MAC changes.

Table 6-15 can be used to estimate operational requirements for IP Telephony.

Table 6-15 IP Telephony Operational Requirements

Management Area	Operational Requirement	Responsible Organization	Required Hours or Head Count
Fault Management	NOC monitoring including CallManager, network fault, application availability, device monitoring, syslog monitoring	NOC group	
	Tier 2 break/fix support 24x7 and on-site support	Data networking tier 2 support	
	Tier 3 escalation	Data networking tier 3 support	
	Fault management tools	NMS	
	Metrics and reporting	NMS	

Table 6-15 IP Telephony Operational Requirements (continued)

Management Area	Operational Requirement	Responsible Organization	Required Hours or Head Count
Configuration Management	Major MAC changes	Data networking tier 2 support	
	Phone MACs	telecom operations	
	Change Management control	IS change management	
	Change Management planning	Data networking tier 2 support	
	CallManager Administration including backups/recovery	NT operations	
	IP/DHCP/TFTP/DNS management	NT operations	
	IP addressing Plan	Data networking tier 2 support	
	Dial Plan Management	Telecom operations	
	Configuration consistency	Data networking tier 2 support	
	Configuration management tools	NMS	
	Cable plant and data hardware	Data network operations	
Inventory Management	Billing – chargeback	Telecom	
	Inventory management tools	NMS	
	Device inventory	NMS	
	Fraud detection	Telecom	
Security Management	User ID/password admin	Security operations	
	Physical security	Facilities	
Performance Management	Performance management tools	NMS	
	CallManager utilization	NT Operations	
	Data network baselining	Data networking tier 3 support	
	Change management what-if analysis	Data networking tier 3 support	
	Performance exception analysis and resolution	Data networking tier 3 support	
	Application profiling	Data networking tier 3 support	

Documenting and Approving the Operations Support Plan

The final step of the operational support planning process will be to document and approve the operational support plan. The plan should include all support responsibilities, parties involved, reporting requirements, expected service levels and service level agreements. It can be expected that in the first few months of the IP Telephony pilot and operation that several amendments will be made to the operational support plan. Cisco recommends weekly meeting for initial deployment periods for revisions and additions.

Service level agreements may also be included in the operational support plan. Expected service level agreements for IP Telephony deployments include Voice availability, voice quality, network performance, fault repair times by priority and MAC service level expectations.

Network Management

The most commonly used reference model in describing network management is the ISO Network Management Model. The model outlines five functional areas in dealing with various aspects of managing a network infrastructure. Those functional areas include fault, configuration, accounting, performance, and security, which are referred to FCAPS in the industry. The model and its functional areas allow well-defined scope and objectives to be defined in the evaluation and implementation of a network management solution. Users involved in the process of implementing management system can focus on a functional area or a combination of them.

Functional Areas of Network Management

This section provides a brief overview of each of the functional areas in the reference model:

- Fault Management
- Configuration Management
- Accounting Management
- Performance Management
- Security Management

Fault Management

Fault management detects problems on network elements in the infrastructure. Hardware and software problems can lead to disruption or degradation of network services. Employees in the NOC rely on a management system to inform them of faults that are detected on network elements. A properly configured network element is capable of forwarding system messages and notifications to a management system. You can take appropriate action to minimize the impact on network availability based on the severity of faults reported by the elements. Fault management should be implemented properly to ensure the effectiveness of fault detection and the timely resolution of network-related issues.

Configuration Management

Configuration management manages configuration files, software, addresses, and detailed inventory information of network elements. An up-to-date configuration management system can significantly reduce the amount of time spent troubleshooting network activities. Complete and detailed inventory information also provides tremendous value in the planning and budget allocation stages of a network rollout.

Accounting Management

With increasing user and application traffic in the network, it is important for an organization to track the use of network resources. A thorough understanding of traffic profiles allows network planners to prioritize and allocate sufficient bandwidth for different applications. Critical and delay-sensitive applications should receive a higher priority over regular user traffic to satisfy their time and bandwidth requirements.

Accounting data collected from network elements typically ranges from simple to detailed records on traffic statistics. This data can be used for planning, or as an input to the billing system for enterprises who need to implement chargeback to internal and external entities.

Performance Management

Performance management measures the performance levels of different components of an IT infrastructure. Satisfactory performance levels are dependent on network, system, and application components of the overall infrastructure. Measuring the performance of different components is crucial and can be accomplished by first defining specific metrics and then collecting them on a regular basis.

You can measure the collected performance data against performance objectives or a service level agreement (SLA) established within the organization. Historical performance data also serves as a baseline of normal operating characteristics and utilization of network elements and end systems. Performance data gathered on an ongoing basis provides network engineering with the ability to effectively plan for growth in the infrastructure.

Security Management

Security management involves the various aspects of controlling access to resources in the infrastructure. You can institute security measures to ensure that only authorized users have access to network platforms, systems, and sensitive business information. You should have a security policy in place before performing any technical evaluation of security protocol or product. The proposed security feature or solution should satisfy the requirements outlined in a security policy. Thoroughly test the implementation of security features before actual deployment.

Network Management Solutions

CiscoWorks2000

The CiscoWorks2000 product line provides network management solutions focused on two key administrative areas: wide-area and local-area switched network management. Each of these areas integrates new and existing Cisco applications into an enhanced web-based management solution.

For information on how to use Internet technology to integrate management tools and integration, go to the following location: http://www.cisco.com/warp/public/cc/pd/wr2k/ent/tech/bmi_wi.htm.

The CiscoWorks2000 product line includes the solutions listed below. To view detailed information about any of these products, go to the following location: http://www.cisco.com/warp/public/cc/pd/wr2k.

- VPN/Security Management Solution
  - VPN Monitor
  - Resource Manager Essentials
  - CiscoView
  - Cisco Secure Policy Manager
- LAN Management Solution
  - Campus Manager
  - Device Fault Manager
  - Content Flow Monitor
  - CiscoView
  - Resource Manager Essentials
- Routed WAN Management Solution
  - Access Control List Manager
  - Internetwork Performance Monitor
  - CiscoView
  - Resource Manager Essentials
- Service Management Solution
  - Service Level Manager
  - Management Engine 1100
  - CiscoView
- Campus Bundle for AIX or HP-UX

In addition to the above mentioned solutions, the CiscoWorks2000 product line also features the following advanced applications:

- User Registration Tool (URT)
- CiscoWorks2000 Voice Manager (CVM)

See Table 6-23 for a list of available network management tools, and the IP Telephony components supported for each tool.

Third-party Applications

Network Management Platforms

Network management platforms discover network devices and poll them for information. Some common network management platforms include:

- HP OpenView Network Node Manager
- Computer Associates Unicenter

- Tivoli NetView
- Aprisma Spectrum

Each network device is represented by a graphical element on the management platform's console. Different colors on the graphical elements represent the current operational status of network devices. Network devices can be configured to send notifications called Simple Network Management Protocol (SNMP) traps to network management platforms. Upon receiving the notifications, the graphical element representing the network device changes to a different color depending on severity of notification received. The notification, usually called an event, is placed in a log file.

It is particularly important that the most current Cisco Management Information Base (MIB) files be loaded on the SNMP platform to ensure that the various alerts from Cisco devices are interpreted correctly. Cisco publishes the MIB files for managing various network devices on the cisco.com website. at the following location: http://www.cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml.

Event Correlation Tools/Manager of Managers (MOMs)

With an increasing number of network elements and complexity of network issues, you may want to consider an event management system capable of correlating different network events such as SYSLOG, TRAP, and log files. This architecture behind an event management system is sometimes referred to as a Manager of Managers (MOM) system because it aggregates information from the other network management system stations on the network and can display the NOC to be more proactive and effective in detecting and diagnosing network issues. Event prioritization and suppression allow network operation personnel to focus on critical network events. Correlation tools currently available include:

- Cisco InfoCenter
- MicroMuse Netcool
- Veritas Nerve Center
- Smarts In Charge
- HP OpenView ECS
- Cisco Device Fault Manger/Voice Health Monitor
- Aprisma Spectrum

Performance Tools

Various solutions are available in the marketplace to address the needs of performance management for enterprise environments. These systems are capable of collecting, storing, and presenting data from network devices and servers. The web-based interface on most products allows the performance data accessible from anywhere in the enterprise. Some commonly deployed performance management solutions include:

- CiscoWorks2000 Internetwork Performance Monitor (IPM)

  http://www.cisco.com/warp/public/cc/pd/wr2k/nemo/index.shtml
- Concord Network Health

  http://www.concord.com/products/ehealth/nethealth/nethealth.htm
- InfoVista VistaViews

  http://www.infovista.com/Pages/products/intr_vv.html
- SAS IT Service Vision

  http://www.sas.com/products/itsv/index.html

- Trinagy TREND
  http://www.trinagy.com/products/trend.asp
- MultiRouter Traffic Grapher (MRTG)
  http://mrtg.hdl.com/mrtg.html

Network Management Architecture

The successful implementation of a network management system depends on having a well-defined architecture. Implementing a system with an architecture in place leads to a successful deployment. Network management architecture is the collection of user requirements, processes, and tools for performing certain operation functions. With the increasing number and complexity of networking technologies, it is important to fully understand the components of the network management architecture in the early phase of an implementation. (For a reference architecture that Cisco Systems believes should be the minimal solution for managing a data network, see "NMS Reference Architecture" section on page 6-62.)

A network management system is capable of performing management functions associated with operating a network infrastructure. The initial step in implementing a system starts with a scope and a clear definition of the function that needs to be addressed by the management tool.

Identifying the specific network elements and technologies that need to be managed along with expected inputs/outputs from the tool will increase the chance of building an efficient and effective network management system.

Collecting user requirements is important to ensure that the new management system meets the expectations of end users and is capable of helping them to complete network-related tasks. Well-defined processes also need to be identified and defined to facilitate the communication between multiple user groups within the organization.

Selection Criteria for Management Tools

Network management tools are offered by a wide range of vendors for managing components of the IT infrastructure. Equipment vendors typically offer an element management system (EMS) to manage vendor-specific devices. You can expect an EMS to work seamlessly with equipment from a particular vendor.

A general-purpose network management platform from various vendors augments the functionality available on a vendor-specific EMS. A network management platform is capable of interacting with equipment from multiple vendors by having support for a vendor-specific management information base.

The choice of operating systems (OS) for an element management system and a network management platform depends on several internal and external factors. In-house expertise on a particular operating system is often required to perform routine maintenance tasks such as backup and restore of a system.

The software release schedules and support for different OS platforms vary between equipment vendors. It is important to know and understand the OS platforms to minimize the risk of having network elements deployed without any tools to manage them.

Integration

With the acceptance of web technologies for application building, a number of new and innovative options for integration of NMS tools have emerged: Extensible Markup Language (XML), Lightweight Directory Access Protocol (LDAP), and other new database techniques which are now open standards.

For more information on integrating CiscoWorks2000 with NMS platforms, go to the following location: http://www.cisco.com/warp/public/cc/pd/wr2k/tech/cwnms_tb.htm.

For information on how to use Internet technology to integrate management tools and information, go to the following location: http://www.cisco.com/warp/public/cc/pd/wr2k/ent/tech/bmi_wi.htm.

High Availability for Network Management Systems

Fault tolerance or the ability to recover automatically from software, hardware, and network failures is critical to achieving a highly available network management system. Each operating system and vendor has specific implementations to guarantee high system availability. Disk mirroring and clustering are often used to provide non-stop services for applications and data availability. A hot standby system automatically resumes the tasks of a failed system upon detecting the lack of heartbeat messages from the main system.

Organizations with requirements for a highly available management system should work closely with the system, software, and equipment vendors to understand the feasibility of such a system for their environment.

Network Management Platform Base Requirements

Network elements and end systems provide remote manageability by supporting several management features. One of the most common management features available on an element or system is SNMP. A network management system can issue administrative commands using SNMP to remotely perform administrative and troubleshooting tasks on network elements and end systems. Remote manageability significantly reduces the cost of running the IT infrastructure within an organization.

The management of SNMP data from the network is typically done using an SNMP platform such as HP OpenView Network Node Manager or Tivoli NetView. For the platform to correctly interpret the vendor-specific data and messages from the network, the vendor's MIBs must be loaded into the system.

It is particularly important that the most current Cisco MIB files be loaded on the SNMP platform to ensure that all Cisco devices can be monitored by the SNMP platform. The SNMP platform will report errors indicating unprocessed messages if all of the correct MIBs are not loaded.

Table 6-16 lists the MIBs that must be loaded in order to manage the Cisco CallManager. The first six MIBs in the table must be loaded before CISCO-CCM-MIB.my or the system will generate an error when the CallManager MIB is loaded.

Table 6-16 Minimum MIB Requirements for Managing Cisco CallManager

MIB	Description
SNMPv2-SMI.my	SNMPv2 structure of management information
CISCO-TC.my	Cisco MIB textual conventions
SNMPv2-TC.my	SNMPv2 textual conventions
SNMP-FRAMEWORK-MIB.my	SNMP framework MIB from RFC-2271
SNMPv2-CONF.my	SNMPv2 conformance MIB file
CISCO-SMI.my	Cisco enterprise structure of management information

Table 6-16 Minimum MIB Requirements for Managing Cisco CallManager (continued)

MIB	Description
CISCO-CCM-MIB.my	Cisco CallManager MIB file
CISCO-CDP-MIB.my	Cisco Discovery Protocol MIB file

Table 6-17 lists additional MIBs that are useful to ensure that Cisco devices can be monitored by the SNMP platform.

Table 6-17 Additional MIBs for Monitoring Cisco Devices

MIB	Description
CISCO-ENVMON-MIB.my	environmental monitoring
CISCO-FRAME-RELAY-MIB.my	Frame Relay monitoring
CISCO-PROCESS-MIB.my	CPU and processing monitoring
CISCO-RTTMON-MIB-120_5_T.my	jitter, delay, and packet loss monitoring
CISCO-STACK-MIB.my	Catalyst switch monitoring
CISCO-STP-EXTNESIONS-MIB.my	spanning tree monitoring
CISCO-SYSTEM-MIB.my	systems monitoring

Syslog Base Requirements

System message logging (syslog) is an additional method of managing network elements. Syslog is one of the best methods for determining the health and status of the network and to proactively recognize issues before they become a problem. The system messages provide information on the operational status of devices at the device level, protocol level, and interface level.

Any effective network management system must include a syslog implementation. Each system message is defined by a severity level, and includes a text description of the problem encountered. Corrective actions can be taken by network operations upon receiving messages with the highest severity. A severity level code is a single digit from 0 to 7. The lower the number, the more serious the situation.

All Cisco devices in the network, with the exceptions of IP phones, support syslog messages. Devices should be configured to send their default system messages (usually level 6) to the CiscoWorks2000 server.

To ensure a simple and manageable syslog implementation, any additional syslog servers such as Cisco's Network Analysis Toolkit (NATkit) or other UNIX servers should remote-mount the CiscoWorks2000 server to access the syslog files.

To use a standalone syslog server, follow the configurations guidelines which are defined at the following location:
http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/12cgcr/fun_c/fcprt3/fctroubl.htm - xtocid1767124.

Network Time Protocol

The Network Time Protocol (NTP) provides a common time base for networked routers, servers, and other devices. A synchronized time allows for correlating syslog and Cisco IOS debug output to specific events. For example, call records for specific users can be found within one millisecond. NTP is a User Datagram Protocol/Transmission Control Protocol (UDP/TCP)-based protocol documented in RFC-1305.

An NTP network usually gets its time from an authoritative time source, such as a radio clock or an atomic clock attached to a time server. NTP then distributes this time across the network. NTP is extremely efficient; no more than one packet per minute is necessary to synchronize two machines to within a millisecond of one another. An economical time source can be an Ethernet-attached global positioning satellite (GPS) receiver.

It is highly recommended that the NTP and Service Time Stamp features of Cisco IOS be implemented on the network in order to facilitate coordination of syslog and SNMP events.

For an explanation of NTP, go to the following location: http://www.eecis.udel.edu/~ntp/.

For Cisco IOS configuration examples that incorporate NTP, go to the following location: http://www.cisco.com/univercd/cc/td/doc/product/software/ios11/cbook/csysmgmt.htm#31645.

NTP should be implemented in a hierarchical fashion using client server modes much like routing protocols in order to reduce the load of UDP messages on the network and to ensure orderly maintenance of the protocol. For example:

- The main NTP server should be the source for the core network devices
- The core devices should be the source for distribution devices
- The distribution devices should be the NTP source for access layer devices

Also, the NTP protocol is optimized when three NTP sources are peered together and provide clock down to the network electronics (and NMS). Currently, many UNIX variants support NTP client and server modes if correlation of server events is necessary.

Managing Voice Over IP Network and Element Layers

This section demonstrates specific examples of ways in which various tools may be employed in a network management system that includes Voice over IP (VoIP) management. This section concentrates on the following *network* and *element* management products:

- Resource Manager Essentials 3.1
- Campus Manager 3.0
- Internetwork Performance Monitor 2.2

The "Managing The Application Layer with CallManager" section on page 6-51 covers application and voice management using Cisco CallManager. To ensure you have the most up-to-date information, first refer to the release notes of any specific product under consideration for implementation.

Fault Management

Fault management is necessary to ensure timely failure detection at the interface level, device level, and protocol level. There are a few methods of detecting faults on a network consisting of Cisco routers, LAN switches, and CallManagers. The most common methods are syslog messages and SNMP traps. Most Cisco devices are capable of sending syslog messages to a syslog server. Syslog messages are basically system messages from routers or switches and CallManagers describing different conditions on the devices. SNMP traps forwarded by devices notify you of faulty conditions on a device, such as if an interface goes up or down. The following subsections suggest ways to use syslog and SNMP for managing the network layer of your network. CiscoWorks2000 can help make this job easier.

CiscoWorks2000 Base Requirements

Two of the most important details to ensure a successful CiscoWorks2000 installation are:

- adequate hardware
- correct installation order

Note Installation order is particularly important on NT installations.

Switch and Router Base Configuration Requirements

Switches and routers need a base configuration to allow them to function with CiscoWorks2000 products. SNMP and syslog need to be configured to get full functionality of CiscoWorks2000.

Note If the implementation is to include a Catalyst switch with additional layer 3 cards, and other special purpose line cards, the base configuration must be included for all of the cards in order to communicate with the entire switch.

Catalyst Switch Base Configuration

```
set logging server x.x.x.x (IP address of your CiscoWorks2000 Server)
set logging timestamp enable
set snmp community public read-only
set snmp community private read-write
set snmp trap enable
set snmp trap x.x.x.x (IP address of your SNMP Platform) public
```

Router Base Configuration

```
service timestamps log datetime
logging x.x.x.x (IP address of your CiscoWorks 2000 Server)
snmp-server community public ro
snmp-server community private rw
snmp-server enable traps
snmp-server host x.x.x.x (IP address of your SNMP Platform) public
```

Using Syslog for Fault Management

The CiscoWorks2000 Resource Manager Essentials (Essentials) product allows for capturing syslog files *and* creating custom daily reports to specifically search for network and VoIP errors. The most effective network management implementations review the system log files on a daily basis. Typically, the syslog messages are noted by NOC personnel and escalated to the appropriate level for further action.

Customizing a Daily Syslog Report for VoIP

Figure 6-2 shows a customized daily syslog report for VoIP. This type of report should be accessed every day in CiscoWorks2000 Essentials from a web browser as a 24 hour syslog report. Additional messages can be added to a report based on the environment.

Figure 6-2 Define Custom Report Window

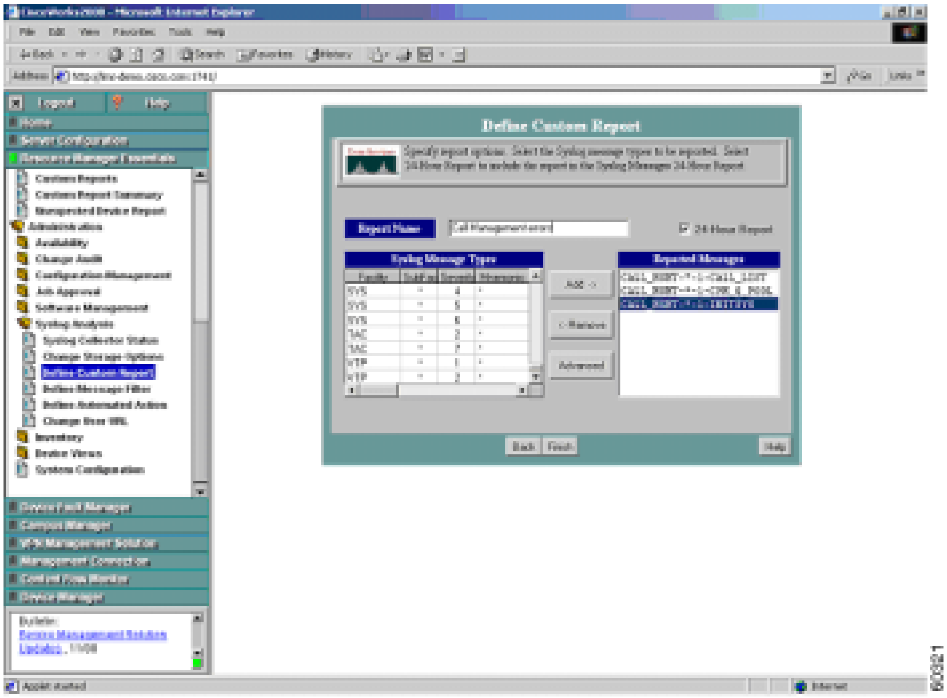

Define Custom Report

Specify report options. Select the Syslog message types to be reported. Select 24 Hour Report to include the report in the Syslog Messages 24 Hour Report.

Report Name: Call Management errors ☑ 24 Hour Report

Syslog Message Types

Facility	SubFacility	Severity	Mnemonic
SYS	*	4	*
SYS	*	5	*
SYS	*	6	*
TAC	*	2	*
TAC	*	7	*
VTP	*	1	*
VTP	*	2	*

Reported Messages

CALL_MGMT-*-1-CALL_LIST
CALL_MGMT-*-1-CPM_Q_POOL
CALL_MGMT-*-1-INITSYS

Table 6-18 lists the call management subsystem error messages used to define the custom report shown in Figure 6-2.

Note System messages are published on the website of each version of the Cisco IOS software. Refer to the Cisco IOS software configuration documentation at the following location: http://www.cisco.com/univercd/cc/td/doc/product/software/.

Table 6-18 Call Management Subsystem Error Messages

Error Message	Explanation
%CALL_MGMT-1-CALL_LIST: [chars]	The specific message text is supplied by the call management software, indicating that internal data was corrupted because of a software error.
%CALL_MGMT-1-CPM_Q_POOL: [chars]	The specific message text is supplied by the call management software, indicating a memory exhaustion condition.
%CALL_MGMT-1-INITSYS: [chars]	The specific message text is supplied by the call management software, indicating an initialization failure. When this message occurs, the call management subsystem is not operational.

The customized report parses the syslog for the voice system error messages (listed in Table 6-18) when the Call Management Errors report, shown in Figure 6-3, is selected. Cisco Systems recommends that this report be run daily by operations personnel and reported to engineering.

Figure 6-3 Syslog 24-hour Report Window

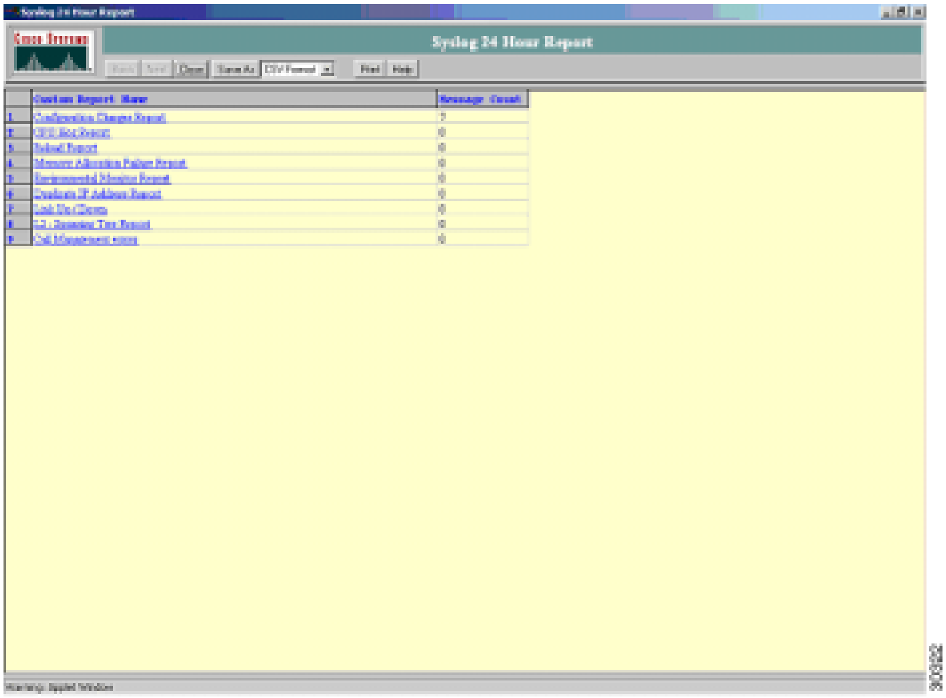

	Custom Report Name	Message Count
1	Configuration Changes Report	7
2	CPU Hog Report	0
3	Reload Report	0
4	Memory Allocation Failure Report	0
5	Environmental Monitor Report	0
6	Duplicate IP Address Report	0
7	Link Up/Down	0
8	L2 - Spanning Tree Report	0
9	Call Management errors	0

Using SNMP Polling and Traps

To verify the operational state of the Cisco devices configured with SNMP, poll the devices. In addition, the devices are capable of sending traps to management stations when events occur. By configuring network devices to send SNMP traps, events can be detected quickly. As a result, a user can diagnose and determine fault conditions and potential problems in the network. Cisco CallManager 3.x can be polled to find out many different types of status information including:

- the number of phones registered
- trunk status
- gateway status

For information on how to poll specific MIBs and set the thresholds, refer to the documentation for your SNMP platform. Cisco's High Availability Services team has additional documentation on MIB recommendations for specific network technologies such as:

- Open Shortest Path First (OSPF)
- Asynchronous Transfer Mode (ATM)
- Integrated Services Digital Network (ISDN)

The above mentioned MIB recommendations should be requested through your Cisco Systems engineer or sales representative.

Using Cisco's Device Fault Manager

The CiscoWorks2000 Device Fault Manager (DFM) provides real-time, detailed fault analysis for Cisco devices. The function of DFM is to:

- monitor Cisco-based networks for a variety of fault conditions

- analyze the fault conditions
- send “intelligent Cisco traps” to notify the user when a problem has occurred that requires attention

DFM can be configured to send the Cisco traps as follows:

- forward to other multi-device, multi-vendor event management systems installed in the network
- send to e-mail/pager gateways
- display in the DFM alarm window

Figure 6-4 shows the DFM Monitoring Console window.

Figure 6-4 DFM Monitoring Console Window—Event Message Description

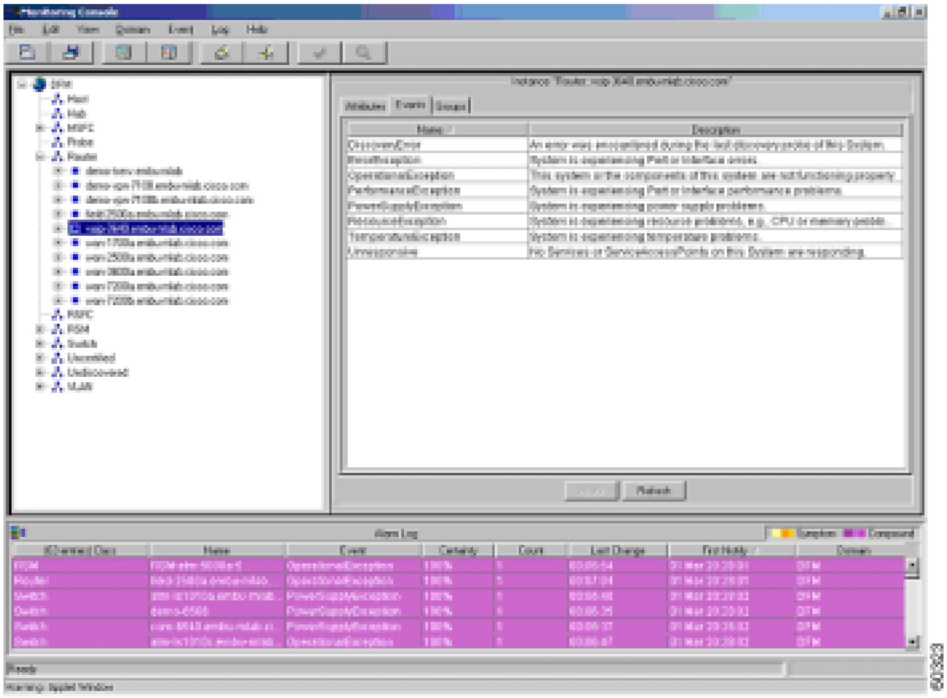

Instance: "Router: [illegible]"

Name	Description
DiscoveryError	An error was encountered during the last discovery probe of this System.
ErrorsException	System is experiencing Port or Interface errors.
OperationalException	This system or the components of this system are not functioning properly
PerformanceException	System is experiencing Port or Interface performance problems.
PowerSupplyException	System is experiencing power supply problems.
ResourceException	System is experiencing resource problems, e.g., CPU or memory proble
TemperatureException	System is experiencing temperature problems.
Unresponsive	No Services or ServiceAccessPoints on this System are responding.

Alarm Log

[illegible] Class	Name	Event	Certainty	Count	Last Change	First Notify	Domain
RSM	[illegible]	OperationalException	100%	1	[illegible]	01 Mar 20:20:01	DFM
Router	[illegible]	OperationalException	100%	1	[illegible]	01 Mar 20:20:01	DFM
Switch	[illegible]	PowerSupplyException	100%	1	[illegible]	01 Mar 20:20:02	DFM
Switch	demo-6506	PowerSupplyException	100%	1	[illegible]	01 Mar 20:20:02	DFM
Switch	[illegible]	PowerSupplyException	100%	1	[illegible]	01 Mar 20:20:02	DFM
Switch	[illegible]	OperationalException	100%	1	[illegible]	01 Mar 20:20:02	DFM

> **Note** When you first open the DFM Monitoring Console window, the right pane is empty. To view descriptions of particular event messages, as shown in the example above, select a device type from the list of devices being managed by DFMand then select the device.

Configuration File Management

It is common practice to make changes to configuration files on network devices such as routers and switches. It is important to have change management procedures in place to prevent a misconfiguration that may result in major outages. A backed up version of the configuration files should be available in case of a failure while performing a change. It is also important to keep track of changes made to existing configuration files.

A Change Summary report consists of the detailed information of changes made to the configuration file. The detailed information includes user name, time, and changes made. This report is useful for tracking changes and accountability. With the availability of web-based configuration management tools in CiscoWorks2000 Essentials, a number of tasks for managing configuration files can be fully automated. The configuration management tool is capable of performing the following tasks:

- Push configuration files from the Essentials configuration archive to a device or multiple devices
- Pull the configuration from the device to the Essentials archive
- Extract the latest configuration from the archive and write it to a file
- Import the configuration from a file and push the configuration to devices
- Compare the last two configurations in the Essentials archive
- Delete configurations older than a specified date or version from the archive
- Copy the startup configuration to the running configuration

Figure 6-5 shows an example of a configuration report for a Cisco Catalyst switch with IP phones connected.

Figure 6-5 Device Configuration Summary Report Example

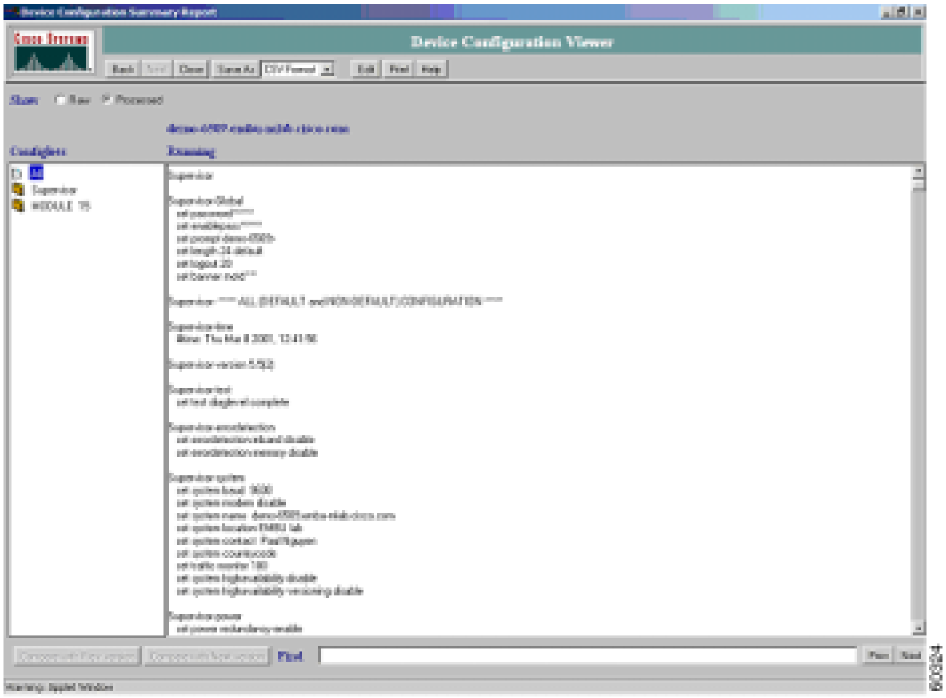

```
Supervisor

Supervisor-Global
  set password *****
  set enablepass *****
  set prompt demo-6509>
  set length 24 default
  set logout 20
  set banner motd ^C^C

Supervisor ****ALL (DEFAULT and NON-DEFAULT) CONFIGURATION****

Supervisor-time
  #time: Thu Mar 8 2001, 12:41:56

Supervisor-version 5.5(2)

Supervisor-test
  set test diaglevel complete

Supervisor-errordetection
  set errordetection inband disable
  set errordetection memory disable

Supervisor-system
  set system baud 9600
  set system modem disable
  set system name demo-6509.embu-mlab.cisco.com
  set system location EMBU lab
  set system contact Paul Nguyen
  set system countrycode
  set traffic monitor 100
  set system highavailability disable
  set system highavailability versioning disable

Supervisor-power
  set power redundancy enable
```

Software Management

Software upgrades to network devices can be a time-consuming and arduous task. Upgrading a router or LAN switch involves several related tasks before the device can be successfully loaded with a different software image. Reviewing the release notes from a vendor is the first step to ensure that all hardware requirements and software dependencies are taken into consideration.

The next step involves downloading the correct images from the vendor's web site or other distribution mechanism. A change management process needs to be in place to ensure that the upgrade has no impact to network services. Cisco has significantly simplified the process of upgrading software by having a management tool to automate the various tasks of the software upgrade. The Software Image Management tool (SWIM) in CiscoWorks2000 automates various tasks associated with software management.

Figure 6-6 shows an example of the software Image Library Summary window in CiscoWorks2000 Essentials.

Figure 6-6 Image Library Summary Window

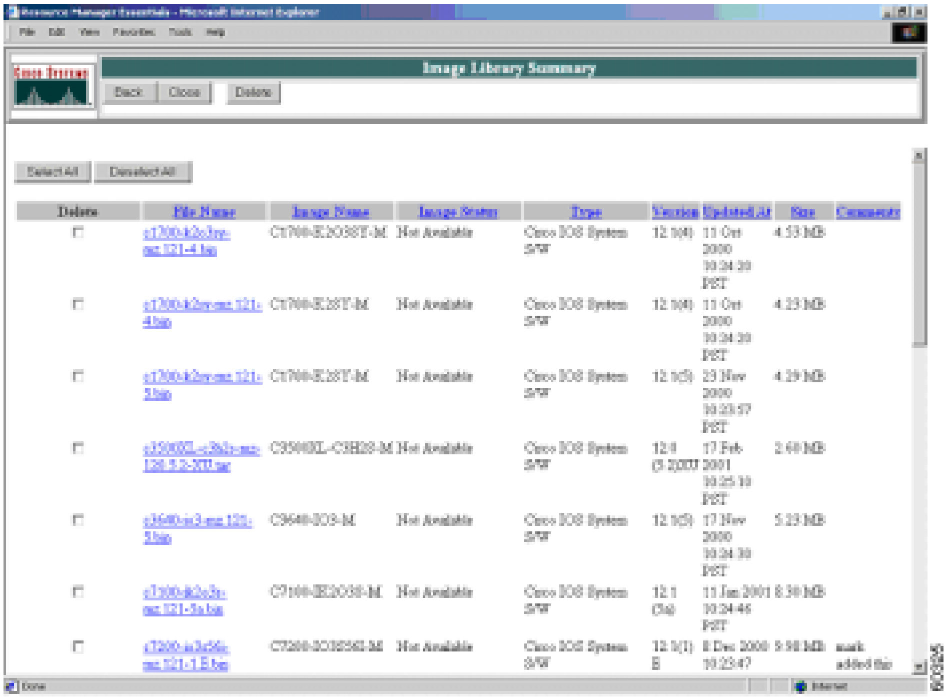
Resource Manager Essentials - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Cisco Systems
Image Library Summary
Back
Close
Delete
Select All
Deselect All
Done
Internet
60265

Delete	File Name	Image Name	Image Status	Type	Version	Updated At	Size	Comments
□	c1700-k2o3sy-mz.121-4.bin	C1700-K2O3SY-M	Not Available	Cisco IOS System S/W	12.1(4)	11 Oct 2000 10:24:20 PST	4.53 MB	
□	c1700-k2sy-mz.121-4.bin	C1700-K2SY-M	Not Available	Cisco IOS System S/W	12.1(4)	11 Oct 2000 10:24:20 PST	4.23 MB	
□	c1700-k2sy-mz.121-5.bin	C1700-K2SY-M	Not Available	Cisco IOS System S/W	12.1(5)	23 Nov 2000 10:23:57 PST	4.29 MB	
□	c3500XL-c3h2s-mz-120.5.2-XU.tar	C3500XL-C3H2S-M	Not Available	Cisco IOS System S/W	12.0 (5.2)XU	17 Feb 2001 10:25:10 PST	2.60 MB	
□	c3640-is3-mz.121-5.bin	C3640-IS3-M	Not Available	Cisco IOS System S/W	12.1(5)	17 Nov 2000 10:24:30 PST	5.23 MB	
□	c7100-ik2o3s-mz.121-5a.bin	C7100-IK2O3S-M	Not Available	Cisco IOS System S/W	12.1 (5a)	11 Jan 2001 10:24:46 PST	8.30 MB	
□	c7200-is3s56i-mz.121-1.E.bin	C7200-IO3S56I-M	Not Available	Cisco IOS System S/W	12.1(1) E	8 Dec 2000 10:23:47	9.98 MB	mark added this

Accounting Management

Accounting management has many meanings. In a traditional networking environment, accounting usually means tracking traffic flows based on IP source and destination addresses for the purpose of either capacity planning or departmental chargeback. In the voice world, accounting usually means call detail records (CDRs) that provide information on called numbers, call times, and so forth. For detailed information about call detail records, go to the following location: http://www.cisco.com/warp/public/788/AVVID/ts_ccm_301_sec7.htm.

The following subsections provide a brief summary of some of the accounting management tools.

Cisco NetFlow and IP Accounting

The first step toward appropriate accounting management is to measure the utilization of all important network resources. A usage-based accounting and billing system is an essential part of any service-level agreement (SLA), as it provides both a practical way of defining obligations under a SLA and clear consequences for behavior outside the terms of the SLA.

Cisco NetFlow and Cisco IP Accounting are tools for measuring network resource utilization. An analysis of the data gathered through the use of these tools provides insight into current usage patterns. The data can be collected by means of probes or by using Cisco NetFlow. The Cisco NetFlow FlowCollector with Network Data Analyzer application gathers and analyzes data from routers and Catalyst switches. Shareware applications such as cflowd can also be used to gather NetFlow data.

For more information about cflowd, go to the following location:
http://www.caida.org/tools/measurement/cflowd/.

An ongoing measurement of resource use can yield billing information, as well as information used to assess continued fair and optimal resources. Some commonly deployed accounting management solutions include:

- Cisco Network Data Analyzer
- Cisco NetFlow Collector
- Apogee Networks

  http://www.apogeenetworks.com
- XACCT Technologies, Inc.

  http://www.xacct.com
- Telemate

  http://www.telemate.net

NetFlow Activation and Data Collection Strategy

NetFlow (network flow) is an input side-measurement technology that allows for capturing the data required for network planning, monitoring, and accounting applications. NetFlow should be deployed on edge and aggregation router interfaces for service providers or WAN access router interfaces for enterprise customers.

Cisco Systems recommends careful planning of a NetFlow deployment with NetFlow services activated on these strategically located routers. NetFlow can be deployed incrementally (interface by interface) and strategically (on well chosen routers), rather than deploying NetFlow on every router on the network. Cisco personnel will work with customers to determine on which key routers and key interfaces NetFlow should be activated based on the customer’s traffic flow patterns, network topology, and architecture.

Key deployment considerations include:

- NetFlow services should be utilized as an edge metering and access list performance acceleration tool and should not be activated on “hot” core/backbone routers or routers running at very high CPU utilization rates.
- Understand application-driven data collection requirements. Accounting applications may only require originating and terminating router flow information whereas monitoring applications may require a more comprehensive (data intensive) end-to-end view.
- Understand the impact of network topology and routing policy on flow collection strategy. For example, avoid collecting duplicate flows by activating NetFlow on key aggregation routers where traffic originates or terminates and not on backbone routers or intermediate routers which would provide duplicate views of the same flow information.
- Service providers in the “transit carrier” business (meaning they carry traffic neither originating nor terminating on their network) may utilize NetFlow Export data for measuring transit traffic usage of network resources for accounting and billing purposes.

IP Accounting Configuration

Cisco IP accounting support provides basic IP accounting functions. By enabling IP accounting, users can see the number of bytes and packets switched through Cisco IOS software on a source and destination IP address basis. Only transit IP traffic is measured and only on an outbound basis. Traffic generated by the software or terminating in the software is not included in the accounting statistics. To maintain accurate accounting totals, the software maintains two accounting databases: an active and a check-pointed database.

Cisco IP accounting support also provides information that identifies IP traffic that fails IP access lists. Identifying IP source addresses that violate IP access lists signals possible attempts to breach security. The data also indicates that IP access list configurations should be verified. To make this feature available to users, enable IP accounting of access list violations using the **ip accounting access-violations** command. Users can then display the number of bytes and packets from a single source that attempted to breach security against the access list for the source destination pair. By default, IP accounting displays the number of packets that have passed access lists and were routed.

To enable IP accounting, use one of the following commands for each interface in interface configuration mode:

Command	Purpose
ip accounting	Enable basic IP accounting
ip accounting access-violations	Enable IP accounting with the ability to identify IP traffic that fails access lists.

To configure other IP accounting functions, use one or more of the following commands in global configuration mode:

Command	Purpose
ip accounting-threshold t*hreshold*	Set the maximum number of accounting entries to be created.
ip accounting-list *ip-address wildcard*	Filter accounting information for hosts.
ip accounting-transits *count*	Control the number of transit records that will be stored in the IP accounting database.

Performance Management

The performance of routers and LAN switches can be analyzed using SNMP management protocols supported on the devices. Various performance metrics are defined in the standard protocols, as well as the Cisco-specific management information base (MIB). A full explanation of performance management is beyond the scope of this document. Minimally, these performance indicators from routers and switches should be collected at the device level, interface level, and protocol level. Device level performance metrics for routers include:

- CPU utilization
- Memory utilization
- System buffers

Various solutions are available in the marketplace to address the needs of performance management for the IP Telephony environment. These systems are capable of collecting, storing, and presenting data from network devices and servers. Some of the commonly deployed performance management solutions include:

- Concord Network Health

  http://www.concord.com
- InfoVista VistaView

  http://www.infovista.com

The SNMP platform can also poll and store the specific MIB Object Identifiers (OIDs) for detailed analysis and threshold monitoring. It is highly recommended that customers work with Cisco to identify the correct performance OIDs for their particular environment. Thresholds should also be reviewed with Cisco since the defaults in many current NMS products are only a starting point and may not be correct for every situation.

- For more detailed information about performance management, and other management topics, go to the following location: http://www.cisco.com/warp/public/126/index.shtml.
- For information about network monitoring and event correlation guidelines, go to the following location: http://www.cisco.com/warp/public/cc/pd/wr2k/tech/cnm_rg.pdf.
- To purchase a reference book on managing Cisco device performance, go to the following location: http://www1.fatbrain.com/asp/bookinfo/bookinfo.asp?theisbn=1578701805&vm=.

Voice Quality Metrics

VoIP is susceptible to network behaviors that can degrade the voice application to the point of being unacceptable to the average user. These network behaviors include:

- Delay—The amount of time taken from point to point in a network. Delay can be measured either in one direction or as a round trip. One-way delay calculations require expensive, sophisticated test gear and are beyond the budget and expertise of most enterprise customers. However, measuring round-trip delay is easier and requires less expensive equipment.

  A sufficient method of measuring a one-way delay is to measure the round-trip delay and divide the result in half. VoIP can typically tolerate delays up to approximately 150ms before the quality of the call is unacceptable.
- Jitter—The variation in delay over time from point to point. If the delay of transmissions varies too widely in a VoIP call, the call quality is greatly degraded. The amount of jitter tolerable on the network is affected by the depth of the jitter buffer on the network equipment in the voice path. The more jitter buffer available, the more the network, in general, can reduce the effects of jitter.
- Packet loss—The loss of packets along the data path, which severely degrades the voice application.

Prior to deploying VoIP applications, it is very important to assess the delay, jitter and packet loss on the data network in order to determine if the voice applications will work. The jitter, delay, and packet loss measurements can then aid in the correct design and configuration of traffic prioritization and buffering parameters in the data network equipment.

Service Assurance Agent and Round Trip Time Monitor

Service Assurance Agent and the Round Trip Time Monitor (RTTMON) MIB are Cisco IOS features that allow for the testing and collection of delay, jitter, and packet loss statistics on the data network.

Internet Performance Monitor (IPM) is a Cisco network management application that can configure the Cisco IOS features and monitor the RTR and RTTMON data. The RTR and RTTMON Cisco IOS features can be used to measure delay, jitter, and packet loss by deploying small Cisco IOS routers as agents to simulate customer end stations or IP phones. The routers are referred to as delay/jitter probes. Additionally, the delay/jitter probes can be configured with RMON Alarm and Event triggers to monitor the network after the baseline values have been determined. This allows the delay/jitter probes to monitor the network for predetermined delay and jitter service levels and alert NMS stations when a threshold is exceeded.

Jitter Calculations in Cisco SAA Delay Jitter Probes

This section shows how Cisco IOS calculates jitter, delay, and packet loss in Cisco SAA delay jitter probes. It is calculated based on sending and receiving time stamps of consecutive packets sent out.

Sender	Responder
T1 send packet1	
T2	receive packet1
T3	send back reply for packet1
T4 receive reply for packet1	
T5 send packet2	
T6	receive packet2
T7	send back reply for packet2
T8 receive reply for packet2	

For the two packets in the example above:

- Jitter from source to destination (JitterSD) = (T6-T2) - (T5-T1)
- Jitter from destination to source (JitterDS) = (T8-T4) - (T7-T3)

Jitter is calculated using time stamps of every two consecutive packets. For example:

Router1 sends packet1	T1 = 0
Router2 receives packet1	T2 = 20ms
Router2 sends back packet1	T3 = 40ms
Router1 receives packet1 response	T4 = 60ms
Router1 sends packet2	T5 = 60ms
Router2 receives packet2	T6 - 82ms

Router2 sends back packet2	T7 = 104ms
Router1 receives packet2 response	T8 = 126ms

Jitter from source to destination (JitterSD) = (T6-T2) - (T5-T1)

Jitter from source to destination (JitterSD) = (82ms - 20ms) - (60ms - 0ms) = 2ms positive jitter SD

Jitter from destination to source (JitterDS) = (T8-T4) - (T7-T3)

Jitter from destination to source (JitterDS) = (126ms - 60ms) - (104ms - 40ms) = 2ms positive jitter DS

Deploying Echo/Jitter Agent Routers

Delay and jitter can be measured by deploying Cisco 17xx, or larger, routers with Cisco IOS version 12.05T or later and configuring the Cisco IOS Service Assurance Agent features. The routers should be placed in the campus networks alongside hosts to provide statistics for end-to-end connections. It is not practical to measure every possible voice path in the network. Therefore, the probes should be placed in typical host locations to provide a statistical sampling of typical voice paths. Some examples include:

- A local campus-to-campus path
- A local campus to a remote campus path via 384kbs Frame Relay circuit
- A local campus to remote campus via asynchronous transfer mode private virtual circuit (ATM PVC)

In the case of a VoIP deployment using traditional phones connected to Cisco routers using FXS station ports, the router to which the phones are connected to should serve as the delay/jitter probes. Once deployed, the probe collects statistics and populates SNMP MIB tables in the router. The data can then be accessed either through the Cisco IPM application, through command line, or through SNMP polling tools. Additionally, after baseline values have been established, RTR can be configured to send alerts to an NMS station if thresholds for delay, jitter, and packet loss are exceeded.

Simulating a Voice Call

One of the strengths of using SAA as the testing mechanism is that a voice call can be simulated. For example, a typical G.711 voice call adheres to the following:

- It uses RTP/UDP ports 14384 and above
- It is approximately 64kb/s

G.711 voice traffic can be simulated by setting up the SAA delay/jitter probe as described below.

The jitter operation would need to adhere to the following in order to provide 64kb/s for 20 seconds:

- Send the request to RTP/UDP port number 14384
- Send 492 byte packets (480 payload *and* 12 byte RTP header size) and 28 bytes (IP and UDP)
- Send 1000 packets for each frequency cycle
- Send every packet 20 milliseconds apart for a duration of 20 seconds and sleep for 40 seconds before starting the next frequency cycle

```
((1000 datagrams * 480 bytes per datagram)/ 60 seconds))  *  8 bits per byte =  64kb/s
```

The configuration of the router would then appear as follows:

```
rtr 1
type jitter dest-ipaddr 172.18.179.10 dest-port 14384 num-packets 1000
request-data-size 492
```

```
frequency 60
rtr schedule 1 life 2147483647 start-time now
```

Note IP and UDP are not considered in the request-data-size: The router adds them automatically to the size internally.

Figure 6-7 shows routers set up to simulate voice calls of a 20-second duration every 60 seconds and record delay, packet loss, and jitter in both directions. The delay calculations in this example are round-trip times and must be divided by two to get the one-way delay.

Figure 6-7 Delay/Jitter Probe Deployment Example

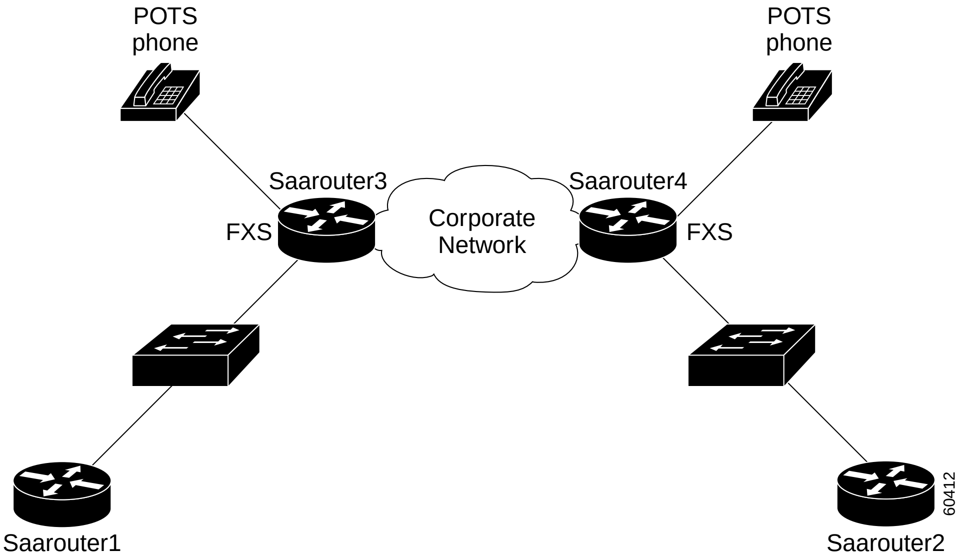

Delay/Jitter Probe Deployment Example

```
saarouter1#
rtr responder
rtr 1
type jitter dest-ipaddr 172.18.179.10 dest-port 14384 num-packets 1000
request-data-size 492
frequency 60
rtr schedule 1 life 2147483647 start-time now

saarouter2#
rtr 1
type jitter dest-ipaddr 172.18.178.10 dest-port 14385 num-packets 1000
request-data-size 492
rtr schedule 1 life 2147483647 start-time now

saarouter3#
rtr 1
```

```
type jitter dest-ipaddr 172.18.179.100 dest-port 14385 num-packets 1000
request-data-size 492
frequency 60
rtr schedule 1 life 2147483647 start-time now

saarouter4#
rtr 1
type jitter dest-ipaddr 172.18.178.100 dest-port 14385 num-packets 1000
request-data-size 492
frequency 60
rtr schedule 1 life 2147483647 start-time now
```

Sample Data Collections

Polling the MIB Tables

The delay/jitter probes begin collecting data that is subsequently placed in SNMP MIB tables. There are two tables of particular interest:

- rttMonStats—Provides a one-hour average of all the jitter operations for the last hour.
- rttMonLatestJitterOper—Provides the values of the last operation completed.

For general statistics on delay and jitter, the rttMonStats table should be polled every hour. For more granular statistics, the rttMonLatestJitterOper table should be polled at an interval larger than the jitter operation frequency which is 60 seconds in this example. That is, if the delay/jitter probe is calculating jitter every five minutes, you would not poll the MIB at any interval less than five minutes.

Figure 6-8 shows data from the rttMonJitterStatsTable gathered from an HP OpenView (HPOV) Network Node Manager MIB poll.

Figure 6-8 HPOV Network Node Manager MIB Poll

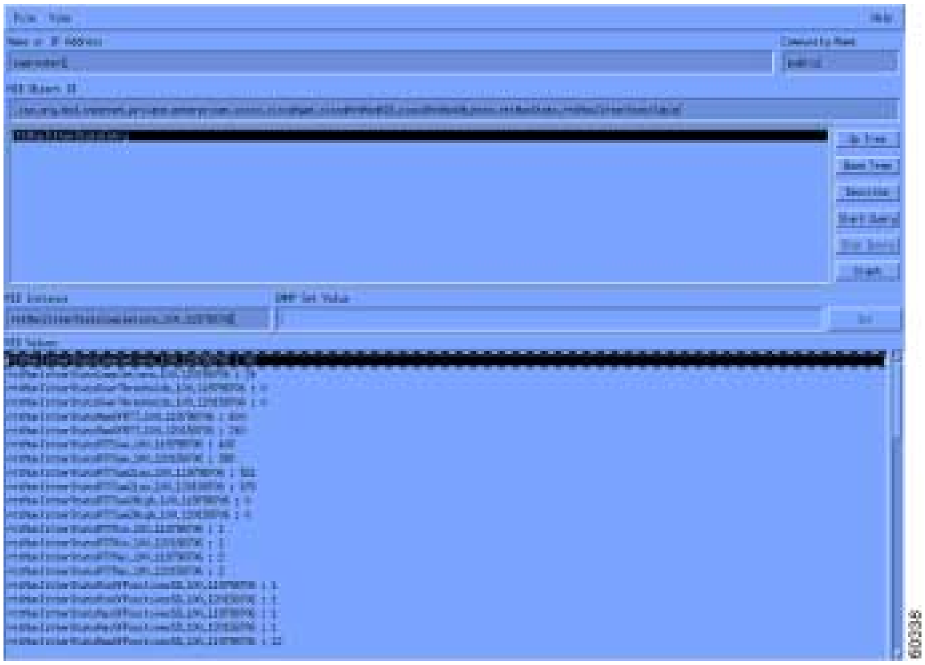

Example Report

Figure 6-9 shows an example of SAA data that was collected via SNMP using HPOV Network Node Manager and exported into a Microsoft Excel spreadsheet. The graph is a compilation of delay, jitter, and packet loss data points over an eight-hour period for one pair of delay/jitter probes.

Figure 6-9 Delay, Jitter, and Packet Loss Graph

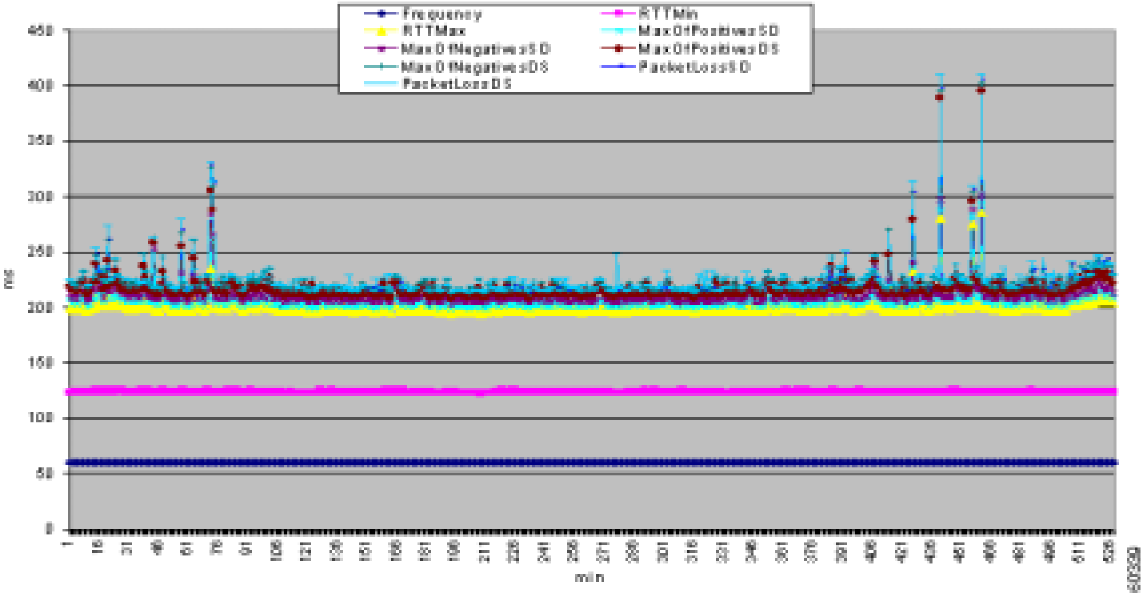

The Internetwork Performance Monitor tool in CiscoWorks2000 simplifies the monitoring and reporting of jitter, delay, and packet loss. The IPM application automatically sets up the agent routers and tracks data for the reports.

Figure 6-10 shows an example of a Cisco IPM report.

Figure 6-10 Cisco IPM Historical Statistics Report

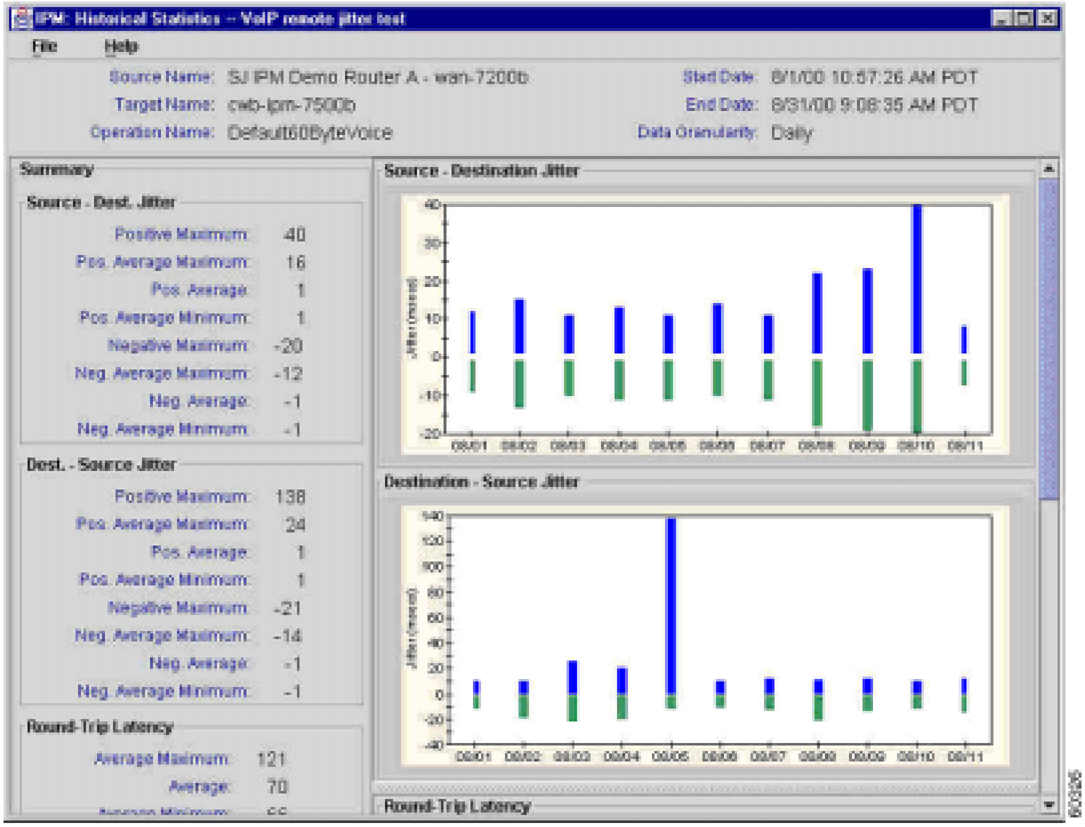

IPM: Historical Statistics -- VoIP remote jitter test

Source Name: SJ IPM Demo Router A - wan-7200b
Target Name: cwb-ipm-7500b
Operation Name: Default60ByteVoice
Start Date: 8/1/00 10:57:26 AM PDT
End Date: 8/31/00 9:08:35 AM PDT
Data Granularity: Daily

Summary

Source - Dest. Jitter

Positive Maximum:	40
Pos. Average Maximum:	16
Pos. Average:	1
Pos. Average Minimum:	1
Negative Maximum:	-20
Neg. Average Maximum:	-12
Neg. Average:	-1
Neg. Average Minimum:	-1

Dest. - Source Jitter

Positive Maximum:	138
Pos. Average Maximum:	24
Pos. Average:	1
Pos. Average Minimum:	1
Negative Maximum:	-21
Neg. Average Maximum:	-14
Neg. Average:	-1
Neg. Average Minimum:	-1

Round-Trip Latency

Average Maximum:	121
Average:	70
Average Minimum:	66

While the above mentioned methods of monitoring jitter, delay, and packet loss are not an exhaustive test of network behavior, they are an acceptable indicator of network characteristics and a predictor of how time-sensitive applications such as VoIP will perform on your network.

Security Management at the Device Level

The goal of security management is to control access to network resources according to local guidelines so that the network cannot be sabotaged (intentionally or unintentionally). A security management subsystem, for example, can monitor users logging on to a network resource, refusing access to those who enter inappropriate access codes.

Security management is a very broad subject; therefore this area of the document only covers security as related to SNMP and basic device access security. Detailed information on advanced security can be found at the following locations:

- http://www.cisco.com/univercd/cc/td/doc/cisintwk/ics/cs003.htm
- http://www.cisco.com/warp/public/cc/pd/wr2k/vpmnso/

A good security management implementation starts with sound security policies and procedures in place. It is important to create a platform-specific minimum configuration standard for all routers and switches that follow industry best practices for security and performance.

There are various methods of controlling access on Cisco routers and switches. Some of these methods include:

- Access Control List (ACL)
- User IDs and passwords local to the device
- Terminal Access Controller Access Control System (TACACS)

TACACS is an Internet Engineering Task Force (RFC 1492) standard security protocol that runs between client devices on a network and against a TACACS server. TACACS is an authentication mechanism that is used to authenticate the identity of a device seeking remote access to a privileged database. TACACS was enhanced by Cisco into TACACS+ (also referred to as CiscoSecure), the AAA architecture that separates authentication, authorization, and accounting functions.

TACACS+ is used by Cisco to allow a finer control over who can access the Cisco device in non-privileged and privileged mode. Multiple TACACS+ servers can be configured for fault tolerance. With TACACS+ enabled, the router prompts the user for a user name and a password. With TACACS+ enabled, the router and switch prompts the user for a user name and password. Authentication can be configured for login control or to authenticate individual commands.

Authentication

Authentication is the process of identifying users, including:

- Login and password dialog
- Challenge and response
- Messaging support

Authentication is the way a user is identified prior to being allowed access to the router or switch. There is a fundamental relationship between authentication and authorization. The more authorization privileges a user receives, the stronger the authentication should be.

Authorization

Authorization provides remote access control, including one-time authorization and authorization for each service that is requested by the user. On a Cisco router, the authorization level range for users is 0 to 15 with 0 being the lowest level and 15 the highest.

Accounting

Accounting allows for the collecting and sending of security information used for billing, auditing, and reporting, such as user identities, start and stop times, and executed commands. Accounting enables network managers to track the services that users are accessing as well as the amount of network resources they are consuming.

Table 6-19 lists basic sample commands for using TACACS+, authentication, authorization, and accounting on a Cisco router and a Catalyst switch. For more in-depth commands, go to the following location: http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/12cgcr/secur_r/srprt1/index.htm.

For more information on how to configure AAA to monitor and control access to the command-line interface on the Catalyst enterprise LAN switches, go to the following location: http://www.cisco.com/univercd/cc/td/doc/product/lan/cat5000/rel_5_4/config/authent.htm.

Table 6-19 Cisco IOS Commands

Cisco IOS Command	Purpose
Router	
aaa new-model	Enable authentication, authorization, accounting (AAA) as the primary method for access control.
aaa accounting *{system | network | connection | exec | command level} {start-stop | wait-start | stop-only} {tacacs+ | radius}*	Enable accounting with the global configuration commands.
aaa authentication login default tacacs+	Set up the router so that connections to any terminal line configured with the login default will be authenticated with TACACS+, and will fail if authentication fails for any reason.
aaa authorization exec default tacacs+ none	Set up the router to check if the user is allowed to run an EXEC shell by asking the TACACS+ server.
tacacs-server host *tacacs+ server ip address*	Specify the TACACS+ server that will be used for authentication with the global configuration commands.
tacacs-server key *shared-secret*	Specify the shared secret that is known by the TACACS+ servers *and* the Cisco router with the global configuration command.
Catalyst Switch	
set authentication login tacacs enable *[all | console | http | telnet] [primary]*	Enable TACACS+ authentication for normal login mode. Use the console or Telnet keywords to enable TACACS+ only for console port or Telnet connection attempts.
set authorization exec enable *{option} fallback option} [console | telnet | both]*	Enable authorization for normal login mode. Use the console or Telnet keywords to enable authorization only for console port or Telnet connection attempts.
set tacacs-server key *shared-secret*	Specify the shared secret that is known by the TACACS+ servers *and* switch.
set tacacs-server host *tacacs+ server ip address*	Specify the TACACS+ server that will be used for authentication with the global configuration commands.
set accounting commands enable *{config | all} {stop-only} tacacs+*	Enable accounting of configuration commands.

SNMP Security

The SNMP protocol can be used to make configuration changes on routers and Catalyst switches similar to those issued from the command-line interface (CLI). Proper security measures should be configured on network devices to prevent unauthorized access and change via SNMP. Community strings should follow the standard password guidelines for length, characters, and difficulty of guessing. It is important to change the community strings from their public and private defaults.

All SNMP management host(s) should have a static IP address and be explicitly granted SNMP communication rights with the network device by that predefined by IP address and ACL. Cisco IOS and Cisco Catalyst software provides security features that ensure that only authorized management stations are allowed to perform changes on network devices.

Router Security Features

- SNMP Privilege Level—Limits the types of operations that a management station can have on a router. There are two types of privilege level on routers: Read-Only (RO) and Read-Write (RW). The RO level only allows a management station to query the router data. It does not allow for configuration commands such as rebooting a router and shutting down interfaces to be performed. Only the RW privilege level can be used to perform such operations.
- SNMP ACL—Can be used in conjunction with the SNMP privilege feature to limit specific management stations from requesting management information from routers.
- SNMP View—Limits specific information that can be retrieved from routers by management stations. It can be used with the SNMP privilege level and ACL features to enforce restricted access of data by management consoles. For configuration samples of SNMP View, go to the following location: http://www.cisco.com/univercd/cc/td/doc/product/software/ios111/mods/1mod/1rbook/1rsysmgt.htm#xtocid27380181.
- SNMP Version 3—SNMPv3 provides secure exchanges of management data between network devices and management stations. The encryption and authentication features in SNMPv3 ensure high security in transporting packets to a management console. SNMPv3 is supported in Cisco IOS version 12.0(3)T and later. For a technical overview of SNMPv3, go to the following location: http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/120newft/120t/120t3/snmp3.htm.
- ACL on Interfaces—Provides security measures in preventing attacks such as IP spoofing. The ACL can be applied on incoming or outgoing interfaces on routers.

Catalyst LAN Switch Security Features

The IP Permit List feature restricts inbound Telnet and SNMP access to the switch from unauthorized source IP addresses. Syslog messages and SNMP traps are supported to notify a management system when a violation or unauthorized access occurs.

Use the **set ip permit** command set to enable or disable the IP permit list and to specify IP addresses to be added to the IP permit list. The IP permit list is disabled by default.

```
set ip permit {enable | disable}
set ip permit {enable | disable} [telnet | ssh | snmp]
set ip permit addr [mask] [telnet | ssh | snmp | all]
```

The syntax description for the **set ip permit** command is as follows:

enable—Keyword to enable the IP permit list

disable—Keyword to disable the IP permit list

telnet—(optional) Keyword to specify the Telnet IP permit list

ssh—(optional) Keyword to specify the SSH IP permit list

snmp—(optional) Keyword to specify the SNMP IP permit list

addr—IP address to be added to the IP permit list (an IP alias or host name that can be resolved through DNS can also be used)

mask—(optional) Subnet mask of the specified IP address

all—(optional) Keyword to specify all entries in the IP permit list be removed

Managing The Application Layer with CallManager

The CallManager application in the IP Telephony architecture performs a number of call processing functions for the IP phones. The media convergence server (MCS) hardware hosts the CallManager software and is fully manageable from a web interface and SNMP. Recent enhancements to the software includes support for syslog, CallManager MIBs, and other Cisco-specific MIBs. CallManager management can be fully integrated into existing network management systems using the management capabilities on the system.

To ease the administration of network and application management, CiscoWorks2000 RME and CallManager 3.x can be integrated at the web interface level as shown in Figure 6-11.

Note Details on configuring RME and CallManager 3.x are available on the web CiscoWorks 2000 documentation on your CiscoWorks 2000 RME server.

Figure 6-11 CiscoWorks2000 RME and CallManager Integration

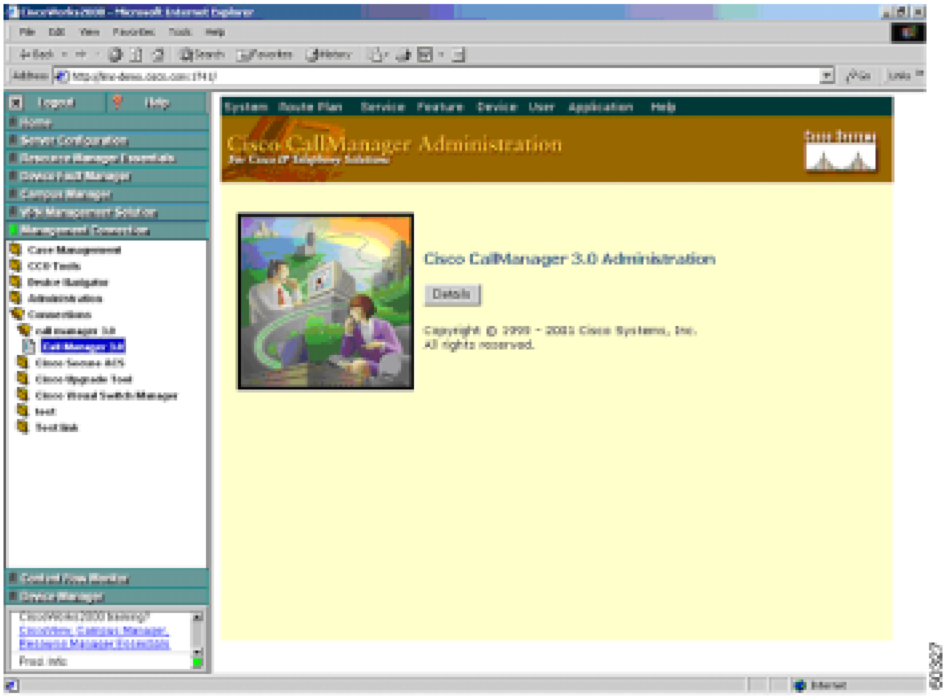
CiscoWorks2000 - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Back Search Favorites History
Address
Go Links
Logout Help
Home
Server Configuration
Resource Manager Essentials
Device Fault Manager
Campus Manager
VPN Management Solution
Management Connection
Case Management
CCO Tools
Device Navigator
Administration
Connections
call manager 3.0
CallManager 3.0
Cisco Secure ACS
Cisco Upgrade Tool
Cisco Virtual Switch Manager
test
Test link
Content Flow Monitor
Device Manager
CiscoWorks2000 training?
CiscoWorks Campus Manager
Resource Manager Essentials
Prod. info
System Route Plan Service Feature Device User Application Help
Cisco CallManager Administration
For Cisco IP Telephony Solutions
Cisco Systems
Cisco CallManager 3.0 Administration
Details
Copyright © 1999 - 2001 Cisco Systems, Inc.
All rights reserved.
Internet
60327

Inventory Management

Inventory details of CallManager can be obtained from the software using network management protocols (SNMP). CallManager 3.x software supports Cisco Discovery Protocol (CDP) and can be auto-discovered by CiscoWorks2000 LAN Management Solution (LMS) software. The most effective NMS implementations have a complete inventory of CallManagers, along with other elements in the infrastructure.

Troubleshooting efforts and upgrade processes require detailed and up-to-date inventory information to assist network administrators in understanding the following:

- Software dependencies

- Feature requirements
- Hardware requirements

Web-based inventory tools can significantly improve ease of access to device information from anywhere in the organization. CallManager provides excellent inventory and element management capabilities. Figure 6-12 shows an example of how CallManager software displays an IP phone inventory list. In addition to the description, CallManager displays the device pool membership and device name of each IP phone.

Figure 6-12 CallManager Inventory/Element Management Report

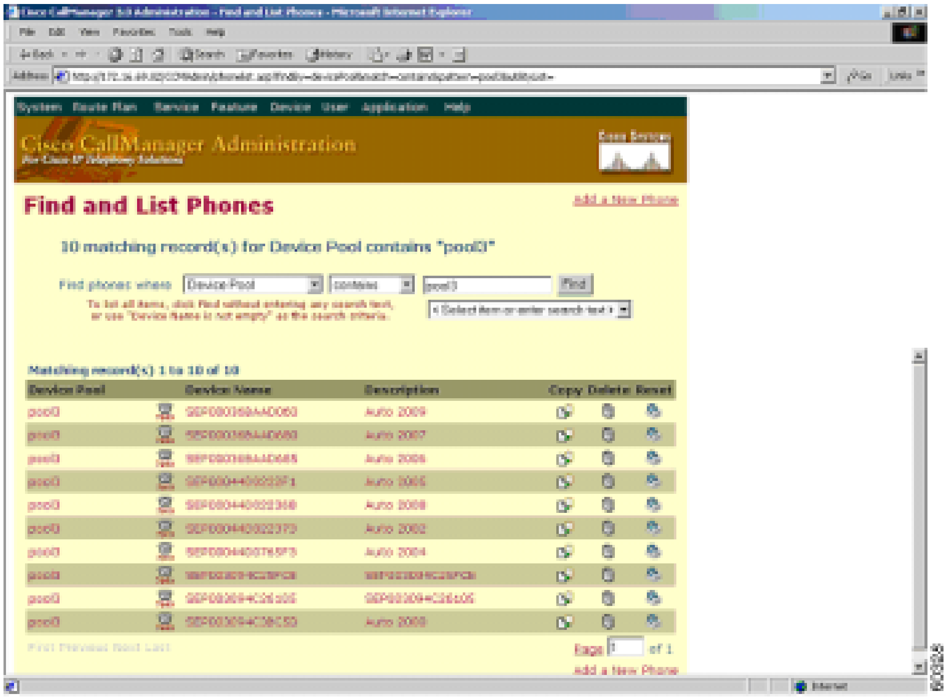
System Route Plan Service Feature Device User Application Help
Cisco CallManager Administration
For Cisco IP Telephony Solutions
Cisco Systems
Add a New Phone
Find phones where
Device Pool
contains
pool0
Find
< Select item or enter search text >

Find and List Phones

10 matching record(s) for Device Pool contains "pool0"

To list all items, click Find without entering any search text, or use "Device Name is not empty" as the search criteria.

Matching record(s) 1 to 10 of 10

Device Pool	Device Name	Description	Copy	Delete	Reset
pool0	SEP00036BAAD060	Auto 2009			
pool0	SEP00036BAAD680	Auto 2007			
pool0	SEP00036BAAD685	Auto 2006			
pool0	SEP0004400222F1	Auto 2005			
pool0	SEP000440022368	Auto 2008			
pool0	SEP000440022373	Auto 2002			
pool0	SEP00044007E5F3	Auto 2004			
pool0	[illegible]	[illegible]			
pool0	SEP003094C26105	SEP003094C26105			
pool0	SEP003094C38C50	Auto 2000			

You can also manage CallManager from the CallManager application. Figure 6-13 shows the CallManager Administration screen from which you can add, change, delete, and restart CallManagers.

Figure 6-13 Cisco CallManager Configuration Window

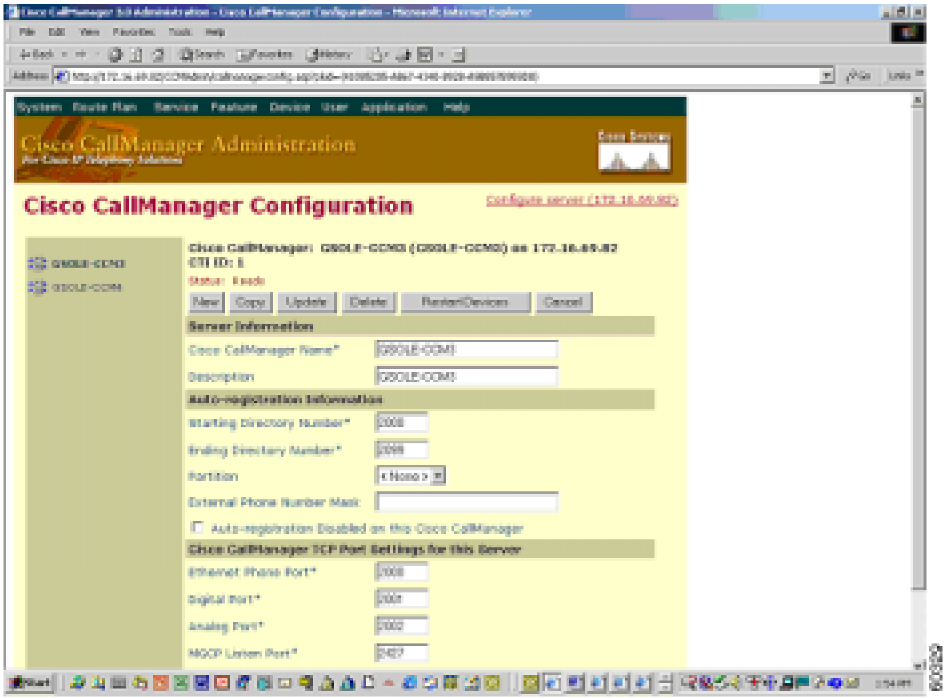

Cisco CallManager Configuration

Cisco CallManager: GSOLE-CCM3 (GSOLE-CCM3) on 172.16.69.82
CTI ID: 1
Status: Ready

Server Information

Cisco CallManager Name*: GSOLE-CCM3
Description: GSOLE-CCM3

Auto-registration Information

Starting Directory Number*: 2000
Ending Directory Number*: 2099
Partition: < None >
External Phone Number Mask:
☐ Auto-registration Disabled on this Cisco CallManager

Cisco CallManager TCP Port Settings for this Server

Ethernet Phone Port*: 2000
Digital Port*: 2001
Analog Port*: 2002
MGCP Listen Port*: 2427

Provisioning and Bulk Administration

Provisioning an IP phone involves adding user-specific configuration parameters in Cisco CallManager. Adding a large number of users in CallManager can be accomplished using the Bulk Administration Tool (BAT). Pre-defined user attributes are defined in a file and loaded into the software, which reduces the overhead associated with manually adding individual users. The tool can also be used to simultaneously delete a number of users from the database.

Figure 6-14 shows the Cisco CallManager BAT window.

Figure 6-14 Cisco CallManager Bulk Administration Tool Window

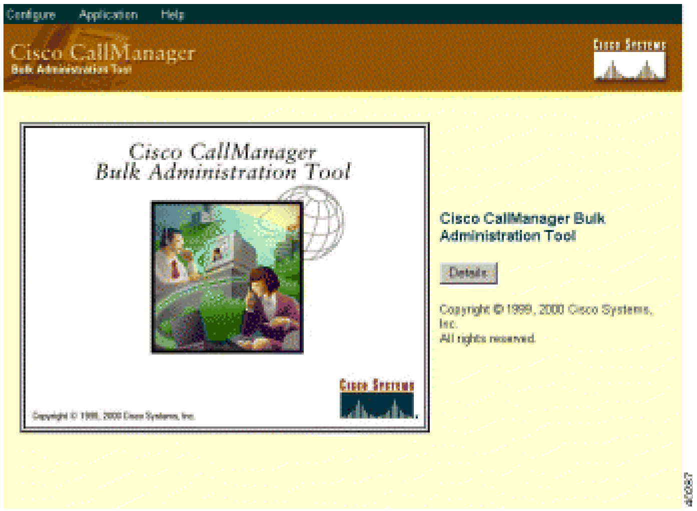

Copyright © 1999, 2000 Cisco Systems, Inc.
All rights reserved.

Copyright © 1999, 2000 Cisco Systems, Inc.

For detailed information about BAT, go to the following location: http://www.cisco.com/univercd/cc/td/doc/product/voice/sw_ap_to/admin/index.htm.

Gateway and Gatekeeper Management

The CallManager software uses a voice gateway when communication needs to be established with a user in the PSTN. Gateways registered with CallManager provide the signaling and call management functions for calls to a non-IP environment. Several types of gateways are available to meet the number of site-specific requirements in deploying IP Telephony. These gateways range from standalone systems to several models that can be integrated into the Catalyst LAN switches. The operational status of gateways and trunks registered with a CallManager can be obtained from a network management system containing Cisco CallManager MIB.

Standard bodies and vendors are currently defining MIBs for additional management of VoIP gateways. Manageability of VoIP gateways will be significantly improved once these MIBs are finalized and adopted by the industry. Refer to Table 6-21 for specific SNMP MIBs for gateway and trunk status.

The current version of CallManager supports three different types of gateway protocols:

- Skinny Station Protocol
- Media Gateway Control Protocol (MGCP)
- H.323

Cisco offers VoIP gateways for connection to the IP Telephony network to the PSTN. For detailed information about VoIP gateways, go to the following locations:

- http://www.cisco.com/warp/public/cc/pd/ga/prodlit/gatwy_wp.htm
- http://www.cisco.com/univercd/cc/td/doc/product/voice/c_callmg/3_0/admin_gd/3_0_5/p6gatewy.htm

Figure 6-15 shows the Cisco CallManager Gateway Administration window.

Figure 6-15 Cisco CallManager Gateway Administration Window

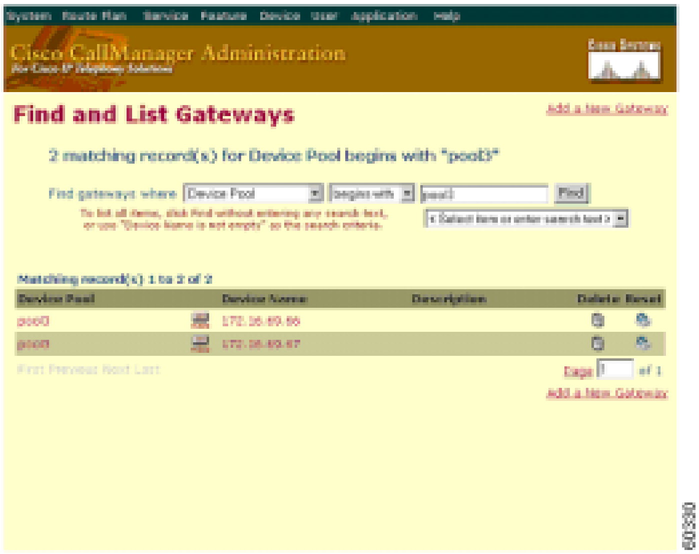
System Route Plan Service Feature Device User Application Help
Cisco CallManager Administration
For Cisco IP Telephony Solutions
Cisco Systems
Add a New Gateway
Find
First Previous Next Last
Add a New Gateway
60330

Find and List Gateways

2 matching record(s) for Device Pool begins with "pool3"

Find gateways where Device Pool begins with pool3

To list all items, click Find without entering any search text, or use "Device Name is not empty" as the search criteria.

< Select item or enter search text >

Matching record(s) 1 to 2 of 2

Device Pool	Device Name	Description	Delete	Reset
pool3	172.16.69.66			
pool3	172.16.69.67			

Page 1 of 1

A gatekeeper, also known as a Cisco Multimedia Conference Manager (MCM), is a device that supports the H.225 RAS message set used for the following:

- Call admission control
- Bandwidth allocation
- Dial pattern resolution

Only one gatekeeper device can be configured per Cisco CallManager cluster. Figure 6-16 shows the Cisco CallManager Gatekeeper Configuration window.

Figure 6-16 Cisco CallManager Gatekeeper Configuration Window

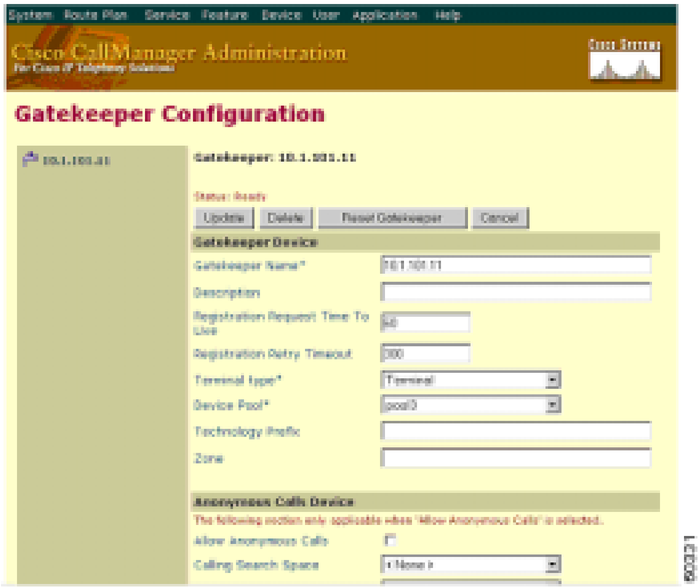
System Route Plan Service Feature Device User Application Help
Cisco CallManager Administration
For Cisco IP Telephony Solutions
Cisco Systems
10.1.101.11
Update
Delete
Reset Gatekeeper
Cancel
60331

Gatekeeper Configuration

Gatekeeper: 10.1.101.11

Status: Ready

Gatekeeper Device

Gatekeeper Name*	10.1.101.11
Description	
Registration Request Time To Live	60
Registration Retry Timeout	300
Terminal type*	Terminal
Device Pool*	pool3
Technology Prefix	
Zone	

Anonymous Calls Device

The following section only applicable when 'Allow Anonymous Calls' is selected.

Allow Anonymous Calls	☐
Calling Search Space	< None >

Voice Port Management

CiscoWorks2000 Voice Manager 2.0 (CVM) is a web-based voice management and reporting solution. The application provides enhanced capabilities to configure and provision voice ports, and to create and modify dial plans on voice-enabled Cisco routers for VoIP, Voice over Frame Relay (VoFR) and Voice over ATM (VoATM) network deployments.

For information about the features in CVM 2.0, go to the following location:
http://www.cisco.com/warp/public/cc/pd/wr2k/cw2kvm/prodlit/cvm2_ds.htm.

Fault and Performance Management

The ability to detect faults in the CallManager is essential to ensure minimal disruption of services to users. Database and call processing redundancy features, combined with distributed call processing capabilities in the software, offers a high level of system availability. The operational status of CallManager software can be obtained by integrating CallManager MIBs into a network management system. MIBs can be polled in the SNMP platform and compared to the recommended thresholds. When those thresholds are exceeded, the SNMP platform can send an alert message to the appropriate personnel.

Another method of managing CallManager faults is by using Cisco Voice Health Monitor (VHM). Voice Health Monitor is a CiscoWorks2000 application that monitors the health of a Cisco environment by proactively monitoring voice elements to minimize network downtime. The VHM application demonstrates converged system reliability by synthetically testing the network, providing real-time status reports on Cisco CallManagers, switches, and router gateways, and assisting with quick troubleshooting.

Some of the features of Cisco VHM include:

- Health Dashboard—Provides any fault indicators in the network
- Status indicators for health of system, environment, applications for media convergence servers
- Detailed device view—Specific information about a single network element
- System Information—A list of applications running on the PC
- Gateways registration and status
- Fault event browser
- VoIP-specific fault alarm view
- Network-wide fault alarm view

The graphical user interface-based Voice Health Monitor application allow for ease of viewing the various faults and VoIP components.

Figure 6-17 shows an example of the fault overview window. Network operations center personnel can use this window to monitor the overall voice health of the network in real-time.

Figure 6-17 Voice Health Monitor Fault Window

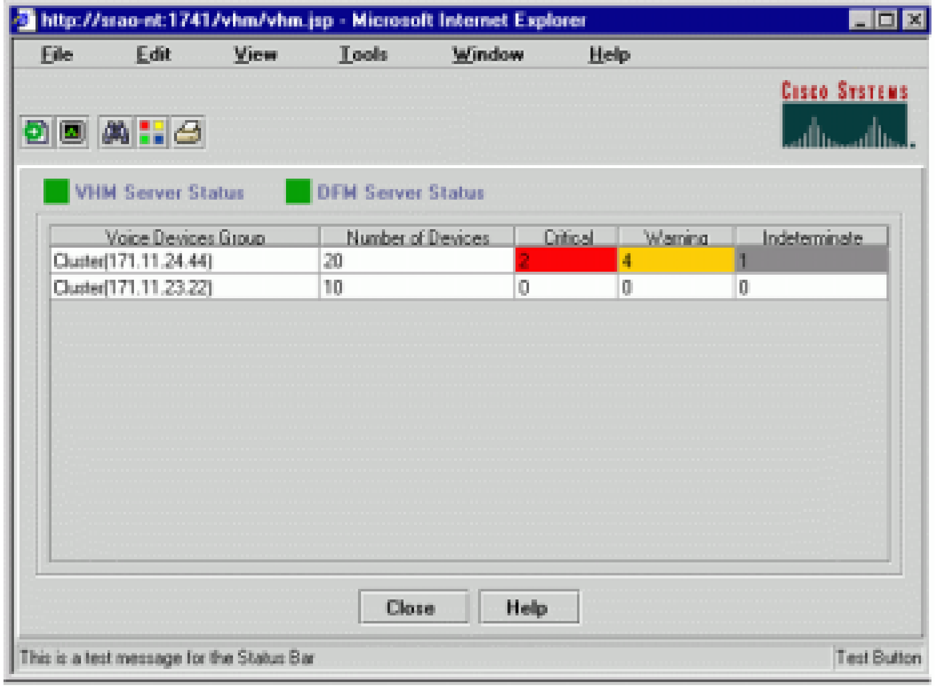

http://srao-nt:1741/vhm/vhm.jsp - Microsoft Internet Explorer

Voice Devices Group	Number of Devices	Critical	Warning	Indeterminate
Cluster(171.11.24.44)	20	2	4	1
Cluster(171.11.23.22)	10	0	0	0

This is a test message for the Status Bar

Figure 6-18 shows an example of a real-time status view of the VoIP components.

Figure 6-18 Voice Health Monitor Status View Window

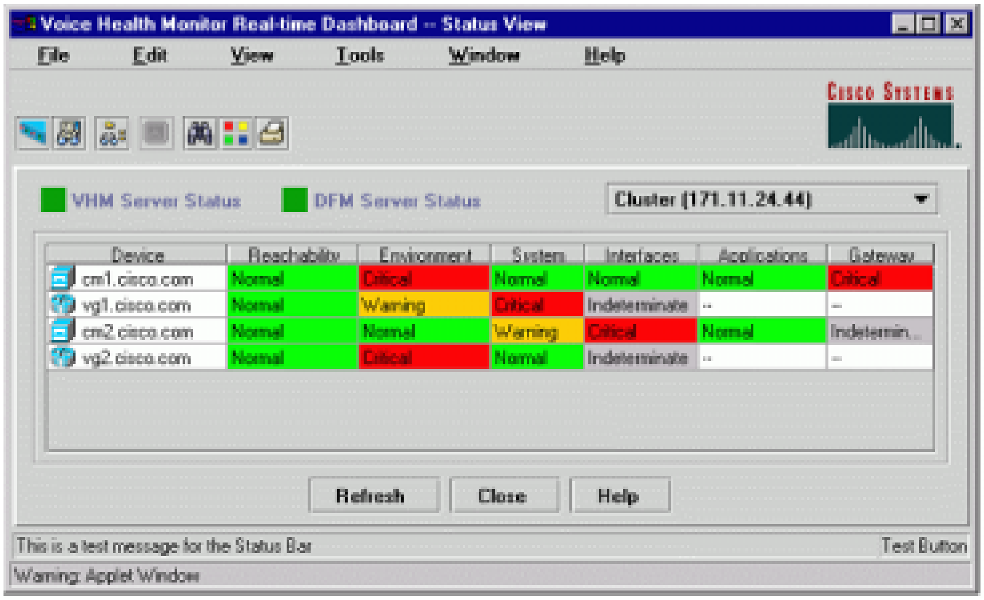

Voice Health Monitor Real-time Dashboard -- Status View

Cluster (171.11.24.44)

Device	Reachability	Environment	System	Interfaces	Applications	Gateway
cm1.cisco.com	Normal	Critical	Normal	Normal	Normal	Critical
vg1.cisco.com	Normal	Warning	Critical	Indeterminate	--	--
cm2.cisco.com	Normal	Normal	Warning	Critical	Normal	Indetermin...
vg2.cisco.com	Normal	Critical	Normal	Indeterminate	--	--

This is a test message for the Status Bar

Warning: Applet Window

The device fault manager (DFM) can display real-time device-specific information based on the CISCO-CCM-MIB and COMPAQ server MIB.

Voice Health Monitor displays real-time dashboards per device. The device detail window, shown in Figure 6-19, would be useful as a daily report of the health of the CallManager server. The CallManager system administrator could print this type of report daily to track resources and predict future trends.

Figure 6-19 Voice Health Monitor Device Detail Window

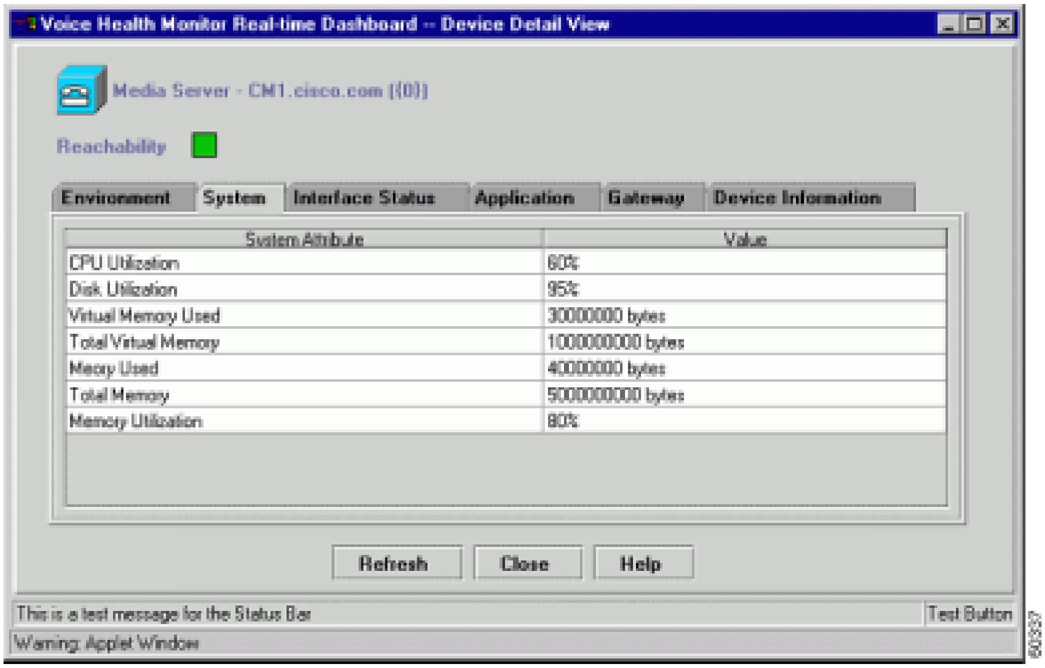

System Attribute	Value
CPU Utilization	60%
Disk Utilization	95%
Virtual Memory Used	30000000 bytes
Total Virtual Memory	1000000000 bytes
Meory Used	40000000 bytes
Total Memory	5000000000 bytes
Memory Utilization	80%

This is a test message for the Status Bar

Warning: Applet Window

Call Detail Record Management

Call detail records in the CallManager software allow detailed information to be captured and used for a variety of purposes. These records facilitate troubleshooting and provide the necessary data for implementing billing and performance management. Records are generated and written into a standard query language (SQL) database. You can access captured data using Open Database Connectivity (ODBC).

Refer to the *Cisco IP Telephony Troubleshooting Guide* for detailed information regarding call record fields. To view this document, go to the following location:
http://www.cisco.com/warp/public/788/AVVID/ts_ccm_301_sec1.htm.

Voice Quality Management with Cisco ART

The Administrative Reporting Tool (ART) for Cisco CallManager, a web-based reporting application, generates the following reports that provide information regarding voice quality and generates reports on the gateway performance:

- Quality of service
- Traffic details
- User call details
- Billing details
- Gateway details

- Call detail records

For detailed information about the Administrative Reporting Tool, go to the following location: http://www.cisco.com/univercd/cc/td/doc/product/voice/sw_ap_to/admin/admin_rp/index.htm.

Software Management

The data in Cisco CallManager can be backed up and restored using called Cisco MCS Backup and Restore utility. Regular backups of data in CallManager is strongly recommended to ensure data availability in the event of a system failure. Procedures for backing and restoring operating systems, hard drive, software, and data are available in Cisco CallManager documentation set or can be found at the following location: http://www.cisco.com/univercd/cc/td/doc/product/voice/c_callmg/3_0/index.htm.

Note It is important that operations personnel have the above mentioned documents easily accessible to ensure that problems are quickly resolved.

Security Management on MCS and CallManager

Access to MCS and CallManager is protected by passwords. Console and remote access to the MCS system and CallManager software are authenticated at the system level. The passwords for accessing the system and software should only be distributed to those responsible for managing the system. The default SNMP community strings should also be changed upon installation in order to reduce the risk of unauthorized access.

Troubleshooting

Cisco CallManager offers a number of built-in troubleshooting capabilities such as generating trace file (SDI, SDL, and CCM traces), diagnostics call detail records, and system events. The system events include system-level events generated by the MCS server and support for syslog messages in CallManager.

Refer to the *Cisco IP Telephony Troubleshooting Guide* for detailed troubleshooting information. To view this document, go to the following location: http://www.cisco.com/warp/public/788/AVVID/ts_ccm_301_sec1.htm.

Managing Cisco IP Phones

Addressing and Registration of IP Phones

Cisco IP phones provide the functionality available in regular analog phones currently deployed in the enterprise network. The IP phones also support features such as:

- Call forwarding
- Speed dialing
- Call transfers
- Conference calls
- Redialing

You must ensure that several components are working properly before you deploy the IP phones. Some of these components include:

- Dynamic Host Control Protocol (DHCP)

- Trivial File Transfer Protocol (TFTP)
- Domain Name Systems (DNS) servers

These servers provide the required services for automatically registering IP phones and their configuration settings. The configuration file containing information on phone and network settings is downloaded to the IP phone during the registration process.

Figure 6-20 shows the CallManager Phone Configuration window.

Figure 6-20 CallManager Phone Configuration Window

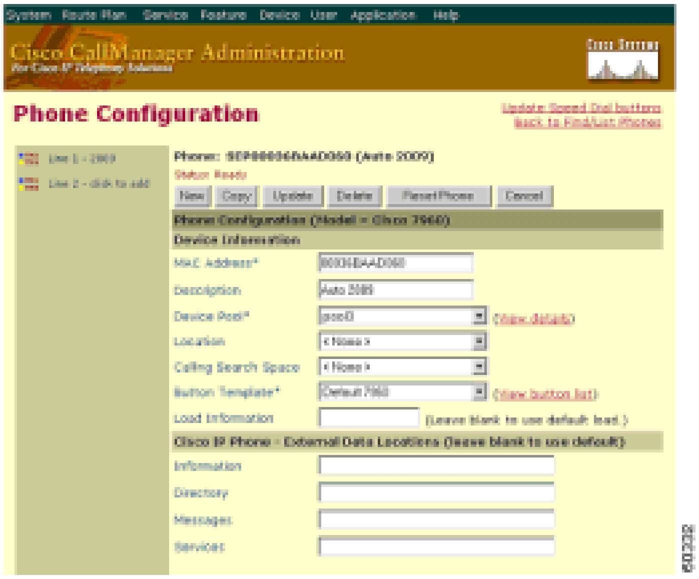

Phone Configuration

Phone: SEP00036BAAD060 (Auto 2009)

Status: Ready

Phone Configuration (Model = Cisco 7960)

Device Information

MAC Address* 00036BAAD060

Description Auto 2009

Device Pool* pool0 (View details)

Location < None >

Calling Search Space < None >

Button Template* Default 7960 (View button list)

Load Information (Leave blank to use default load.)

Cisco IP Phone - External Data Locations (leave blank to use default)

Information

Directory

Messages

Services

For additional information on how to set up and manage Cisco IP phones, go to the following location: http://www.cisco.com/univercd/cc/td/doc/product/voice/c_ipphon/7900/admingd/index.htm.

Power Management

Power for IP phones can be derived from three sources:

- External power source
- Switching module with inline power
- Inline power patch panel

You can integrate the management of switching modules into the network management system using SNMP and syslog support available in LAN switches hosting those modules. Syslog messages are supported by the switches to report power availability, port operational status, and link status. The LEDs on the switching module provide additional diagnostics and status indications on the hardware. Extensive SNMP support on the LAN switches allows you to remotely perform management functions.

For additional information on OIDs for polling the inline power status on Catalyst switches, go to the following location: http://www.cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml. (From this web page, select the OIDs link to view a listing of OIDs currently available.)

Troubleshooting and User Tracking

Cisco network management tools, such as Campus Manager, provide the added capability to trace voice calls placed on an IP network. The web-based application is capable of tracing data paths as well as troubleshooting signaling paths that VoIP calls use in the network. Layer 2 and layer 3 devices, along with IP phones, are displayed in a topology map showing individual hops taken by voice calls.

Figure 6-21 shows an example of a voice trace path as captured in Campus Manager.

Figure 6-21 Voice Trace Path Example

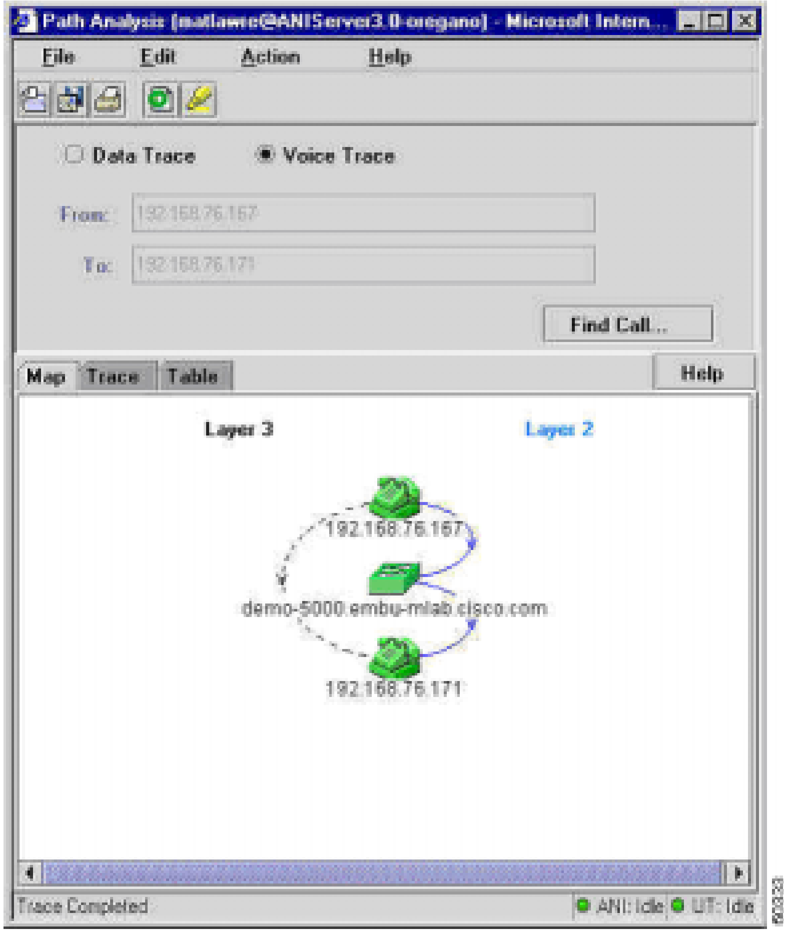
Path Analysis [mattawe@ANIServer3.0-oregano] - Microsoft Intern...
File
Edit
Action
Help
Find Call...
Map
Trace
Table
Help
Layer 3
Layer 2
192.168.76.167
demo-5000.embu-mlab.cisco.com
192.168.76.171
60333

○ Data Trace ● Voice Trace

From: 192.168.76.167

To: 192.168.76.171

Trace Completed ANI: Idle UT: Idle

The enhancement in the Campus Manager User Tracking software allows the application to display information associated with users and their phones. For troubleshooting purposes, MIBs supported by the CallManager provide detailed information such as:

- IP address
- Phone status
- User name
- E.911 location
- Port
- Device connection

For a list of SNMP MIB variables for managing CallManager and Cisco IP phones, go to the following location: http://www.cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml. (From this web page, select the SNMP v1 MIBs link to view a listing of OIDs currently available.)

Figure 6-22 shows an example of a user tracking phone table.

Figure 6-22 User Tracking Phone Table Example

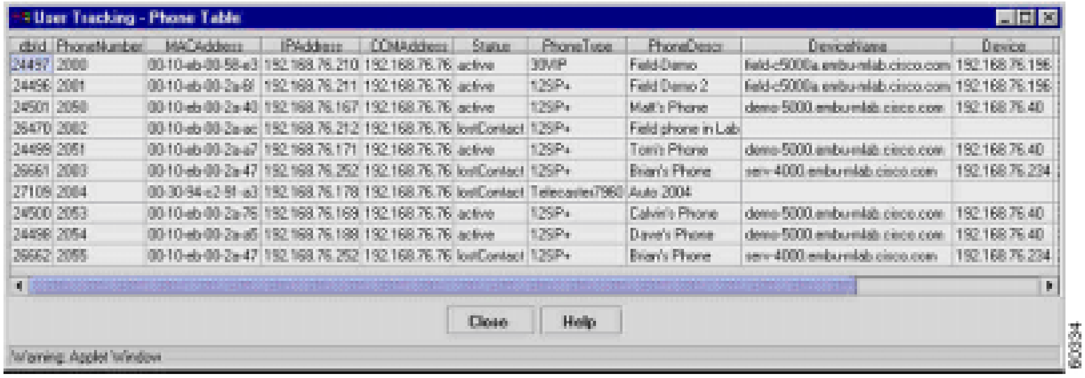

dbid	PhoneNumber	MACAddress	IPAddress	CCMAddress	Status	PhoneType	PhoneDescr	DeviceName	Device
24457	2000	00-10-eb-00-58-e3	192.168.76.210	192.168.76.76	active	30VIP	Field-Demo	field-c5000a.embu-mlab.cisco.com	192.168.76.196
24496	2001	00-10-eb-00-2a-8f	192.168.76.211	192.168.76.76	active	12SP+	Field Demo 2	field-c5000a.embu-mlab.cisco.com	192.168.76.196
24501	2050	00-10-eb-00-2a-40	192.168.76.167	192.168.76.76	active	12SP+	Matt's Phone	demo-5000.embu-mlab.cisco.com	192.168.76.40
26470	2002	00-10-eb-00-2a-ac	192.168.76.212	192.168.76.76	lostContact	12SP+	Field phone in Lab		
24499	2051	00-10-eb-00-2a-a7	192.168.76.171	192.168.76.76	active	12SP+	Tom's Phone	demo-5000.embu-mlab.cisco.com	192.168.76.40
26661	2003	00-10-eb-00-2a-47	192.168.76.252	192.168.76.76	lostContact	12SP+	Brian's Phone	serv-4000.embu-mlab.cisco.com	192.168.76.234
27109	2004	00-30-94-c2-91-a3	192.168.76.178	192.168.76.76	lostContact	Telecaster7960	Auto 2004		
24500	2053	00-10-eb-00-2a-76	192.168.76.169	192.168.76.76	active	12SP+	Calvin's Phone	demo-5000.embu-mlab.cisco.com	192.168.76.40
24498	2054	00-10-eb-00-2a-a5	192.168.76.188	192.168.76.76	active	12SP+	Dave's Phone	demo-5000.embu-mlab.cisco.com	192.168.76.40
26662	2055	00-10-eb-00-2a-47	192.168.76.252	192.168.76.76	lostContact	12SP+	Brian's Phone	serv-4000.embu-mlab.cisco.com	192.168.76.234

Refer to the *Cisco IP Telephony Troubleshooting Guide* for detailed CallManager troubleshooting information. To view this document, go to the following location: http://www.cisco.com/warp/public/788/AVVID/ts_ccm_301_sec1.htm.

NMS Reference Architecture

Figure 6-23 shows a reference architecture that Cisco Systems believes should be the minimal solution for managing a data network. The architecture includes the following:

- SNMP platform for fault management
- Performance monitoring platform for long-term performance management and trending
- CiscoWorks2000 server with LMS bundle for configuration management, syslog collection, and hardware and software inventory management

Figure 6-23 NMS Reference Architecture

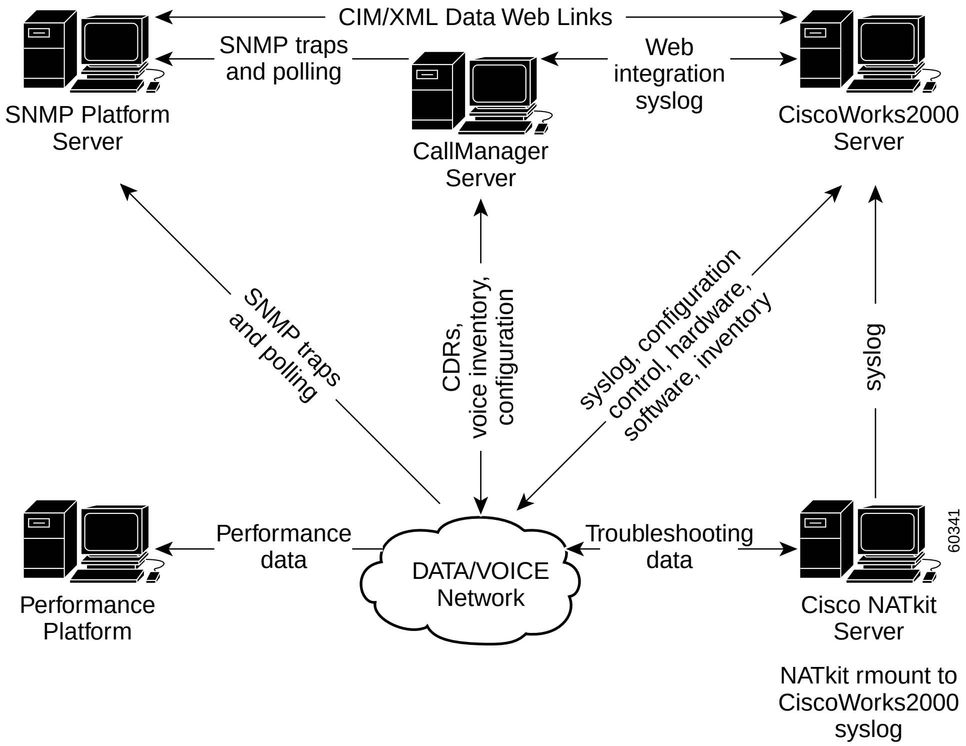

Some SNMP platforms can directly share data with the Cisco Works2000 server using CIM/XML methods. Cisco's Network Supported Accounts (NSA) customers can also include Cisco's NATkit server for additional proactive monitoring and troubleshooting support from their NSA engineer. The NATkit server would either have a remote disk mount (rmount) or ftp access to the data residing on the CiscoWorks2000 server. This will reduce the number of syslog receivers configured in your Cisco devices and reduce the CPU load caused by sending duplicate messages.

- For information about Cisco Network Supported Accounts, go to the following location: http://www.cisco.com/warp/public/cc/serv/mkt/sup/ent/nsa/nsa_pl.htm.
- For information about Cisco NATkit, go to the following location: http://www.cisco.com/univercd/cc/td/doc/product/natkit/natk2010/5_nkitov.htm.

In addition to hardware and software components, an effective NMS implementation must include using well-defined processes and procedures. The following subsections describe suggestions based on conversations and observations between Cisco personnel and Cisco customers.

- "Planning, Design, and Implementation"
- "Daily Operations"

Planning, Design, and Implementation

1. Include NMS in your initial network design. Ensure that CPU, memory, and other hardware specifications include the added overhead of running NMS protocols. Also, identify what NMS configuration commands will be in the network devices and build a template for NMS configurations that will be used as your company standard.

2. Design an implementation and acceptance test plan for the NMS system using a checklist. Ensure that all features work as expected and document them.
3. Identify all possible important events and alert conditions that the network operations group will receive (see Table 6-20 for an example). Write Service Level Agreements (SLAs) per event and alert that include:
   - Time to react
   - Time to respond
   - Time to repair
   - Escalation procedures
   - Personnel
4. Periodically test NOC response by sending test alerts to make sure that the processes are being followed.

Table 6-20 Network Event Notification (SNMP) Requirements

Priority	Category/Requirement	SNMP Get Support[1]	SNMP Trap Support[1]	NOC Action
Frame Relay				
1	State change for DLCI	Yes—RFC1315-MIB.my	Yes	Tier 1 action Tier 2 Escalation/Debug
2	FECNs	Yes—CISCO-FRAME-RELAY-MIB.my	No	Tier 2 Debug
2	BECNs	Yes—CISCO-FRAME-RELAY-MIB.my	No	Tier 2 Debug
3	Dropped or lost packets	Yes—CISCO-FRAME-RELAY-MIB.my	No	Tier 2 Debug
3	Discard eligible inbound	Yes—CISCO-FRAME-RELAY-MIB.my	No	Tier 2 Debug
3	Discard eligible outbound	Yes—CISCO-FRAME-RELAY-MIB.my	No	Tier 2 Debug
X.25				
1	State change for X.25 SVC (up/down indication)	Yes—CISCO-CALL-HISTORY-MIB.my	No	Tier1 Action Tier2 Escalation/Debug
2	Max. call setup based on the HTC setting for the interface	Yes—RFC1382-MIB.my	No	Tier 2 Debug
2	Retransmission attempts exceeded		No	Tier 2 Debut
General Information				
2	CRC errors	Yes—CISCO-INTERFACES-MIB.my	No	Tier 2 Debug
3	Carrier transitions	Yes—CISCO-INTERFACES-MIB.my	No	Tier 2 Debug
3	Input errors	Yes—CISCO-INTERFACES-MIB.my	No	Tier 2 Debug
3	Output buffer failures	Yes—CISCO-INTERFACES-MIB.my	No	Tier 2 Debug

[1] Support for the Cisco 7206 VXR

- Priority—1 = Critical, 2 = Debug, 3 = Informational

- SNMP Trap—Threshold trigger point generation
- CISCO 7206 VXR SNMP get support
  - SNMP MIB extension for X.25 LAPB (RFC 1381)
  - SNMP MIB extension for the X.25 packet layer (RFC 1382)
  - Exempted; The LAPB XID table, X.25 cleared circuit table, X.25 call parameter table
- NOC Action—Tier 1 = Help desk team, Tier 2 = Escalation to network administration team

Daily Operations

1. Assign at least one operations person to proactively review syslogs on a daily basis and forward important messages to appropriate personnel.
2. Assign at least one operations person to review performance data on a daily basis to identify trends, and study network operational characteristics.
3. Ensure that operations personnel have access to logins and mission-critical applications and devices such as Call Manager, IP phones, CiscoWorks2000, and mainframe applications. It has been shown by numerous Cisco customers that those whose operations personnel study how applications function on the network are more able to identify problems and correct actions.
4. Assign at least one engineer-level person to study performance data on weekly, monthly, and yearly time lines to identify trends and trouble characteristics.
5. Assign at least one engineer-level person to proactively study and document network characteristics and behaviors of a section of the network that is considered by its users to have acceptable response time and stability. This provide a baseline of well-behaving parts of the network.

Managing Cisco CallManager with CISCO-CCM-MIB

The Cisco CallManager MIB table, shown in Table 6-21, lists the most important MIB objects to monitor to ensure that the CallManager, gateways, trunks, and phones are working. These MIB objects (OIDs) may be polled by your NMS platform at the recommended interval. The thresholds are general recommendations that may need adjustment on your network. These particular OIDs fall under fault management and when they exceed the indicated threshold this indicates a major system problem and should be acted upon immediately by network operations personnel. In order for your NMS platform, HPOV, Tivoli, and so forth, to poll these objects, the CISCO-CCM-MIB.my file must be loaded into the system.

Refer to your SNMP platform vendor’s documentation to find out how to begin proactive polling and threshold monitoring.

Table 6-21 SNMP MIBs for Fault Management

CISCO-CCM-MIB Object Name	Object Description	OID	Poll Interval	Threshold
ccmStatus	The current status of the CallManager.	.1.3.6.1.4.1.9.9.156.1.1.2.1.5	30 mins.	≠ 2
ccmPhoneStatus	The state of the phone. The state of the phone changes from Unknown to Active when it registers itself with the local CallManager.	.1.3.6.1.4.1.9.9.156.1.2.1.1.7	30 mins.	≠ active
ccmGatewayStatus	The state of the gateway. The gateway status changes from Unknown to Registered when the gateway registers itself with the local CallManager.	.1.3.6.1.4.1.9.9.156.1.3.1.1.5	30 mins.	≠ registered
ccmGatewayTrunkStatus	The state of the trunk. The Trunk status changes from Unknown to Up when it registers itself withe the local CallManager.	.1.3.6.1.4.1.9.9.156.1.4.1.1.5	30 mins.	≠ up
ccmActivePhones	The number of phones connected to this CallManager and actively in communication (by means of keepalives) with this CallManager.	.1.3.6.1.4.1.9.9.156.1.5.1	1 hour	any unplanned delta
ccmInActivePhones	The number of phones that are registered with the CallManager but have lost contact with the CallManager. The phones are said to have lost contact with the CallManager if the CallManager does not receive any keepalives.	.1.3.6.1.4.1.9.9.156.1.5.2	15 mins.	>0
ccmActiveGateways	The number of gateways configured with this CallManager and actively in communication (by means of keepalives) with the CallManager.	.1.3.6.1.4.1.9.9.156.1.5.3	1 hour	any unplanned delta
ccmInActiveGateways	The number of gateways that are registered with the CallManager but have lost contact with the CallManager. The gateways are said to have lost contact with the CallManager if the CallManager does not receive any keepalives.	.1.3.6.1.4.1.9.9.156.1.5.4	15 mins.	>0

Table 6-22 lists the SNMP trap messages that are automatically added to your NMS platform when you load the CISCO-CCM-MIB.my file. The CallManager server will send these error messages to your NMS station when you set up the CallManager to send SNMP traps.

Table 6-22 CallManager SNMP Traps

SNMP Trap	OID
ccmCallManagerFailed	.1.3.6.1.4.1.9.9.156.2.0.1
ccmPhoneFailed	.1.3.6.1.4.1.9.9.156.2.0.2
ccmPhoneStatusUpdate	.1.3.6.1.4.1.9.9.156.2.0.3
ccmGatewayFailed	.1.3.6.1.4.1.9.9.156.2.0.4
ccmOutOfResource	.1.3.6.1.4.1.9.9.156.2.0.5
ccmGatewayLayer2Change	.1.3.6.1.4.1.9.9.156.2.0.6

Summary of IP Telephony Network Management Products

Table 6-23 lists the available network management tools, and the IP Telephony components supported for each tool.

Table 6-23 IP Telephony Network Management Products

Management Tool/System	Supported IP Telephony Component	Function
Infrastructure Management		
CiscoWorks2000		
Resource Manager Essentials	Routers Catalyst LAN switches VoIP gateways IP phones	Configuration file management Syslog management Software image management Inventory management
Campus Manager	Routers Catalyst LAN switches VoIP gateways IP phones	Discovery and topology mapping End station and handset tracking Layer 2/layer 3 path analysis VLAN/LAN and ATM configuration
Traffic Director	Routers Catalyst LAN switches	Traffic analysis for RMON/RMON2 Trending Real-time analysis
Internetwork Performance Monitor	Routers VoIP gateways	Performance indicators using the Service Assurance Agent (SAA) feature
Service Level Manager and Management Engine 1110	Routers VoIP gateways	End-to-end service monitoring
CiscoView	Routers Catalyst LAN switches VoIP gateways	Device-level management and troubleshooting
Cisco Voice Manager	VoIP gateways	Dial plan management

Table 6-23 IP Telephony Network Management Products (continued)

Management Tool/System	Supported IP Telephony Component	Function
QoS Policy Manager	Routers Catalyst LAN switches VoIP gateways	Defines QoS policies on devices
User Registration Tracking	Catalyst LAN switches	Creates user registration policy bindings
Cisco Secure ACS	Routers Catalyst LAN switches VoIP gateways	Security server for authentication, authorization, and accounting
Other Infrastructure Components		
Syslog server	Routers Catalyst LAN switches VoIP gateways CallManagers	Syslog service
SNMP trap receiver	Routers Catalyst LAN switches VoIP gateways	Receives SNMP traps
NTP server	Routers Catalyst LAN switches VoIP gateways	Time synchronization
CallManager Management		
Microsoft Event Viewer	CallManager	Displays system-level messages
Microsoft Performance Viewer	CallManager	Displays performance of system (CPU, disks)
Compaq Insight Manager	Media Convergence Server	Provides service-specific statistics
IP Phone Management		
Network Registrar	IP phones	DNS and DHCP services

Securing IP Telephony Networks

This section provides a guide to designing and implementing secure IP Telephony networks. It adheres to the Cisco secure blueprint for enterprise networks (SAFE) security architecture. The information in this section is not meant to be all encompassing, but rather a starting point for network managers, system administrators, and systems engineers to use when building the IP Telephony network. It is important to keep in mind that securing a network is a continual process that requires keeping abreast of the latest vulnerabilities that may exist in network infrastructure components, server operating systems, and applications deployed throughout the enterprise network.

The subject of securing voice communications, which is typically a sensitive topic for communications architects, has received even greater visibility as network convergence becomes the accepted design model. With the advent of IP Telephony, which uses IP data network devices for voice communication, the potential exists for call processing components and telephony applications to be compromised by malicious attacks.

As computer technology has become a vital part of our lives, the requirement to protect against security vulnerabilities becomes ever more important. The challenge facing many network and systems engineers is that the hacking community almost instantly publicizes new system vulnerabilities, as well as the tools used to crack the systems, making vulnerabilities easily exploitable for even inexperienced programmers. The convergence of the voice, video and data networks combined with the availability of advanced *cracking* tools, accessible to novice Internet users, further highlights the requirement for initiating a sound security design practice.

Security Policy Best Practices

The first step in any security implementation is to establish a security policy. This topic is not covered in this document; instead, a series of recommended *best practices* is outlined. Implementation of these best practices is dependent on the organization’s security policy. The following best practices are described in this section:

- Establishing Physical Security

  Creating a secure physical boundary for critical communications equipment is a fundamental foundation in building secure networks. Network designs and software configurations cannot protect a network whose assets are not physically protected from potential malicious threats.

- Protecting the Network Elements

  Routers, Ethernet switches and VoIP gateways define network boundaries and act as gateway interfaces to all networks. Securing these vital pieces of voice and data networks is a requirement for securing the data, voice, and video application running across the infrastructure.

- Designing the IP Network

  Understanding and following sound IP network design principles not only allows the network to scale and perform, but also increases the security of all attached devices.

- Securing the CallManager Server

  Securing the actual voice call processing platform and installed applications is the last, and perhaps most vital step in securing IP Telephony networks.

Each of these steps is analyzed and configuration examples are listed below.

Establishing Physical Security

As with most computing devices, Cisco routers and switches, servers, and other infrastructure components are not designed to provide protection against penetration or destruction by an attacker with direct physical access. Reasonable steps must be taken to prevent physical access by unauthorized personnel.

Common precautions include restricting access to wiring closets and data centers to staff that are considered *trusted*; generally, such staff already have direct or indirect logical access to the devices being protected and therefore gain no advantage from physical access. In data centers where non-trusted staff may be present, the practice of separate locking cabinets for individual items or racks that have more stringent security requirements can be followed. When using keyed or electronic locks on doors, consider any facilities, security, and janitorial staff that may have the ability or clearance to bypass the locks.

The policy defining who has physical access to the organization’s infrastructure, communications, and power equipment should be as carefully evaluated as who has system or network administrator login permissions. Physical access logs should be carefully maintained. Additionally, the computer which maintains the physical access logs should use a synchronized time source, such as the Network Time Protocol (NTP), so the logs can be accurately cross checked with logs of other computers, network elements, and communications equipment.

An additional step is to add IP video surveillance equipment into each data center and wiring closet. The video recordings can be used to track all access into a given secure environment by linking with the electronic access system. Video archives can be searched by time of day or even changes to physical locations. For instance, the video logs can be searched for all recorded footage of a systems cabinet opening. Additionally, a real-time alarming capability can be enabled. These recording tools can help catch intruders, track down attack origination, or be used as evidence in legal proceedings.

Protecting the Network Elements

Once physical security has been established, the next step is to secure the actual routers, switches, and VoIP gateways that make up the communications network. These network elements provide both the physical and logical connectivity of the entire enterprise network and can be a major target of a well-informed hacker.

Refer to the *Improving Security on Cisco Routers* tech note for up-to-date information for securing Cisco routers. This document can be found at the following location:
http://www.cisco.com/warp/public/707/21.html

Telnet Access

Configuration from the Command Line Interface (CLI) by means of a telnet session is still the most popular method of managing routers and switches. Therefore, the first step in securing these network elements is to limit which subnets are able to access the router and switch virtual terminal sessions. Limiting virtual console access to the IP address range(s) of operations staff and network management hosts is a useful method to prevent unauthorized users from accessing network devices, even if a password is discovered.

Cisco IOS Routers

Example 6-1 defines a Cisco IOS router access list that permits only hosts on the network admin subnet 172.21.167.0 to connect to the virtual console.

Example 6-1 Cisco IOS Router Access—Hosts to Virtual Console

```
access-list 12 permit 172.21.167.0 0.0.0.255
line vty 0 4
   access-class 12 in
   login
   password g0+$k1
```

Catalyst Ethernet Switches

Example 6-2 defines an IP permit list that permits only hosts on network 172.21.167.0 to connect to the virtual console of a Catalyst Ethernet switch.

Example 6-2 IP Permit List—Hosts to Virtual Console

```
set ip permit 172.21.167.0 0.0.0.255 telnet
set ip permit enable telnet
```

The functionality of the IP permit list can also be achieved with VLAN access control lists (VACLs). Because VACLs are processed by hardware (Policy Feature Card [PFC]), VACL processing is considerably faster than IP permit list processing.

Example 6-3 defines a VACL mapped to VLAN 10 that permits hosts on network 172.21.167.0 to use Telnet to access the Catalyst Ethernet switch virtual console at IP address 172.21.167.1 and blocks all other Telnet access to the virtual console while allowing all other traffic.

Example 6-3 VACL Mapped to VLAN—Telnet Access From Hosts

```
set security acl ip ACCESS permit tcp 172.21.167.0 0.0.0.255 host 172.21.167.1 eq 23
set security acl ip ACCESS deny tcp any host 172.21.167.1 eq 23
set security acl ip ACCESS permit any any
commit security acl ACCESS
set security acl map ACCESS 10
```

Figure 6-24 shows the result of limiting which subnets are able to access the router and switch virtual terminal sessions.

Figure 6-24 Telnet Access and Restrictions

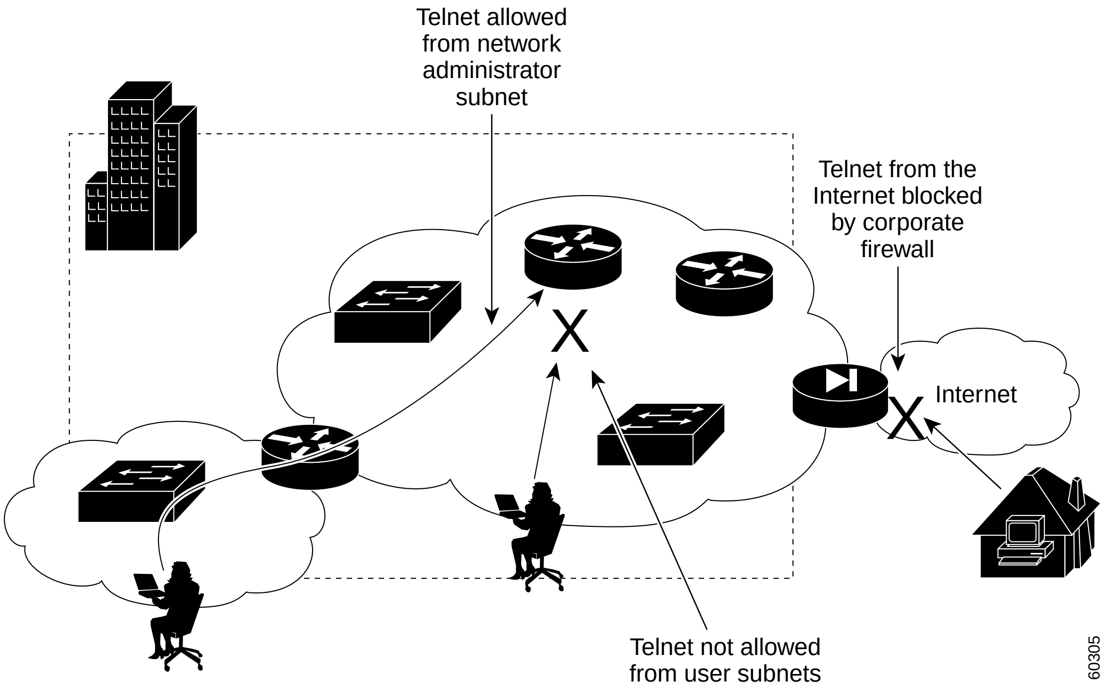

Passwords and Authentication Using TACACS+ and RADIUS

Passwords are another important line of defense against unauthorized access to routers and switches. The best way to control user passwords is to maintain them on a TACACS+ or RADIUS authentication server. However, many routers will still have a locally configured last-resort password for privileged access during required authentication server maintenance periods. To ensure that the maximum precautions are taken for limiting router password exposure, all passwords should be encrypted using the **service password-encryption** command. Additionally, it is recommended that the enable secret be used to further hide configuration access. The **service password-encryption** command uses a weak algorithm, the simple Vigenere cipher, to encrypt the enable password. A competent cryptographer can reverse this cipher fairly quickly. However, the **enable secret** command uses a strong MD5 one-way hashing algorithm to encrypt the secret password. It is important to note that the MD5 hashes can be subject to dictionary attacks, so always use passwords that include non-alphanumeric symbols.

Choosing Good Passwords

Ideally, you should use one-time password systems such as SofToken, SecurID, or DES Gold Cards to prevent an attacker from reusing trusted users' passwords. However, many organizations consider these tools to be cumbersome or too expensive. If you use normal passwords, they should be chosen so they cannot be guessed by using dictionary words in sequence or by observing their proper user. For maximum security, there should be numbers or punctuation symbols, as well as mixed case letters. Included below are examples of the recommended maintenance password configurations. In Example 6-4, a user must type in the enable secret password to enter enable mode.

Example 6-4 Enable Password on Cisco IOS Routers

```
enable password Go-5m4LL
enable secret g0-B1g!
service password-encryption
```

AAA with TACACS+ and RADIUS

Once the last-resort maintenance passwords have been securely configured, users' passwords can be configured for everyday maintenance and configuration of the network elements. A sound Authentication, Authorization and Accounting (AAA) scheme should be used for all user access to network elements.

Using either RADIUS or TACACS+ to perform per-user AAA functions is vital to enhance security and accountability for infrastructure devices. These features enable centralized administration of account and password information, eliminating per-device maintenance efforts (and errors) when trusted users are added or removed. Also, CiscoSecure ACS can be configured to require *strong* passwords, password aging, and intrusion lockout.

When using RADIUS or TACACS+ for AAA, each user has their own password for accessing network devices. If the AAA server is unavailable, then a shell can be opened using the local username/password database stored in the router's configuration. In this configuration, physical security is assumed and AAA is not used on the console port. Example 6-5 and Example 6-6 shows AAA enabled on Cisco IOS routers using RADIUS and on Catalyst switches using TACACS+.

Example 6-5 Enable AAA on Cisco IOS Routers using RADIUS

```
username todd password 5y5c0jamS
aaa new-model
aaa authentication login default radius local enable
aaa authorization exec default group radius local none
aaa authorization network default group radius local none
aaa accounting exec default start-stop group radius
aaa accounting connection default start-stop group radius
aaa accounting system default start-stop group radius
aaa accounting update periodic 60
radius-server host 10.21.101.10
radius-server key 2B-$Ecur3
line vty 0 4
login authentication default
```

Example 6-6 Enable AAA on Catalyst Ethernet Switches Using TACACS+

```
set tacacs server 10.21.101.10
set tacacs key 2B-$Ecur3
set tacacs attempts 3
set authentication login tacacs enable all primary
set authorization exec enable tacacs+ deny both
set accounting connect enable stop-only tacacs+
set accounting exec enable stop-only tacacs+
set accounting system enable stop-only tacacs+
set accounting system commands enable all stop-only tacacs+
set accounting update periodic 60
```

Secure Shell

Secure Shell (SSH) is an application that allows a secure, encrypted shell session from a client to a Cisco router. The SSH version 1 server running in select versions of Cisco IOS is compatible with publicly available SSH client software. The Cisco SSH server supports both DES and 3DES encryption and user authentication.

Note Cisco IOS software images with strong encryption (including, but not limited to, 56- and 168-bit data encryption feature sets) are subject to both United States and International government export controls. Contact your sales representative or distributor for more information, or send e-mail to export@cisco.com.

Example 6-7 illustrates the configuration steps required to enable SSH in Cisco IOS.

Example 6-7 Enable SSH in Cisco IOS

```
Hostname secure-router
ip domain-name cisco.com
crypto key generate rsa
ip ssh time-out 60
ip ssh authentication-retries 5
```

A user name or AAA authentication server must be configured for users to access the router via SSH. Also, a valid domain name must be specified and name service look-ups must be part of the router configuration. Currently, Cisco IOS supports only SSH version 1.

Simple Network Management Protocol

SNMP is widely used for monitoring and configuring network elements. SNMP uses an authentication method based on a community string. This community string is essentially a password used for accessing the network element. SNMP version 1 sends this password string in clear-text across the network. SNMP version 2 supports an MD5-based digest authentication scheme and allows for restricted access to various management data. Because nearly all SNMP implementations send the community string as part of the polling transaction process, it is highly recommended to use SNMP version 2 as a minimum if SNMP is utilized. Currently, SNMPv3 is not supported by Cisco CallManager version 3.0(7).

Adhere to the following three basic rules for using SNMP securely:

1. Never use *public* and *private* as community strings
2. Limit SNMP access to only a few specific hosts or subnets
3. Use SNMPv2 if possible

Since nearly all of the information that can be viewed or configured from a virtual console can also be accessed via SNMP, it is essential to restrict this method of access as completely as possible. Only those hosts with a verified need to perform SNMP writes should have full access.

Example 6-8 to Example 6-11 define an access list that permits only hosts on network 10.21.101.0 to perform SNMP reads with the community *ph0oBar* and only the host 10.21.101.10 to perform SNMP writes with the community *ph0oBaz*.

Example 6-8 Access List Permit—Cisco IOS Routers

```
access-list 12 permit 10.21.101.0 0.0.0.255
snmp-server community ph0oBar RO 12
access-list 13 permit host 10.21.101.10
snmp-server community ph0oBaz RW 13
```

Example 6-9 Access List Permit—Catalyst Ethernet Switches

```
Set snmp community read-only ph0oBar
Set snmp community read-write ph0oBaz
```

IP Permit command statements or VACL statements can be used to restrict SNMP access to the switch.

Example 6-10 IP Permit Command Statements

```
set ip permit 10.21.101.0 0.0.0.255 snmp
set ip permit enable snmp
```

Example 6-11 VACL Command Statements

```
set security acl ip ACCESS permit udp 10.21.101.0 0.0.0.255 host 172.21.167.1 eq 161
set security acl ip ACCESS permit udp 10.21.101.0 0.0.0.255 host 172.21.167.1 eq 162
set security acl ip ACCESS deny udp any host 172.21.167.1 eq 161
set security acl ip ACCESS deny udp any host 172.21.167.1 eq 162
set security acl ip ACCESS permit any any
commit security acl ACCESS
set security acl map ACCESS 10
```

HTTP

Web configuration is disabled by default on most platforms; however, novice network administrators often enable it. If disabling the service is not feasible, restrict HTTP access to specific management addresses.

Note If HTTP configuration is not necessary, it should be disabled.

Example 6-12, a router Cisco IOS router example, defines an access list that permits only hosts on network 10.21.101.0 to connect to the HTTP server.

Example 6-12 Access List Permit—Hosts to HTTP Server

```
access-list 12 permit 10.21.101.0 0.0.0.255
ip http access-class 12
```

On Catalyst switches, the HTTP server can be disabled with the **set ip http server disable** command.

Cisco Discovery Protocol

Cisco Discovery Protocol (CDP) is a Layer 2 protocol used for network management. It can be dangerous since it allows any directly connected system to learn that the router is a Cisco device, and to determine the model number and the Cisco IOS software version being run. This information could be used to attack the router. The CDP protocol can be disabled with the **no cdp running** global configuration command or it can be disabled on individual interfaces with the **no cdp enable** command.

On Catalyst switches, CDP can be disabled globally with **set cdp disable** or on individual ports with **set cdp disable mod_num/port_num**.

Warning Banners

The **banner login** command should be used to notify authorized and unauthorized users that their activities are monitored and that unauthorized usage is forbidden and may be prosecuted.

Note Your banner typically should not contain specific information about your organization, the router, or your network.

Specific banner contents vary by organization and location. The following information can be used as a guide:

- A notice that the system is to be logged into or used only by specifically authorized personnel, and perhaps information about who may authorize use.
- A notice that any unauthorized use of the system is unlawful, and may be subject to civil and/or criminal penalties.
- A notice that any use of the system may be logged or monitored without further notice, and that the resulting logs may be used as evidence in court.
- Specific notices required by specific local laws.

Line Configuration

Virtual Terminal Lines (VTY) should be configured to only accept connections with required protocols by using the **transport input** command. If the VTY should only receive Telnet sessions the command is as follows:

transport input telnet

If Telnet and SSH are required the command should be:

transport input telnet ssh

Note If possible, only SSH should be used as the connection protocol, and clear-text Telnet should be disabled.

VTY timeouts should be configured using the **exec-timeout** command. This prevents an idle session from consuming a VTY indefinitely. Also, TCP keepalives can be used for incoming connections using the **service tcp-keepalives-in** command. This can help protect against attacks and *orphaned* sessions.

On Catalyst switches, the session idle timeout default is 20 minutes. This can be configured with the **set logout *timeout*** command.

Finger and TCP/UDP Small Servers

By default, several services are enabled which either allow an attacker to more easily consume device resources, indirectly attack other hosts, or gain information about operations staff currently accessing the network device. You can disable these services to prevent malicious use of these services or the information they may provide. Example 6-13 shows disabled small server and finger services.

Example 6-13 Disabled Services on Cisco IOS Routers

```
no service tcp-small-servers
no service udp-small-servers
no service finger
```

Remote Copy/Remote Shell

Some network administrators use the Berkeley Remote Copy (RCP) command to copy files to a device and the Remote Shell (RSH) command to execute commands without logging in. However, be aware that these services have extremely weak authentication and should not be enabled unless no other options (such as SSH support in Cisco IOS version 12.1T) are available.

Example 6-14 Disabled RCP and RSH Services on Cisco IOS Routers

```
no ip rcmd rcp-enable
no ip rcmd rsh-enable
```

Neighbor Authentication

If you're using a dynamic routing protocol that supports authentication, it's a good idea to enable that authentication. This prevents malicious attacks targeting the routing intelligence of the infrastructure, and can also help to prevent damage caused by misconfigured *rogue* devices on the network. Most common networking protocols provide a means for neighbors to authenticate each other to ensure that unauthorized devices are not allowed to affect the stability or security of the network. These authentication mechanisms prevent casual attempts at disrupting proper operation, but should not be expected to stop a determined attacker.

HSRP

Example 6-15 shows an enabled authentication string “hsRp$af3” for HSRP group 21 on interface Ethernet 2/1.

Example 6-15 Enabled HSRP Authentication

```
interface Ethernet 2/1
standby 21 authentication hsRp$af3
```

Enhanced IGRP (EIGRP)

Example 6-16 shows enabled EIGRP authentication between two interfaces on a Gatekeeper and 3640 router.

Example 6-16 Enabled EIGRP Authentication

```
hostname denlab_gk
key chain gk-2-denver
    key 1
        key-string d3nV3r_gk
    key 2
        key-string d3nV3r_3640
!
interface Ethernet0
    ip address 10.21.1.200 255.255.255.0
    ip authentication mode eigrp 247 md5
    ip authentication key-chain eigrp 247 gk-2-denver


hostname denver_3640
key chain denver-2-gk
    key 1
        key-string d3nV3r_gk
    key 2
        key-string d3nV3r_3640
!
interface FastEthernet0/0
    ip address 10.21.1.1 255.255.255.0
    ip authentication mode eigrp 247 md5
    ip authentication key-chain eigrp 247 denver-2-gk
```

OSPF

Example 6-17 shows enabled OSPF authentication string “0spFmd5” for Area 2 on interface Ethernet 2/1.

Example 6-17 Enabled OSPF Authentication

```
router ospf 1
    area 2 authentication message-digest
!
interface Ethernet 2/1
```

```
   ip ospf message-digest-key 1 md5 0spFmd5
```

BGP

Example 6-18 shows an enabled BGP MD5 authentication string “BgP%md5” for the connection to neighbor 171.70.209.2.

Example 6-18 Enabled BGP Authentication

```
router bgp 1
   neighbor 171.70.209.2 password BgP%md5
```

Setting Device Time, Timestamps, and Logging

Accurate logging is still one of the most valuable tools for tracking down intruders. In order to obtain an accurate log, the following steps must be taken on all network elements:

Step 1 Configure all devices to use an accurate, centralized time source.

Step 2 Enable time stamping on Cisco IOS devices.

Step 3 Designate a syslog server to receive logging information.

The Network Time Protocol

The Network Time Protocol (NTP) is a protocol designed to time-synchronize a network of machines. NTP runs over UDP, which in turn runs over IP. NTP is documented in RFC 1305. An NTP network usually gets its time from an authoritative time source, such as a radio clock or an atomic clock attached to a timeserver. NTP then distributes this time across the network. NTP is extremely efficient; no more than one packet per minute is necessary to synchronize two machines to within a millisecond of one another.

In Figure 6-25, the enterprise designates a single router to query a Stratum 1 clock. In turn, all other network devices query this single router for NTP information. It is recommended to enable NTP authentication.

Figure 6-25 Network Time Protocol

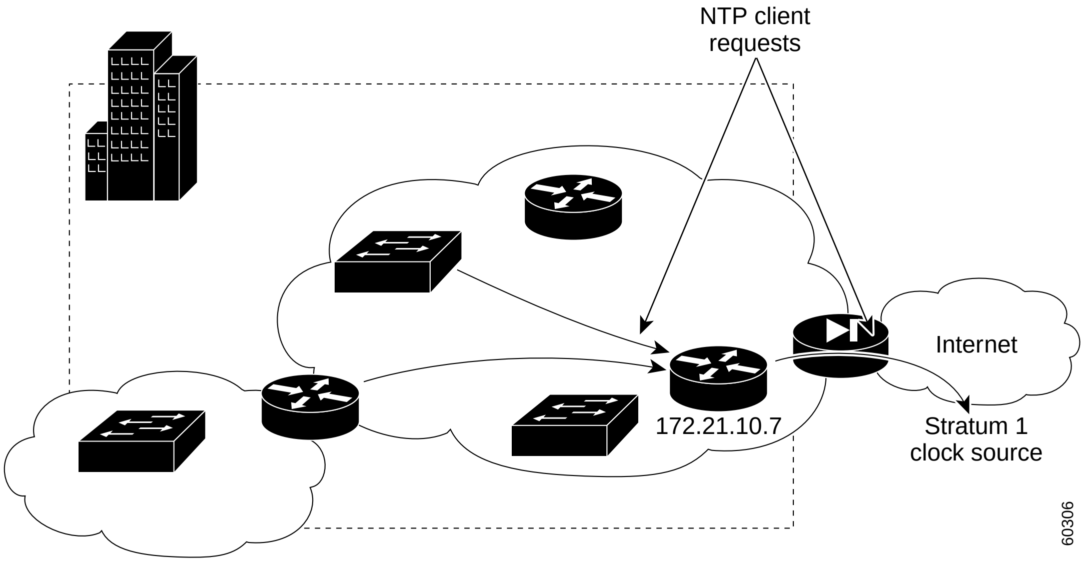

Example 6-19 Network Time Protocol Enabled on Cisco IOS Router

```
clock timezone MST -7
clock summer-time MDT recurring
ntp update-calendar
ntp server 172.21.10.7
ntp authenticate
ntp authentication-key 1 md5 nTp-ru1e$
ntp trusted-key 1
```

Example 6-20 Network Time Protocol Enabled on Catalyst Ethernet Switch

```
set timezone Mountain -7
set summertime enable MDT
set summertime recurring
set ntp server 172.21.10.7
set ntp client enable
set ntp public_my trusted md5 nTp-ru1e$
set ntp authentication enable
```

Time Stamps

Once all devices are configured to use a central time source, accurate time stamping for logs must be enabled on both the routers and switches.

Example 6-21 Time Stamp Enabled on Cisco IOS Router

```
service timestamps debug datetime msec
service timestamps log datetime msec
```

Example 6-22 Time Stamp Enabled on Catalyst Ethernet Switch

```
set logging timestamps enable
```

Syslog Server

Logging all system notices and error messages often provides valuable insight into the operational status of network devices. If access list violations are logged, the logs may also be correlated between devices to determine that the network is being probed or that a device has been compromised.

The following examples show logging enabled to a syslog server at 10.21.101.10.

Example 6-23 Logging Enabled—Cisco IOS Router

```
logging 10.21.101.10
logging facility local5
logging trap 5
```

Example 6-24 Logging Enabled—Catalyst Ethernet Switches

```
set logging server 10.21.101.10
set logging server facility local5
set logging server severity 5
set logging server enable
```

Table 6-24 lists the available logging severity levels.

Table 6-24 Logging Severity Levels

Severity Level	Description
0—emergency	System unusable
1—alert	Immediate action required
2—critical	Critical condition
3—error	Error condition
4—warning	Warning condition
5—notification	Normal bug-significant condition
6—informational	Informational message
7—debugging	Debugging message

Designing the IP Network

Before IP phones can be installed, a sound and secure IP network must be designed. The design must include the following:

- Separate broadcast domains
- A logical association of IP Telephony equipment
- Isolated IP Telephony management servers
- Security relationships
- Secure perimeters from both outside attackers and internal users

When building the secure IP Telephony network, the following rules should be followed:

- Place all CallManagers, IP Telephony application servers, and IP telephones on their own, separate IP networks (VLANs). These VLANs should also be different from any used by the organization’s data networks. Where possible, use RFC 1918 IP address space, which is not Internet routable, to further separate the IP Telephony networks. NAT should be used judiciously and only where required by call center applications, SoftPhone, or WebAttendant.

  The Internet gateway router and firewall should not allow NAT translations for Internet-to-IP Telephony connectivity. Keep in mind that implementing SoftPhone functionality places voice features on computers residing on the data network which impacts the separation of voice and data into distinct networks.
- IP filters or firewall features should be used on the gateway router between the IP Telephony network and the organization’s data network to eliminate well-known, malicious attacks that may originate from within the organization’s network.
- Firewalls should be used at the organization’s Internet connection, partner connections, and in front of the CallManager cluster.
- Intrusion detection tools should be placed at strategic points in the network to monitor for attacks.

The above mentioned rules are detailed below.

Creating and Assigning VLANs and Broadcast Domains

The vast majority of IP security solutions can only be implemented if a packet encounters a Layer 3 (IP) device. Due to protocol architecture, the MAC layer, or Layer 2, offers very little inherent security. Because of this, understanding and establishing broadcast domains is one of the fundamental precepts in designing secure IP networks. Many simple, yet dangerous attacks can be launched if the attacking device resides within the same broadcast domain as the target system.

To mitigate this inherent vulnerability of the IP protocol, the following potential target end-points should always be on their own subnets, separate from the rest of the data network and each other:

- Servers
- CallManager cluster
- IP telephones
- VoIP gateways
- Network management workstations

Additionally, every device should use a separate switched segment to connect to the network. The reason for using separate segments, in other words, a switched Ethernet infrastructure, is to prevent attackers or attacking applications from *snooping* or capturing Ethernet traffic traversing the physical wire. By ensuring that each device connects to the network using a switched infrastructure, packet-sniffing tools are rendered ineffective for capturing other users' traffic. Additionally, the recommended Cisco IP Telephony design model uses separate subnets for the IP phone and attached data PC by using 802.1Q VLAN trunking technology.

Figure 6-26 shows each of the major components that comprise an example enterprise network.

Note All IP Telephony devices reside on various subnets and VLANs in the voice IP network, 10.x.x.x, and all data pieces, such as PCs, e-mail servers, the DHCP server, etcetera, reside on the data IP network, 172.21.x.x. Additionally, this is a 100 percent switched Ethernet environment with every user and device residing on a separate segment.

Figure 6-26 Major Components of an Enterprise Network

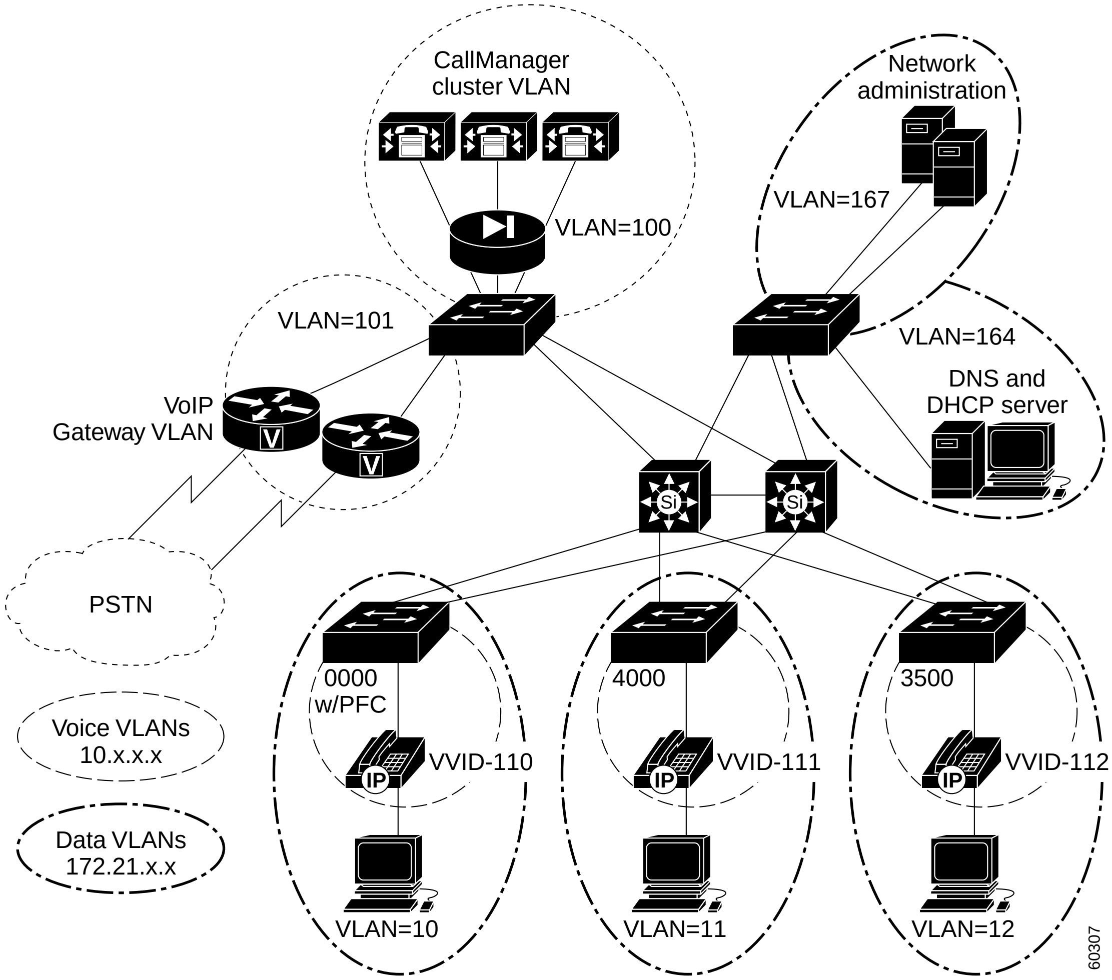

Note Assigning static IP addresses or using a separate, IP Telephony-specific DHCP server, located within the voice network, is a more secure solution for IP telephone IP address management than using the organization's existing DHCP services.

IP Addressing and NAT

Once VLAN assignment is complete, the actual IP addressing takes place. Use of RFC 1918 public address space is recommended for the IP Telephony network.

- IP Telephony deployment is much easier without having to re-IP and re-subnet the existing data network.
- An easy, logical separation of the IP Telephony and data networks is established through the use of a different address space.
- While Network Address Translation (NAT) does not guarantee security, when properly used, it is an additional tool for securing the IP Telephony network.

By using well-conceived IP addressing schemes, which include the use of RFC 1918 addresses and NAT, the IP Telephony network can be separated from the Internet without modifications to the Internet gateway router or Internet firewall.

Figure 6-27 shows the logical IP address separation using NAT.

Figure 6-27 Logical IP Address Separation Using NAT

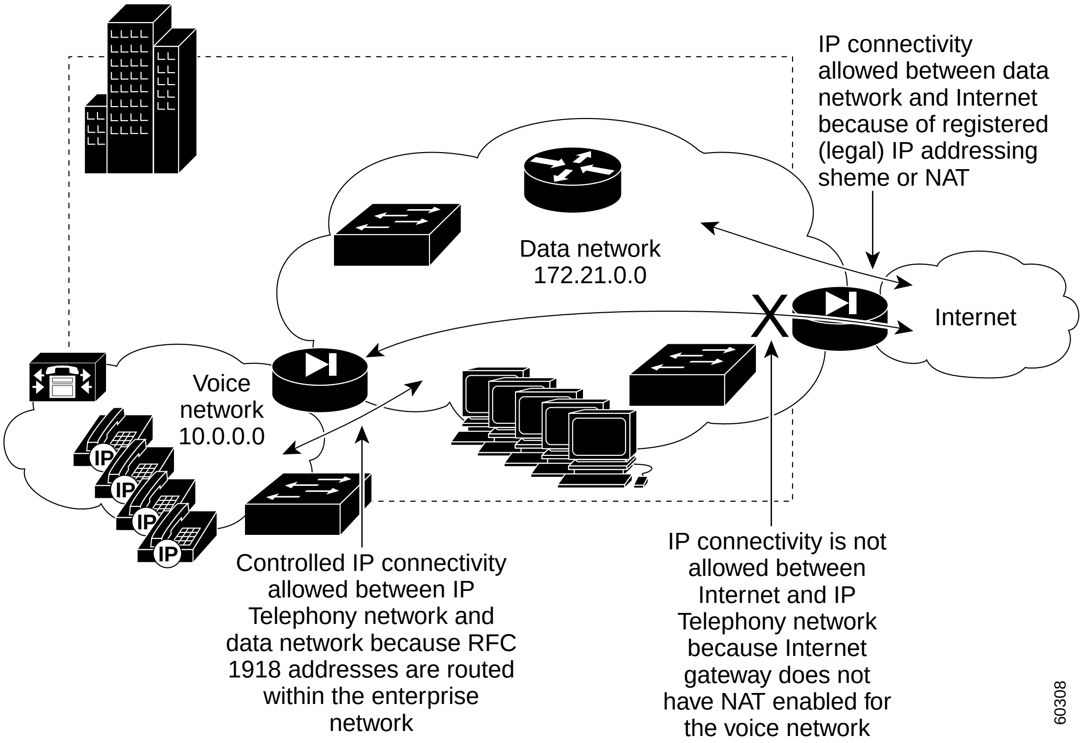
IP connectivity allowed between data network and Internet because of registered (legal) IP addressing sheme or NAT
Data network 172.21.0.0
Internet
Voice network 10.0.0.0
IP
IP
IP
IP
X
Controlled IP connectivity allowed between IP Telephony network and data network because RFC 1918 addresses are routed within the enterprise network
IP connectivity is not allowed between Internet and IP Telephony network because Internet gateway does not have NAT enabled for the voice network
60308

Implementing Packet Filters or Firewall Features

Using IP filters is an integral part of building secure networks. It is highly recommended to use them when securing the IP Telephony network. In most organizations, these filters should be placed on the router or firewall separating the IP Telephony and data networks.

Figure 6-28 shows the placement of IP filters.

Figure 6-28 IP Filters Placement

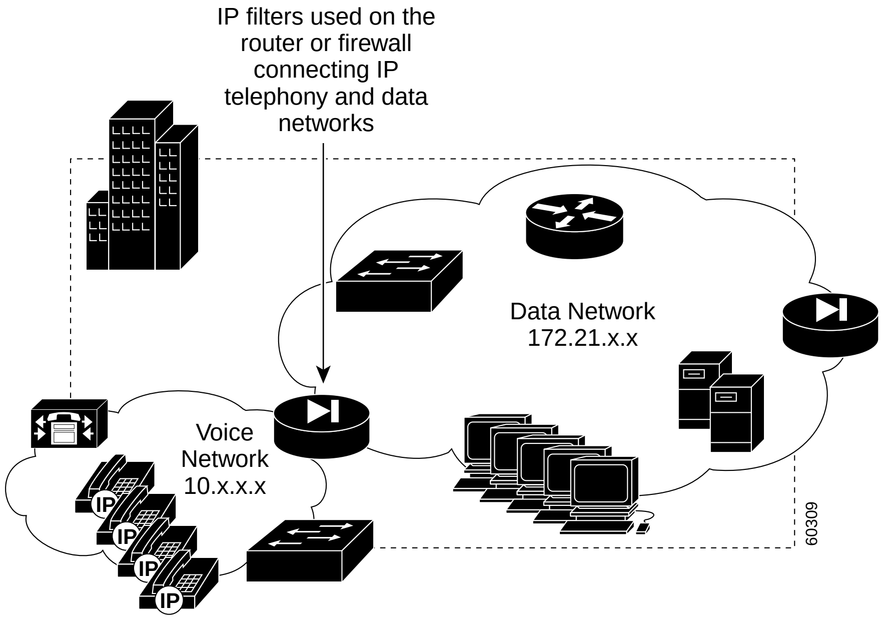
IP filters used on the router or firewall connecting IP telephony and data networks
Data Network
172.21.x.x
Voice
Network
10.x.x.x
IP
IP
IP
IP
60309

Directed Broadcasts

IP directed broadcasts are used in the popular *smurf* denial of service attack and its derivatives. An IP directed broadcast is a datagram which is sent to the broadcast address of a subnet to which the sending machine is not directly attached. The directed broadcast is routed through the network as a unicast packet until it arrives at the target subnet, where it is converted into a link-layer broadcast. Because of the nature of the IP addressing architecture, only the last router in the chain, the one that is connected directly to the target subnet, can conclusively identify a directed broadcast. Directed broadcasts are occasionally used for legitimate purposes, but such use is not common outside the financial services industry.

In a *smurf* attack, the attacker sends Internet Control Message Protocol (ICMP) echo requests from a spoofed legitimate source address to a directed broadcast address, causing all the hosts on the target subnet to send replies to the spoofed source. By sending a continuous stream of such requests, the attacker can create a much larger stream of replies, which can completely inundate the legitimate host whose address is being spoofed.

If a Cisco router interface is configured with the **no ip directed-broadcast** command, directed broadcasts that would otherwise be *exploded* into link-layer broadcasts at that interface are dropped instead. Note that this means that the **no ip directed-broadcast** command must be configured on every interface of every router that might be connected to a target subnet; it is not sufficient to configure only firewall routers. The **no ip directed-broadcast** command is the default in Cisco IOS software version 12.0 and later. In earlier versions, the command should be applied to every LAN interface.

Source-Routed Packet

The IP protocol supports source routing options that allow the sender of an IP packet to control the route that packet will take toward its ultimate destination, and generally the route that a reply will take. These options are rarely used for legitimate purposes in real networks and can be used for attacking hosts on an enterprise network. A Cisco router with the **no ip source-route** command set will never forward an IP packet that carries a source routing option. This command should be used both on the Internet connected router as well as the gateway router connecting the IP Telephony and data networks within the enterprise.

IP Spoofing

Many widespread Denial Of Service attacks rely on the ability of the attacker to send packets with forged (spoofed) source addresses, which makes tracking the true source of the attack very difficult. To help prevent your site from sourcing these types of attacks, you should block any outbound packets outside of your own address space.

Example 6-25 prevents any outbound packets that do not come from the site's imaginary address block of 172.21.0.0 as well as blocking any inbound packets that have a source address matching the internal network. It is also recommended that the service provider network implement RFC 2827 IP address blocking in order to help prevent *spoofed* traffic.

Example 6-25 Access List Permissions

```
access-list 101 permit ip 172.21.0.0 0.0.255.255 any
access-list 101 deny ip any any
access-list 102 deny ip 172.21.0.0 0.0.255.255 any
access-list 102 permit ip any any
interface Serial 2/1
   ip access-group 101 out
   ip access-group 102 in
```

ICMP Redirects

ICMP redirect messages instruct end nodes to use a specific router as its path to an IP destination. In a properly functioning IP network, a router sends redirects only to hosts on its own local subnets, no end node ever sends a redirect, and no redirect ever traverses more than one network hop. However, an attacker may violate these rules. In fact, some attacks are based on this.

To migrate this threat, filter all ICMP redirects on the border router between the IP Telephony network and the data network.

Example 6-26 ICMP Redirect Filtering

```
interface ethernet 0/1
   no ip redirects
```

Note ICMP redirect filtering only prevents redirect attacks initiated by attackers across at least one network hop. If the attacker is connected to the same subnet, it's still possible to launch this attack. This is another reason to ensure that all data devices such as user workstations and data servers are always on a separate network from all voice devices and endpoints.

Permitting IP Telephony and Other Services

Organizations with the most stringent security policies will prohibit any connections between the voice and data networks: This greatly reduces the threat of attacks. However, many organizations will require at least minimal communication between the data network and the IP Telephony network for policy, application, or cost reasons. For example, some organizations may not wish to use a separate DHCP server for voice and data devices because of the increased costs and management complications associated with a second server.

Table 6-25 lists applications and ports that may, need to be *opened* between the data and telephony networks.

Table 6-25 Open Applications and Ports Between Data and Telephony Networks

Application	Port	Direction	Requirement
Routing protocol	Protocol dependent	Direction—Both ways	IP routing
DHCP	UDP 67, 68	Source—Voice network Destination—Data network DHCP server Direction—Both ways	Single enterprise-wide DHCP used for PCs and IP phones. A more secure alternative is to use static IP addresses for IP phones or a separate IP Telephony-specific DHCP server.
ICMP	IP ICMP echo	Source/Destination—Network Admin subnet and voice network Direction—Both ways	Basic troubleshooting. Only allow all Network Admin workstations this privilege.
NTP	TCP 123 UDP 123	Source—Routers, switches, gateways, CallManagers, and management servers Destination—Trusted enterprise NTP server Direction—One way	Synchronized timestamps on all network elements for troubleshooting and tracking.
Telnet	TCP 23	Source—Network Admin subnet Destination—Voice network Direction—One way	Configuration and troubleshooting. SSH should be preferred over Telnet. Only allow Network Admin workstations this privilege.
SSH	TCP 22	Source—Network Admin subnet Destination—All routers on all subnets Direction—One way	Configuration and troubleshooting. SSH provides a more secure method of administering network elements.
FTP	TCP 20, 21	Source—CallManager *publisher* and TFTP server on the voice network Destination—Data network FTP server (behind the organization’s Internet firewall) Direction—One way	When downloading new versions of IP Telephony software from Cisco Connection Online, it is recommended to use a data network FTP server to download the code and then access the software through the additional step of accessing FTP though the IP Telephony firewall.
RADIUS	UDP 1645, 1646, 1812, 1813	Source—Voice network Destination—Organization’s RADIUS server Direction—One way	Router, switch, and VoIP gateway access configuration. Ports depend on whether it's Cisco IOS (1645-46) or CatOS (1812-13).
TACACS+	TCP 49	Source—IP Telephony network Destination—Organization’s TACACS+ server Direction—One way	Router, switch, and VoIP gateway access configuration.

Table 6-25 Open Applications and Ports Between Data and Telephony Networks (continued)

Application	Port	Direction	Requirement
DNS	UDP 53	Source—Voice network Destination—DNS servers Direction—One way	CCM server lookups, TFTP server lookups, FTP server lookups, IP phones, LDAP, gateways, IP Telephony network management servers
LDAP	TCP 389	Source—Voice network Destination—LDAP servers Direction—One way	LDAP functionality
SoftPhone	TAPI/JTAPI = TCP 2748 VoIP media stream = UDP 16384-32767	Source/Destination—Data network and CallManagers Direction—Both ways	SoftPhone residing on user PCs. SoftPhone functionality places voice features on computers residing on the data network which impacts the separation of voice and data into distinct networks.
SoftPhone CCM Directory Lookup	TCP 8404	Source—SoftPhone clients Destination—TCP port 8404 on the CallManager Direction—One way	Using CCM/SoftPhone directory services. **Note** SoftPhone also has an LDAP client for querying the organization's LDAP service using TCP 389.
IP IVR	TAPI/JTAPI = TCP 2748 VoIP media stream = UDP 16384-32767	Source—IP IVR server Destination—Voice network Direction—Both ways	Only required if IP IVR server located on the data network instead of the IP Telephony network. Placing this service on the data network impacts the separation of voice and data into distinct networks.
IPCC	TAPI/JTAPI = TCP 2748	Source—GeoTel IPCC Server Destination—CallManagers Direction—Both ways	Only required if IPCC server is located on the data network instead of the IP Telephony network. Placing this service on the data network impacts the separation of voice and data into distinct networks.
HTTP	TCP 80	Source—Network Admin subnet and potentially all user workstations Destination—CallManagers Direction—One way	Web access to the Voice network.
Skinny Client	TCP 2000	Source—CallManagers Destination—Voice gateway routers Direction—One way	Call setup and control.
Skinny Gateway (analog)	TCP 2001	Source—CallManagers Destination—Voice gateway routers Direction—One way	Call setup and control

Table 6-25 Open Applications and Ports Between Data and Telephony Networks (continued)

Application	Port	Direction	Requirement
Skinny Gateway (digital)	TCP 2002	Source—CallManagers Destination—Voice gateway routers Direction—One way	Call setup and control
H.323 RAS	TCP 1719	Source—CallManagers Destination—Voice gateway routers Direction—One way	Call setup and control
H.323 (H.225)	TCP 1720	Source—CallManagers Destination—Voice gateway routers Direction—One way	Call setup and control
H.323 (H.245)	TCP 11000 to 11999	Source—CallManagers Destination—Voice gateway routers Direction—One way	Call setup and control
MGCP	UDP 2427 and TCP 2428	Source—CallManagers Destination—Voice gateway routers Direction—One way	Call setup and control
SNMP	UDP 161	Source—Network management devices Destination—SNMP-managed devices Direction—One way	SNMP network management
SNMP TRAP	UDP 162	Source—SNMP-managed devices Destination—Network management devices Direction—One way	SNMP network management

Example 6-27 shows a router as the gateway between the voice and data networks.

Example 6-27 Cisco IOS Firewall Configuration

```
ip inspect audit-trail
ip inspect max-incomplete low 150
ip inspect max-incomplete high 250
ip inspect one-minute low 100
ip inspect one-minute high 200
ip inspect udp idle-time 20
ip inspect dns-timeout 3
ip inspect tcp idle-time 1800
ip inspect tcp finwait-time 3
ip inspect tcp synwait-time 15
ip inspect tcp max-incomplete host 40 block-time 0
ip inspect name avvid_firewall tcp timeout 300
ip inspect name avvid_firewall udp
ip inspect name avvid_firewall tftp
ip inspect name avvid_firewall http
```

```
interface FastEthernet0/1
    description This interface connects to the Data Network
    ip address 172.21.100.1 255.255.255.0
    ip access-group secure_avvid in
        ! Controlled access from the Data Network to the Voice Network
    no ip directed-broadcast
    no ip redirects
    ip inspect avvid_firewall out
        ! Allows traffic originating from the voice network to access the data network

ip access-list extended secure_avvid
    permit eigrp any any
        ! Allow routing protocol access to the voice network

    permit tcp 172.21.167.0 0.0.0.255 10.21.100.0 0.0.0.255 eq www
    permit tcp 172.21.167.0 0.0.0.255 10.0.0.0 0.255.255.255 eq telnet
    permit tcp 172.21.167.0 0.0.0.255 10.0.0.0 0.255.255.255 eq 22
    permit icmp 172.21.167.0 0.0.0.255 10.0.0.0 0.255.255.255 echo-reply log
    permit icmp 172.21.167.0 0.0.0.255 10.0.0.0 0.255.255.255 echo log
    permit udp 172.21.167.0 0.0.0.255 10.0.0.0 0.255.255.255 snmp
        ! Allow Network Admin subnet access to the Voice network and CallManager Cluster
        subnet

    permit udp host 172.21.164.40 10.0.0.0 0.255.255.255 eq 67
    permit udp host 172.21.164.40 10.0.0.0 0.255.255.255 eq 68
    ! 172.21.164.40 is the dhcp server – only needed if data network provides dhcp service
    to voice

    permit tcp 172.21.0.0 0.0.255.255 10.21.100.0 0.0.0.255 eq 2748
    ! JTAPI for Softphones, IP-IVR, and IPCC to the CallManager subnet

    permit udp 172.21.0.0 0.0.255.255 10.0.0.0 0.255.255.255 range 16384 32767
    ! RTP for SoftPhone and IP-IVR to communicate to IP Phones and VoIP gateway

    permit tcp 172.21.0.0 0.0.255.255 10.21.100.0 0.0.0.255 eq 8404
    ! SoftPhone Directory service to CallManager cluster

permit tcp 172.21.0.0 0.0.255.255 host 10.21.101.11 eq 389
    ! SoftPhone access to the LDAP server

deny ip any any log
    ! explicit deny all with logging
```

Protecting the VoIP Gateways

The VoIP gateways should only accept call setup attempts from CallManager servers. This can be accomplished by configuring an access list on the VoIP gateway. Example 6-28 shows a VoIP gateway that only accepts call setup attempts from the CallManager cluster subnet.

Note In this example, the VoIP gateway resides on the voice network.

Example 6-28 VoIP Gateway Router Interface and ACL

```
interface fastethernet 1/0
    description VoIP (H.323 & MGCP) Gateway
    ip address 10.21.101.1 255.255.255.0
    ip access-group secure-gw in
    no ip directed-broadcast
    no ip redirects
```

```
ip access-list extended secure-gw
   permit eigrp any any
      ! Allow routing protocol access to the VoIP Gateway

   permit udp host 172.21.164.40 eq domain host 10.21.101.1
      ! Allow DNS lookups to the VoIP Gateway

   permit tcp 172.21.167.0 0.0.0.255 eq 123 host 10.21.101.1
   permit udp 172.21.167.0 0.0.0.255 eq 123 host 10.21.101.1
   permit tcp 172.21.167.0 0.0.0.255 host 10.21.101.1 eq telnet
   permit tcp 172.21.167.0 0.0.0.255 host 10.21.101.1 eq 22
   permit icmp 172.21.167.0 0.0.0.255 host 10.21.101.1  echo-reply log
   permit icmp 172.21.167.0 0.0.0.255 host 10.21.101.1  echo log
   permit udp 172.21.167.0 0.0.0.255 host 10.21.101.1  snmp
      ! Allow Network Admin subnet access to the VoIP gateway

   permit ip 10.21.100.0 0.0.0.255 host 10.21.101.1
      ! Allow CallManager Cluster Subnet full access to the VoIP gateway

   permit udp 10.0.0.0 0.255.255.255 any range 16384 32767
      ! Allow IP Phones access to the VoIP gateway and beyond

   permit udp 172.21.0.0 0.0.255.255 any range 16384 32767
      ! Allow SoftPhones access to the VoIP gateway and beyond

   deny ip any any log
      ! explicit deny all with logging
```

Firewalls

IT departments consider Internet firewalls to be a required piece of the network security infrastructure. Since the purpose of this section is IP Telephony security, not Internet security, it is assumed that an Internet security policy and architecture has already been established using secure network design principles such as SAFE. Sound security policies dictate that external partner connections require additional firewall measures. Once these are in place, and the IP Telephony network is built and connected to the existing IP data network, another firewall should be added between the CallManager cluster and the rest of the organization’s networks.

Figure 6-29 shows the placement of a firewall between the CallManager cluster and the network.

Figure 6-29 Firewall Placement Between CallManager and Network

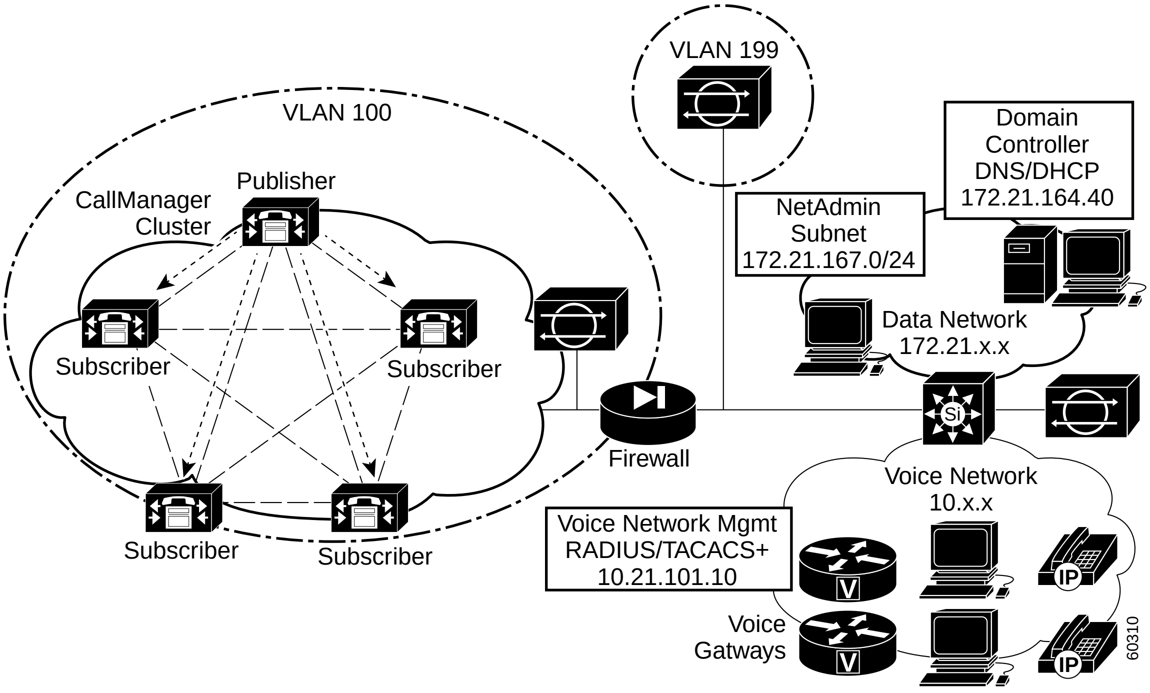

The firewall can either be a router running Cisco Secure Integrated Software (Cisco IOS Firewall), a PIX firewall or a third-party firewall. The primary criteria for choosing a firewall should be as follows:

- The strength of its security features
- The speed at which it handles connections
- High-availability
- How it integrates with the existing network infrastructure

By placing a firewall between the CallManager cluster and the voice and data networks, network designers greatly reduce the exposure of the most critical component in the IP Telephony network—the call processing intelligence. The firewall acts as a trusted proxy between all IP devices and the CallManagers, ensuring that only authorized transactions are allowed.

This firewall sits between the voice network and the CallManager cluster subnet. Please note that Network Address Translation (NAT) must not be used on this firewall and since the firewall is unable to pass EIGRP traffic: Routes to and from the CallManager cluster must be statically configured and redistributed as necessary.

Example 6-29 contains a list of transactions that could be allowed through the firewall to the CallManager cluster.

Example 6-29 PIX Firewall Software Version 5.X Configuration

```
access-list avvid_in permit udp 10.0.0.0 255.0.0.0 10.21.100.0 255.255.255.0 eq tftp
    ! Allow TFTP from the Voice Network to the CallManager Cluster Subnet

access-list avvid_in permit tcp 10.0.0.0 255.0.0.0 10.21.100.0 255.255.255.0 eq 2000
access-list avvid_in permit tcp 10.0.0.0 255.0.0.0 10.21.100.0 255.255.255.0 eq 2001
access-list avvid_in permit tcp 10.0.0.0 255.0.0.0 10.21.100.0 255.255.255.0 eq 2002
    ! Allow Skinny from the Voice Network to the CallManager Cluster Subnet
```

```
access-list avvid_in permit tcp 10.0.0.0 255.0.0.0 10.21.100.0 255.255.255.0 eq 1719
access-list avvid_in permit tcp 10.0.0.0 255.0.0.0 10.21.100.0 255.255.255.0 eq 1720
access-list avvid_in permit tcp 10.0.0.0 255.0.0.0 10.21.100.0 255.255.255.0 range 11000
11999
       ! H.323 access from the Voice Network to the CallManager Cluster Subnet

access-list avvid_in permit udp 10.0.0.0 255.0.0.0 10.21.100.0 255.255.255.0 eq 2427
access-list avvid_in permit tcp 10.0.0.0 255.0.0.0 10.21.100.0 255.255.255.0 eq 2428
    ! MGCP from the Voice Network to the CallManager Cluster Subnet

access-list avvid_in permit tcp 172.21.0.0 255.255.0.0 10.21.100.0 255.255.255.0 eq 2748
    ! CTI (TAPI and JTAPI) for SoftPhone to the CallManager Cluster Subnet

access-list avvid_in permit tcp 172.21.0.0 255.255.0.0 10.21.100.0 255.255.255.0 eq 8404
    ! SoftPhone Directory to the CallManager Cluster Subnet

access-list avvid_in permit tcp 172.21.167.0 255.255.255.0 10.21.100.0 255.255.255.0 eq
www
access-list avvid_in permit tcp 172.21.167.0 255.255.255.0 10.21.100.0 255.255.255.0 eq
telnet
access-list avvid_in permit tcp 172.21.167.0 255.255.255.0 10.21.100.0 255.255.255.0 eq 22
access-list avvid_in permit icmp 172.21.167.0 255.255.255.0 10.21.100.0 255.255.255.0 0
access-list avvid_in permit icmp 172.21.167.0 255.255.255.0 10.21.100.0 255.255.255.0 8
access-list avvid_in permit udp 172.21.167.0 255.255.255.0 10.21.100.0 255.255.255.0 eq
snmp
    ! Allow Network Admin subnet access to CallManager Cluster subnet

access-list avvid_in permit tcp 172.21.0.0 255.255.0.0 10.21.100.0 255.255.255.0 eq www
    ! SoftPhone Telecaster HTTP access to the CallManager Cluster Subnet

ip address outside 10.21.199.2 255.255.255.0
    ! Interface attached to the Voice Network

ip address inside 10.21.100.1 255.255.255.0
    ! Interface attached to the CallManager Cluster

static (inside,outside) 10.21.100.0 10.21.100.0 netmask 255.255.255.0
    ! Do not NAT the CallManager Cluster address across the firewall

access-group avvid_in in interface outside
    ! Apply the access-list to the outside interface of the firewall
```

Intrusion Detection Systems

The last step in designing the network for secure IP Telephony is adding network Intrusion Detection Systems (IDS). IDS examine IP traffic going to strategic servers or networks looking for *attack* profiles or signatures. Once the system detects traffic streams that match the characteristics of an attack, it can either alarm the IDS management station or send configuration calls to a Cisco router to block the attack. Host based intrusion detection should be considered for critical servers including the CallManagers.

Figure 6-30 shows a network intrusion detection system.

Figure 6-30 Network Intrusion Detection System

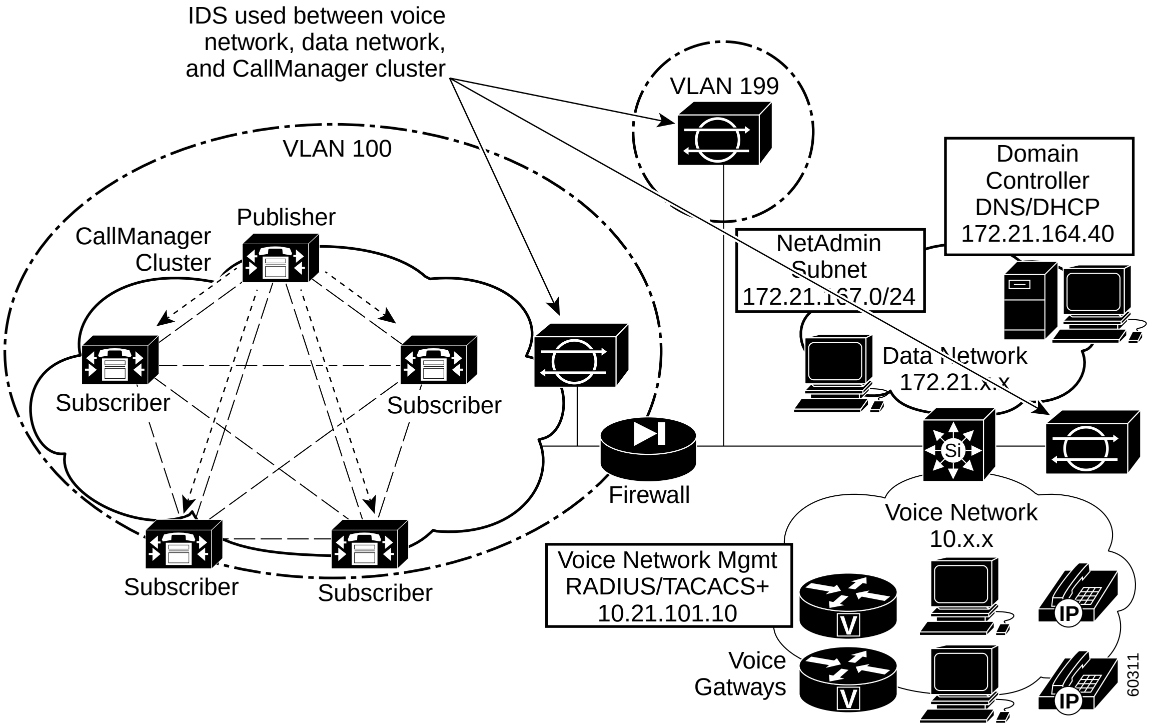
IDS used between voice network, data network, and CallManager cluster
VLAN 199
VLAN 100
Domain Controller DNS/DHCP 172.21.164.40
NetAdmin Subnet 172.21.167.0/24
Publisher
CallManager Cluster
Subscriber
Subscriber
Subscriber
Subscriber
Data Network 172.21.x.x
Si
Firewall
Voice Network 10.x.x
Voice Network Mgmt RADIUS/TACACS+ 10.21.101.10
V
V
IP
IP
Voice Gatways
60311

IDS systems should be placed on all Internet connections as part of the existing security policy. Additional IDS systems should be considered for partner connections, as well. When the IP Telephony network is built, additional IDS systems, situated between the organization's data network and IP Telephony network and on the inside firewall protecting the CallManager cluster, are also recommended.

The IDS appliance or Catalyst module monitoring traffic between the IP Telephony network and the organization's data network will log all suspicious flows with both the syslog server and the IDS management system, providing an audit trail for network managers. The IDS appliance or Catalyst module inside the PIX or Cisco IOS firewall protecting the CallManager cluster will be configured to block attacks destined for the call processing infrastructure. In other words, traffic that is diagnosed as malicious in intent will be *shunned* through the intrusion detection system's near real-time ability to modify the routers connecting the CallManager cluster and IP Telephony network with the organization's data network. The IDS identifies attacks based upon its signature database.

Securing the CallManager Server

Once the network design portion of securing the IP Telephony network is complete, it is time to address the Cisco CallManager servers. The CallManager runs on a Compaq server using the Microsoft Windows 2000 server operating system. Windows 2000 is inherently more secure than any previous Microsoft OS because of the default file permission settings, improved TCP/IP stack, and configuration granularity.

However, as is the nature of all operating systems, a constant vigil must be kept by system administrators to ensure that newly discovered vulnerabilities are quickly eliminated. This ongoing system modification can range from changing the default installation procedures of new software to modifying the kernel through OS patches.

To configure administrator security enhancements on the CallManager, perfom the following steps.

Step 1 Update the Windows 2000 operating system with the latest patches.

Step 2 Turn off unused or unnecessary services and applications running on the CallManager.

Step 3 Apply security settings to the NTFS file structure.

Step 4 Enable system security auditing.

Step 5 Secure IIS settings.

Step 6 Secure the SQL server installation.

Step 7 Remove voice conferencing/MTP software packages installed on the CallManager.

Step 8 Configure CallManager SNMP securely.

Step 9 Shut down the TFTP server running on the CallManager and install a separate TFTP server for the IP telephones and gateways.

Step 10 Move all IP phone Web Services to a separate server.

Each of the above steps is detailed in this section. It is highly recommended to routinely visit the Microsoft security site at http://www.microsoft.com/security for new information to keep the CallManager and other servers and workstations using Microsoft operating systems secure. This should be part of the enterprise security policy and done on a routine basis.

Many of the current security and OS patches have already been added by the operating system installation CD-ROM that comes with CallManager. Since new patches are continually released, it is important to make sure that all of the steps detailed here are followed to ensure a secure system.

Patches

Adding patches to an operating system, and installing and configuring the CallManager server, at the time of implementation is standard procedure for system administrators. As soon as the initial installation is complete, the administrator should verify that Microsoft recommended security patches have been installed and that CallManager functionality has not been impacted. It is important to visit the following website for Microsoft updates: http://www.microsoft.com/technet/security/current.asp.

As new security patches are issued, they should be installed on the CallManager once it has been verified that they will not impact the CallManager’s functionality in the production environment.

Turning Off Unnecessary Services

When the Windows 2000 operating system installs, many applications and services not needed by the CallManager are enabled by default. These services can be disabled so that potential vulnerabilities associated with them can be eliminated.

One of the fundamental principles of securing a server is that each running service exposes potential security vulnerabilities. Since securing every running service is difficult and time-consuming, it is a logical task to disable all services that are not mandatory, even if those services aren't immediately known to have a vulnerability.

Unless otherwise needed on the system, the following services should be stopped and set to **Manual**:

- Alerter Service
- Clipbook Server

- Computer Browser
- Distributed File System
- DHCP Client
- Messenger
- Net Logon
- Network DDE and DDE DSDM
- Network Monitor Agent
- Spooler
- SMTP Service
- NNTP Service
- DHCP Service
- DNS Server Service
- FTP Publishing Service
- Fax Service
- Net Meeting Remote Desktop Sharing

On subscriber CallManagers, the following services should be disabled in addition to those listed above:

- IIS Admin Service
- World Wide Web Publishing Service

All of the web administration takes place on the *publisher*, so there is no reason to keep these services running on the subscribers. The publisher database has read and write access, while all subscribers have only read access.

Securing File System Access

Unknown users should not be able to log into the CallManager system. By default, the Windows 2000 file system is not very secure. However, sanitizing the different user and system accounts and applying the proper permissions to the file system structure can secure the system reasonably well.

Besides stopping unneeded services and applications, there are two key pieces to securing access to the system. First, the NTFS permissions on the file system itself need to be secured. As a general rule, NTFS permissions are cumulative. If someone is a member of two groups that have access to a directory, they will have the higher access of the two groups. The exception is for explicit *Deny* access settings. If there are two groups assigned to a directory, and one group is explicitly denied, anyone in both groups will be denied since an explicit denial overrides everything.

The second piece, which will be covered later, relates to securing the method of access. Accessing the file system via IIS is similar to accessing a file remotely through Windows Explorer. A *share* is set up that allows someone to access a resource. IIS is just another means of *sharing* a series of files or resources. When someone attempts to access a resource through a share, the access they have is the cumulative access of any groups given access to that share.

However, when someone accesses a NTFS secured resource through a share, the most restrictive access applies. If someone has Administrative privileges via IIS, but they only have read access on the file system itself, their total access level is read only.

First, the accounts themselves need to be modified. Disable the Guest account and remove any users from the Guests group. You can perform this modification from the **System Tools > Local Users and Groups** subdirectory, and then selecting **Start > Programs > Administrative Tools > Computer Management**.

All accounts except for the administrator account will become locked if too many incorrect passwords are attempted. The Administrator account never locks up, so many attacks rely on blindly trying to execute commands as Administrator. By changing the name of the Administrator account to CallmgrAdmin or some other logical name, these types of attacks can be mitigated.

Figure 6-31 shows an example of the Local Users and Groups directory.

Figure 6-31 Computer Management Window

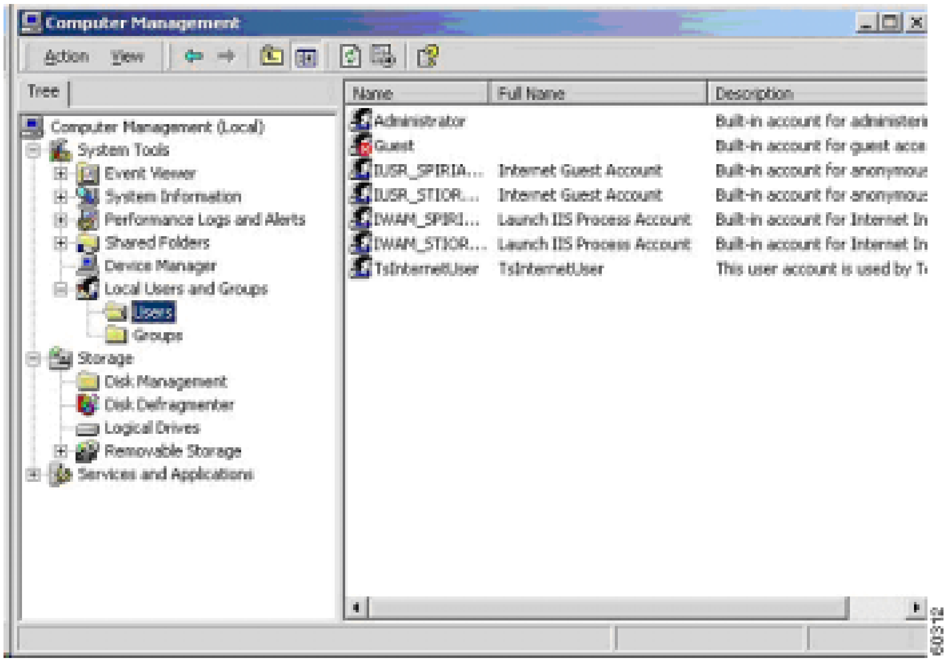

Name	Full Name	Description
Administrator		Built-in account for administeri
Guest		Built-in account for guest acce
IUSR_SPIRIA...	Internet Guest Account	Built-in account for anonymous
IUSR_STIOR...	Internet Guest Account	Built-in account for anonymous
IWAM_SPIRI...	Launch IIS Process Account	Built-in account for Internet In
IWAM_STIOR...	Launch IIS Process Account	Built-in account for Internet In
TsInternetUser	TsInternetUser	This user account is used by T

Another recommended step in Windows 2000 account sanitation is to secure all privileges of the Everyone group. The default of the Everyone group has access to every file. This account must be removed from the root file system. However, setting Everyone to No Access will prevent all users, including the administrator, from accessing the system. To remove the Everyone group from the file system, perform the following steps.

Step 1 Right-click the c: drive in Windows Explorer.

Step 2 Go to Properties, and click the **Security** tab.

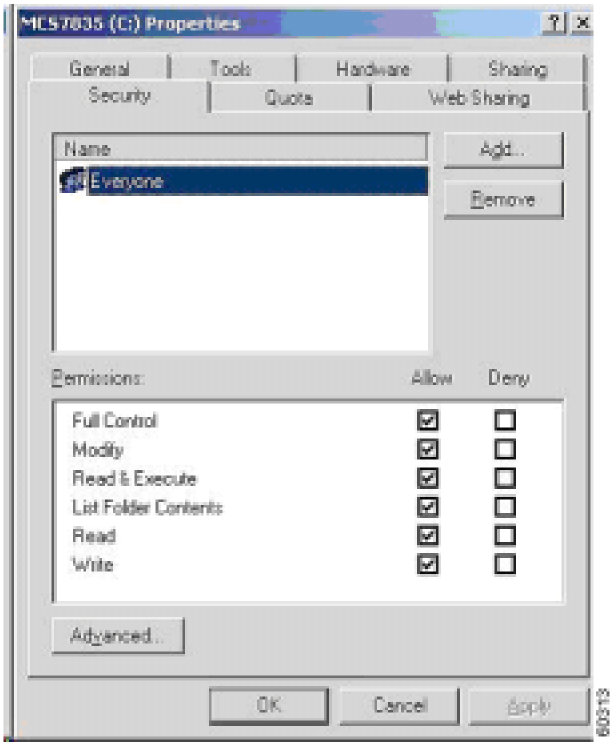

Name

Everyone

Permissions:	Allow	Deny
Full Control	☑	□
Modify	☑	□
Read & Execute	☑	□
List Folder Contents	☑	□
Read	☑	□
Write	☑	□

Step 3 Add the Administrator group.

Step 4 Verify that all permissions are granted full access.

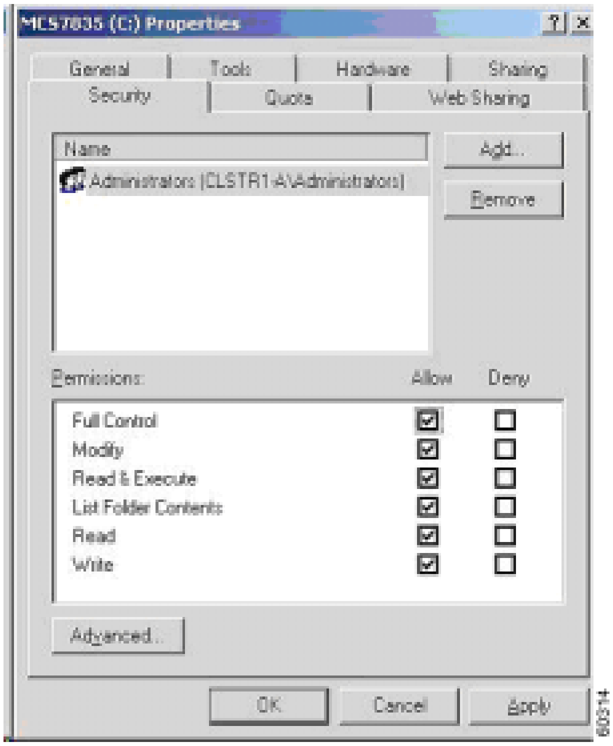

Name

Administrators (CLSTR1-A\Administrators)

Permissions:	Allow	Deny
Full Control	☑	□
Modify	☑	□
Read & Execute	☑	□
List Folder Contents	☑	□
Read	☑	□
Write	☑	□

Step 5 Remove the Everyone group.

System Auditing and Logging

Auditing allows usage tracking for many privileged tasks in Windows 2000. When auditing is enabled, regularly reviewing the Event Viewer may help determine if the system is under attack or has been compromised. Table 6-26 shows the suggested auditing scheme.

Table 6-26 Auditing Scheme

Description	Log Access	Log Failure
Audit Account Login Events	Yes	Yes
Audit Account Management	Yes	Yes
Audit Directory Service Access	Yes	Yes
Audit Login Events	No	Yes
Audit Object Events	Yes	Yes
Audit Policy Change	Yes	Yes
Audit Privilege Use	Yes	Yes
Audit Process Tracking	No	Yes
Audit System Events	Yes	Yes

SQL Server 7.0 provides a very powerful profiler, which allows the analysis of many internal events that occur within SQL Server. SQL Server Profiler works by capturing all the actions performed on the SQL Server and sending them to the SQL Server Profiler, where they can be analyzed. The capture can be viewed by the following methods:

- In real-time on the screen
- Saved to a text file
- Inserted into a SQL Server table

The SQL Server Profiler allows capturing of virtually all events that take place within SQL Server, including:

- Login Failed
- Locking: Deadlock
- Object: Closed
- Stored Procedure: Statement Starting
- Session: Disconnect
- RPC: Completed

This information can provide excellent support to establish event time and origin.

Securing IIS

IIS is a major component in the CallManager server. All administration of the CallManager flows through this service. The difficulty lies in the fact that IIS has been a target of several well-known attacks in the hacking community. Because of this, several steps must be taken to secure IIS.

Enable W3C Extended Logging Format

The default logging mechanism does not record enough information to help determine whether a server is under attack.

Clear Indexing

By indexing source code, it is possible for an attacker to view the content of the web pages. To clear indexing, perform the following steps.

Step 1 Start the IIS Microsoft Management Console (MMC) and go to the Web Site Properties by right-clicking the website entry and selecting **Properties**.

Step 2 Click the **Home Directory** tab.

Step 3 Clear the **Index this Directory** and the **Directory Browsing Allowed** options.

Remove unused script mappings

IIS is preconfigured to support various common filename extensions such as .ASP, .SHTML, and .HTR. Processing of these requests are handled by various DLLs located on the system. By removing the mappings to extensions that are not used, you minimize the potential attack points. To remove the mappings to extensions, perform the following steps.

Step 1 Start the IIS MMC and go to the Web Site Properties by right-clicking on the Web Site entry and selecting **Properties**.

Step 2 Click the **Home Directory** tab.

Step 3 Click the **Configuration** tab.

Step 4 Click the **App Mappings** tab.

Step 5 Remove the necessary mappings.

Removing IIS Virtual Directories

IIS contains the following virtual directories that need to be removed:

- IISAMPWD
- IISSAMPLES
- IISADMIN
- IISHELP

Removing All Sample Application Directories

The following directories contain sample files that you should remove from the system. This will prevent an attacker from exploiting vulnerabilities in one of the sample files to gain access to the system.

- \Inetpub\iisamples
- \Inetpub\scripts\samples
- \Inetpub\wwroot\samples
- \Program Files\Common Files\System\msadc\Samples
- \WINNT\system32\inetsrv\adminsamples
- \WINNT\system32\inetsrv\iisadmin
- \WINNT\system32\inetsrv\iisadminpwd

Setting Appropriate Virtual Directory Permissions for Web Application Space

It is important that you apply the correct permissions to the files available on the web server. These permissions vary depending on the type of files being accessed. Table 6-27 provides a rough guideline to follow.

Table 6-27 Virtual Directory Permissions

File Type	File Extension	ACL
CGI and related files	.EXE, .DLL, .CMD, .PL	Everyone—Execute Administrators and System—Full Control
Script files	.ASP	Everyone—Execute Administrators and System—Full Control
Include files	.INC, .SHTML, .SHTM	Everyone—Execute Administrators and System—Full Control
Static content	.HTML, .GIF, .JPEG	Everyone—Execute Administrators and System—Full Control

Setting Appropriate IIS Log File ACLs

To prevent malicious users from deleting log files to cover their activities, the file permissions on the IIS generated log files (%systemroot%\system32\LogFiles) should be shown in Table 6-28:

Table 6-28 File Permission Setting

File Type	File Extension	ACL
Log files	.LOG	Everyone—Execute Administrators and System—Full Control

Installing Microsoft MDAC 2.1.2.4202.3

On websites that have both IIS and certain versions of MDAC, a visitor could perform privileged actions on the system. You can remove this vulnerability by installing MDAC 2.1 and configuring it to operate in Safe Mode. Change the following registry key as follows:

Hive: HKEY_LOACL_MACHINE\SOFTWARE

Key: \Microsoft\DataFactory\HandlerInfo

Securing the SQL Server

The SQL Server database is a key part of the CallManager. Several simple changes can be made to the SQL Server to setup to help tighten down overall call processing security. It is not recommended to use the *sa* account for daily administration, but rather only for emergencies. Once the *sa* password has been configured, put it in a safe and use an administrative group for all SQL Server administration and configuration.

Using a Separate Group for SQL Server Administration

Instead of using the *sa* account for administration and database configuration, it is recommended to use a Windows 2000 group for administrative privileges. The administrators are granted access to SQL Server through group-wide Windows 2000 permissions.

Setting the SQL Server to Run Under a Local System Account

SQL Server runs as three related processes under Windows 2000:

- MSSQLServer
- SQLServerAgent
- Microsoft Search

In order to match the rest of the CallManager User configuration, each of these services should run under Local System Accounts. Because Microsoft Search must use a Local System Account, MSSQLServer and SQLServerAgent can be configured to run under the same User as the Search Process. Once the SQL Server has been installed, use the SQL Server Enterprise Manager to change the accounts these processes use.

The following permissions must be set for the Local System Account for SQL Server 7.0 to perform its tasks properly:

- Full Control on the SQL Server directory (by default \MSSQL7)
- Full Control on all .mdf, .ndf, and .ldf database files
- Full Control on the registry keys at and under:

  Hive: HKEY_LOCAL_MACHINE\Software

  Key: \Microsoft\MSSQLServer

  Hive: HKEY_LOCAL_MACHINE\System

  Key: \CurrentControlset\Services\MSSQLServer

SQL Server Auditing

Using the SQL Server Enterprise Manager, set the server logging to **ALL**. The auditing information is written to the SQL Server 7.0 error log.

Uninstall Conferencing Software and Services

It is recommended not to install the software-based conferencing services on the CallManager server. The conferencing application terminates RTP/UDP VoIP streams and mixes them together to create a conference call. The risk is that UDP is an inherently insecure protocol and terminating it on the CallManager opens it to exposure of attacks unnecessarily. Using either hardware based conferencing or installing the conferencing software on a separate Windows 2000 server can mitigate this risk.

Configure CallManager SNMP

If the Cisco Works option is chosen during the initial CallManager installation, SNMP is enabled on the CallManager server. Several vulnerabilities are opened through starting the SNMP service on a Windows 2000 server. Because of this, if SNMP is not required, it is recommended to disable the SNMP service and configuration. However, if SNMP is used to manage the CallManager and IP Telephony network, perform the following steps to secure the system.

Step 1 Change the default communities.

By default, Windows 2000 installs the READ community as public. This must be changed to a unique community name. Both the READ and WRITE community strings should be treated as passwords. The same care should be given to choosing these strings as any root password or router login.

Step 2 Configure the IP Telephony network management workstation as the only host able to send and receive traps.

Only a network management server on the IP Telephony network should be allowed to send and receive SNMP TRAPs. It is highly recommended to separate SNMP management servers for the organization's data network and IP Telephony network so no SNMP requests/replies will traverse the firewall.

Using a Separate TFTP Server

TFTP services should not be run on the CallManager server. Because TFTP uses UDP for transferring data, inherent risks are associated with running it on strategic servers. For scalability and security, a separate TFTP server should be configured. Please note, if the TFTP server become unavailable, the IP telephones will simply revert to the configuration and firmware version stored in FLASH memory, so no interruption of the voice network is experienced.

Moving All IP Phone Web Services to a Separate Server

All web proxy and *asp* functionality that is used for services on the IP phones should be moved to a separate server. This secures the CallManager publisher from *proxy* type attacks. It also prevents these operations from consuming processing power on the CallManager.

Troubleshooting IP Telephony Networks

This section describes the tools and applications available for troubleshooting the IP Telephony network, and focuses on the following topics:

- Troubleshooting Tools
- Troubleshooting Cisco CallManager Devices
- Call Detail Records

Troubleshooting Tools

This section focuses on the following topics:

- Cisco CallManager Administration Details
- Microsoft Performance Monitor
- Microsoft Event Viewer
- CallManager Trace
- SDL Trace
- Sniffer Trace
- Call Detail Records (CDR) and Call Management Records (CMR)
- Q.931 Translator

Cisco CallManager Administration Details

Cisco CallManager Administration provides version information for the system, database, and other components. Figure 6-32 illustrates the associated screen that provides this information. On the opening window, click the **Details** button to view the versions in use.

Figure 6-32 Cisco CallManager Administration Version Information Window

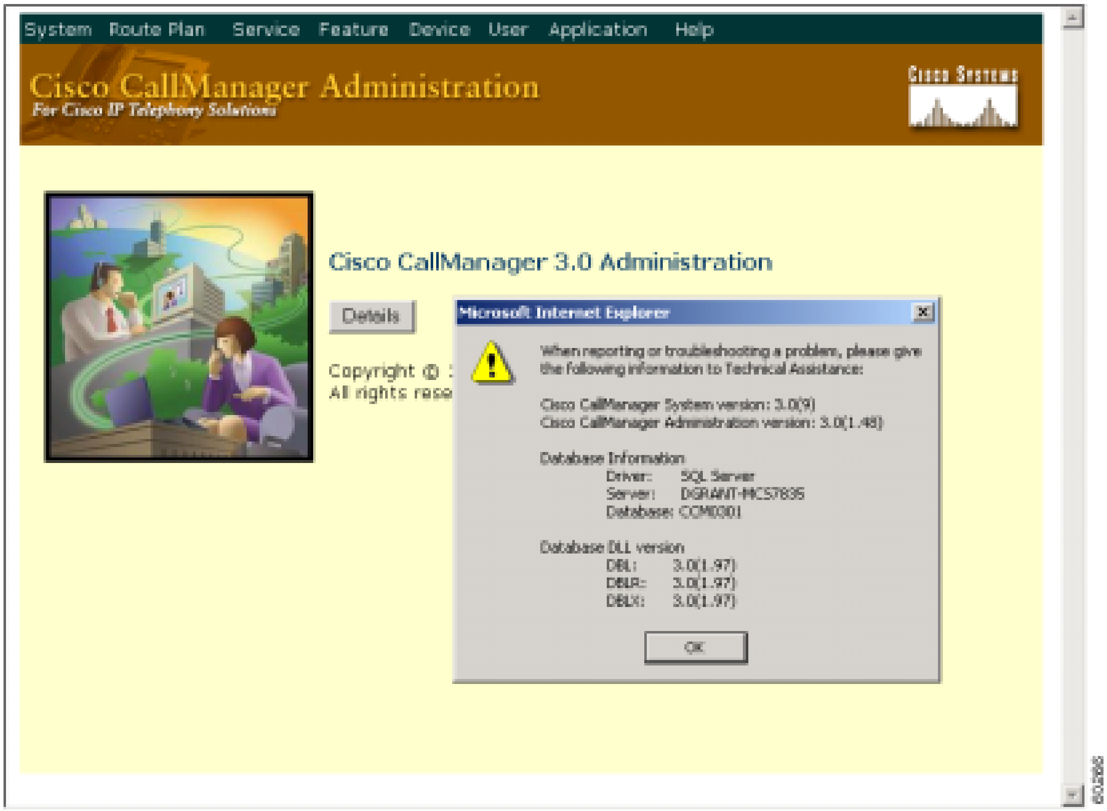

System
Route Plan
Service
Feature
Device
User
Application
Help
Cisco CallManager Administration
For Cisco IP Telephony Solutions
Cisco Systems
Cisco CallManager 3.0 Administration
Details
Copyright © :
All rights rese
Microsoft Internet Explorer
OK
60298

When reporting or troubleshooting a problem, please give the following information to Technical Assistance:

Cisco CallManager System version: 3.0(9)
Cisco CallManager Administration version: 3.0(1.48)

```
Database Information
          Driver:     SQL Server
          Server:     DGRANT-MCS7835
          Database: CCM0301

Database DLL version
          DBL:      3.0(1.97)
          DBLR:     3.0(1.97)
          DBLX:     3.0(1.97)
```

A more detailed explanation of Cisco CallManager Administration is available at the following location: http://www.cisco.com/univercd/cc/td/doc/product/voice/c_callmg/3_0/index.htm.

Microsoft Performance Monitor

Performance Monitor (PerfMon) is a Windows 2000 application that can display the activities and status of your Cisco CallManager system. Figure 6-33 illustrates attributes available via the PerfMon screen. It reports both general and specific information in real time. You can use Windows 2000 Performance Monitor to collect and display system and device statistics for any Cisco CallManager installation. This administrative tool allows you to gain a full understanding of a system without studying the operation of each of its components.

You can use PerfMon to monitor a variety of system variables in real time. After adding the Cisco CallManager parameters, you can define the terms under which Cisco CallManager will display statistics generated by the system. For example, you can monitor the number of calls in progress at any time, or the number of calls currently passing through a specific gateway. Performance shows both general and Cisco CallManager-specific status information in real-time.

Figure 6-33 PerfMon Window

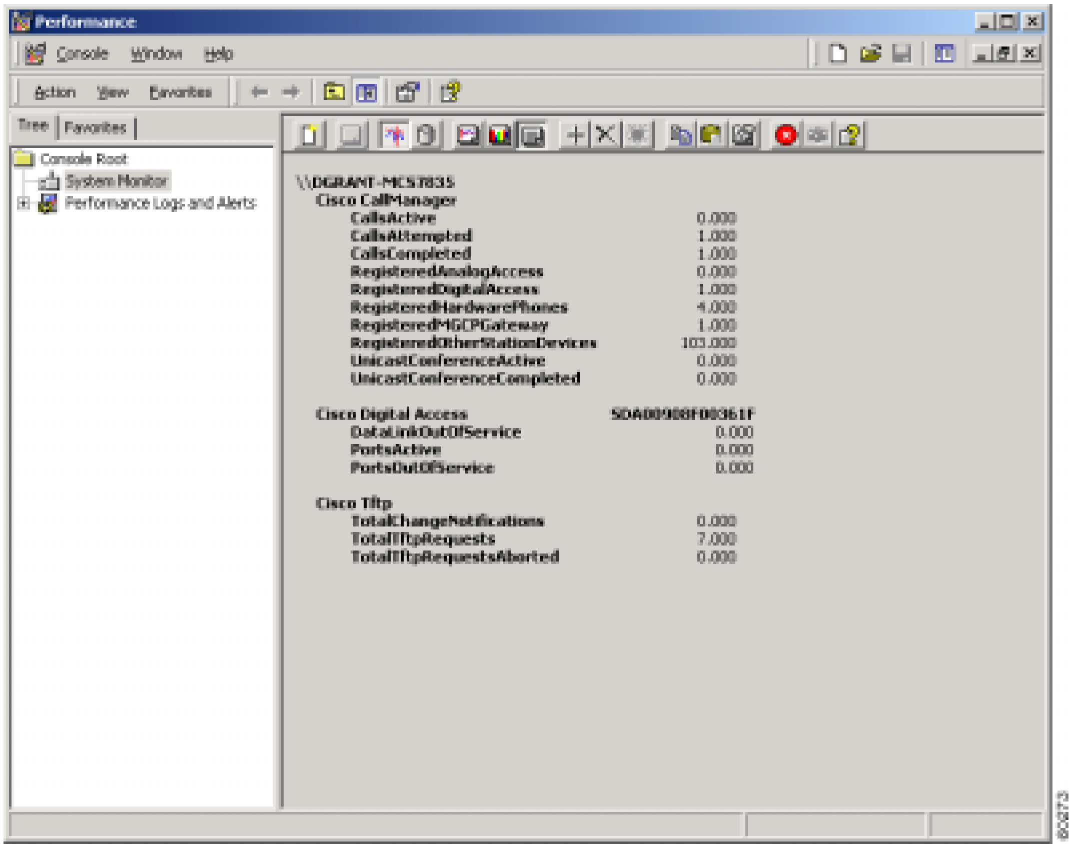

```
\\DGRANT-MCS7835
  Cisco CallManager
      CallsActive                              0.000
      CallsAttempted                           1.000
      CallsCompleted                           1.000
      RegisteredAnalogAccess                   0.000
      RegisteredDigitalAccess                  1.000
      RegisteredHardwarePhones                 4.000
      RegisteredMGCPGateway                    1.000
      RegisteredOtherStationDevices          103.000
      UnicastConferenceActive                  0.000
      UnicastConferenceCompleted               0.000

  Cisco Digital Access                  SDA00908F00361F
      DataLinkOutOfService                       0.000
      PortsActive                                0.000
      PortsOutOfService                          0.000

  Cisco Tftp
      TotalChangeNotifications                 0.000
      TotalTftpRequests                        7.000
      TotalTftpRequestsAborted                 0.000
```

Opening Microsoft Performance

To open Performance on the server running Cisco CallManager, click **Start** > **Settings** > **Control Panel** > **Administration Tools** > **Performance**.

Customizing Performance

The Performance monitor must be customized to view the Cisco CallManager-related parameters that you wish to monitor. Choose the object, counter, and instance you want to include. Please refer to the Remote Serviceability documentation for instructions on how to use objects and counters to customize Microsoft Performance for Cisco CallManager operations. This document can be found at the following location: http://www.cisco.com/univercd/cc/td/doc/product/voice/c_callmg/3_0/service/index.htm.

Microsoft Event Viewer

Microsoft Event Viewer is a Windows NT Server application that displays system, security, and application events (including Cisco CallManager) for the Windows NT Server. If a service (including TFTP) cannot read the database (where it gets the trace configuration), it will add errors to the Event Viewer. The Event Viewer is the only place where these types of errors will appear. Figure 6-34 shows the application logs running on a Windows NT Server.

Figure 6-34 Application Logs Running on Windows NT Servers

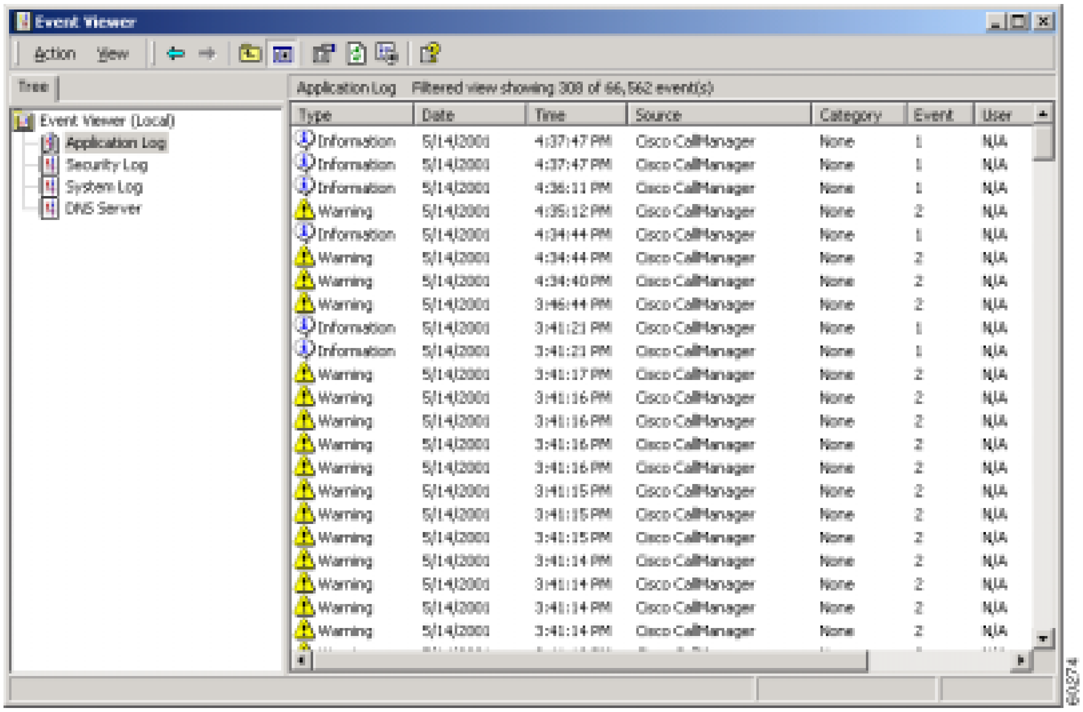

Application Log Filtered view showing 308 of 66,562 event(s)

Type	Date	Time	Source	Category	Event	User
Information	5/14/2001	4:37:47 PM	Cisco CallManager	None	1	N/A
Information	5/14/2001	4:37:47 PM	Cisco CallManager	None	1	N/A
Information	5/14/2001	4:36:11 PM	Cisco CallManager	None	1	N/A
Warning	5/14/2001	4:35:12 PM	Cisco CallManager	None	2	N/A
Information	5/14/2001	4:34:44 PM	Cisco CallManager	None	1	N/A
Warning	5/14/2001	4:34:44 PM	Cisco CallManager	None	2	N/A
Warning	5/14/2001	4:34:40 PM	Cisco CallManager	None	2	N/A
Warning	5/14/2001	3:46:44 PM	Cisco CallManager	None	2	N/A
Information	5/14/2001	3:41:21 PM	Cisco CallManager	None	1	N/A
Information	5/14/2001	3:41:21 PM	Cisco CallManager	None	1	N/A
Warning	5/14/2001	3:41:17 PM	Cisco CallManager	None	2	N/A
Warning	5/14/2001	3:41:16 PM	Cisco CallManager	None	2	N/A
Warning	5/14/2001	3:41:16 PM	Cisco CallManager	None	2	N/A
Warning	5/14/2001	3:41:16 PM	Cisco CallManager	None	2	N/A
Warning	5/14/2001	3:41:16 PM	Cisco CallManager	None	2	N/A
Warning	5/14/2001	3:41:15 PM	Cisco CallManager	None	2	N/A
Warning	5/14/2001	3:41:15 PM	Cisco CallManager	None	2	N/A
Warning	5/14/2001	3:41:15 PM	Cisco CallManager	None	2	N/A
Warning	5/14/2001	3:41:14 PM	Cisco CallManager	None	2	N/A
Warning	5/14/2001	3:41:14 PM	Cisco CallManager	None	2	N/A
Warning	5/14/2001	3:41:14 PM	Cisco CallManager	None	2	N/A
Warning	5/14/2001	3:41:14 PM	Cisco CallManager	None	2	N/A

Opening Event Viewer

To open the Event Log on the server PC running Cisco CallManager, click **Start** > **Settings** > **Control Panel** > **Administrative Tools** > **Event Viewer**. The Event Viewer provides error logs for System, Security, and Applications. Cisco CallManager errors are logged under the Application log.

Detailed Information about Events

You can double-click an event in the log to learn more information about the event, as illustrated in Figure 6-35.

Figure 6-35 Event Properties Description

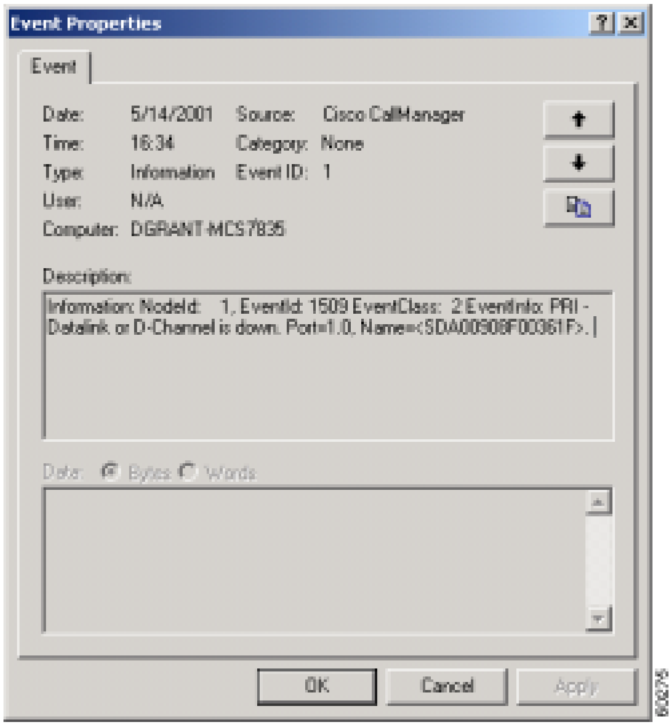

Event Properties

Event

Date: 5/14/2001 Source: Cisco CallManager
Time: 16:34 Category: None
Type: Information Event ID: 1
User: N/A
Computer: DGRANT-MCS7835

Description:

Information: NodeId: 1, EventId: 1509 EventClass: 2 EventInfo: PRI - Datalink or D-Channel is down. Port=1.0, Name=<SDA00908F00361F>.

Data: ● Bytes ○ Words

CallManager Trace

CallManager traces are local log files. The IP address, TCP handle, device name, or the time stamp can be used when reviewing the CallManager trace to monitor the occurrence or the disposition of a request. This device name could be tracked back to the building of the file, which shows the device pool and model. The device pool and model can be tracked back to the building of the configuration file prototype, which will list the network address of the Cisco CallManager(s) and the TCP connection port.

When observing CallManager traces, notice that C++ class and routine names are included with most trace lines. Most routines associated with the serving of a particular request include the thread ID in a standard format.

CallManager traces will be explained in detail in the case studies.

CallManager Trace Output

CallManager traces generate files (for example, CCM000000000) that store traces of Cisco CallManager activities. These traces provide information about the Cisco CallManager initialization process, registration process, KeepAlive process, call flow, digit analysis, and any registered devices such as Cisco IP phones, gateways, gatekeepers, and more. This information can help you isolate problems when troubleshooting Cisco CallManager. To properly track the information you need—and only the information you need—it's important to understand how to set the options on the trace configuration interface.

The trace files are stored in the following default location: C:\Program Files\Cisco\Trace\CCM. A new trace file is started each time Cisco CallManager restarts, or when the designated number of lines has been reached.

Figure 6-36 illustrates the Cisco CallManager Administration trace configuration interface. You must enable the trace, choose the level of detail needed, and check the user mask to obtain the desired level of information.

Figure 6-36 Cisco CallManager Administration Trace Configuration Interface

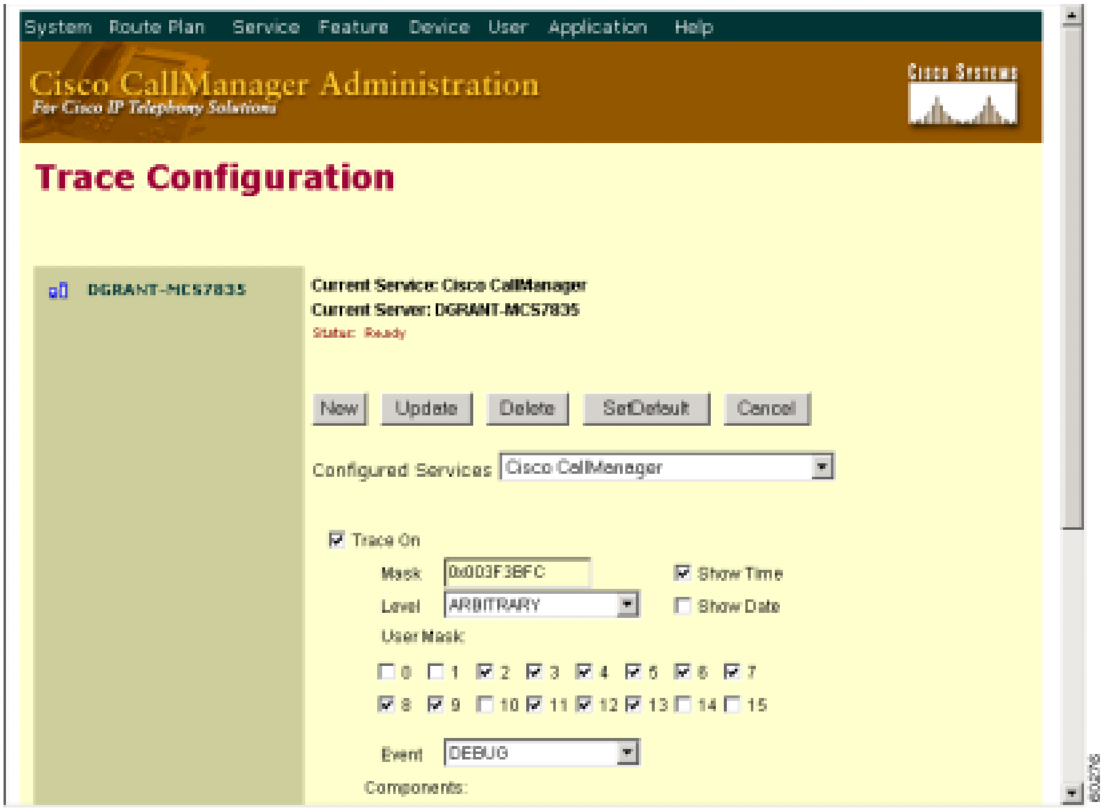

Trace Configuration

Current Service: Cisco CallManager
Current Server: DGRANT-MCS7835
Status: Ready

Configured Services: Cisco CallManager

☑ Trace On

Mask: 0x003F3BFC ☑ Show Time
Level: ARBITRARY □ Show Date

User Mask:

□ 0 □ 1 ☑ 2 ☑ 3 ☑ 4 ☑ 5 ☑ 6 ☑ 7
☑ 8 ☑ 9 □ 10 ☑ 11 ☑ 12 ☑ 13 □ 14 □ 15

Event: DEBUG

Components:

If the trace is not configured properly, it will generate a large amount of information making it very difficult to isolate problems. The following section explains how to properly configure a useful trace.

Configuring Traces

Traces are composed of user mask flags (also known as bits) and trace levels. Open Cisco CallManager Administration. To turn on tracing, set your trace parameters (including configured service, bits, and so on) in the **Service** > **Trace** screen. Refer to the Cisco CallManager documentation for complete information about turning tracing on and off, and for descriptions of the User Masks and Levels for each configured service, and more. This document can be found at the following location: http://www.cisco.com/univercd/cc/td/doc/product/voice/c_callmg/3_0/admin_gd/admin_gd/index.htm.

Following are two examples of trace mask bits that would be enabled based on the particular problem.

- For normal message debugging, turn on subsystem bits 5, 6, 7, 8, 11, and 12
- For debugging gateways, turn on subsystem bits 3, 4, 5, 6, 7, 8, 9, 11, 12, and 13

Following are two examples of desired trace levels based on the particular problem

- For normal debugging, the trace level should be set to ARBITRARY
- For normal running system, the trace level should be set to ERROR

SDL Trace

SDL trace provides a C level interface to trace and alarms. Alarms are used to inform the administrator of unexpected events, such as being unable to access a file, database, Winsock, or being unable to allocate other operating system resources.

Cisco engineers use SDL traces to find the cause of an error. You are not expected to fully understand the information contained in an SDL trace. However, while working with TAC, you may be asked to enable the SDL trace and provide it to the TAC. SDL trace files can be saved to local directories, the Windows NT Event Viewer, and CiscoWorks 2000. To avoid any performance degradation on the server, be sure that you turn off SDL tracing after the trace has been captured.

Enabling SDL Trace

SDL traces are enabled in the **Service** > **Service Parameter** area in Cisco CallManager Administration. Remember that these traces should be turned on only when requested by a TAC engineer. Note the values chosen to turn on the SDL trace in Figure 6-37.

Figure 6-37 Service Parameters Configuration Window

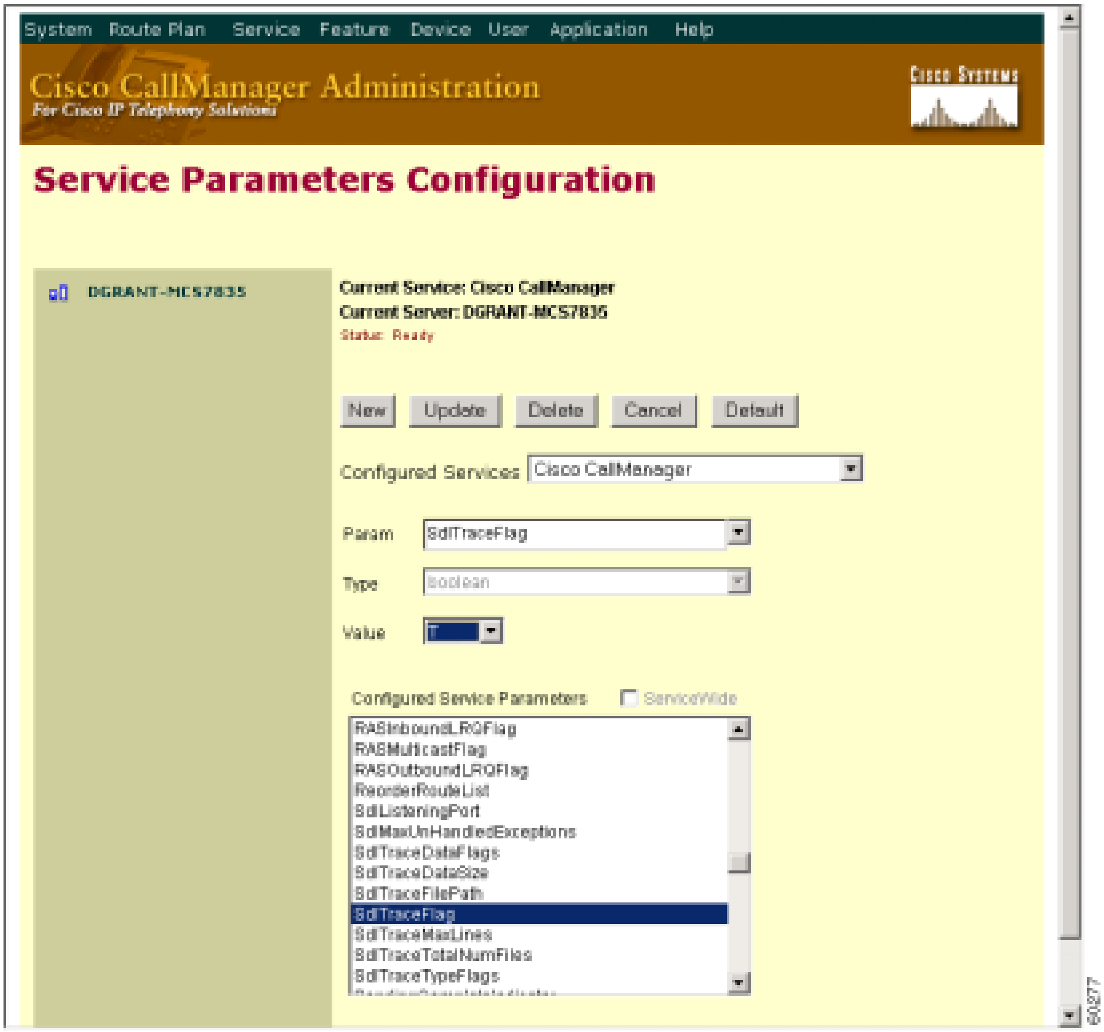

Current Service: Cisco CallManager
Current Server: DGRANT-MCS7835
Status: Ready

Configured Services: Cisco CallManager

Param: SdlTraceFlag

Type: boolean

Value: T

Configured Service Parameters ☐ ServiceWide

RASInboundLRQFlag
RASMulticastFlag
RASOutboundLRQFlag
ReorderRouteList
SdlListeningPort
SdlMaxUnHandledExceptions
SdlTraceDataFlags
SdlTraceDataSize
SdlTraceFilePath
SdlTraceFlag
SdlTraceMaxLines
SdlTraceTotalNumFiles
SdlTraceTypeFlags

Once SDL traces are enabled, collect the traces. If the traces are being sent to the local drive, then you can retrieve them in the Cisco\Trace subdirectory. Alternatively, the trace files can be sent to an event log or to CiscoWorks 2000.

SDL flag bits described in Table 6-32 are set in the **Service** > **Service Parameters** area in Cisco CallManager Administration. Following are two examples of desired values based on the particular problem.

- The recommended value for normal call debugging is SdlTraceTypeFlags=0x00000b04
- The recommended value for low level debugging or debugging gateways is SdlTraceTypeFlags=0x00004b05

Table 6-29 SDL TraceTypeFlag Definitions

SDLTraceTypeFlag	Value	Definition
traceLayer1	= 0x00000001	All Layer 1 trace on
TraceDetailLayer1	= 0x00000002	Detail Layer 1 trace on
TraceSdlLinkAdmin	= 0x00000004	Trace inter-Cisco CallManager links within a cluster
traceUnused	= 0x00000008	Not used
traceLayer2	= 0x00000010	All Layer 2 trace on
traceLayer2Interface	= 0x00000020	Layer 2 interface trace on
traceLayer2TCP	= 0x00000040	Layer 2 TCP trace on
TraceDetailLayer2	= 0x00000080	More detail dump of Layer 2 frames
traceLayer3	= 0x00000100	All Layer 3 trace on
traceCc	= 0x00000200	All call control trace on
traceMiscPolls	= 0x00000400	Trace miscellaneous polls
traceMisc	= 0x00000800	Miscellaneous trace on (database signals)
traceMsgtrans	= 0x00001000	Message Translation signals (TranslateIsdnToSdlReq, TranslateIsdnToSdlRes TranslateSdlToIsdnReq, TranslateSdlToIsdnRes)
traceUuie	= 0x00002000	UUIE output trace on
traceGateway	= 0x00004000	Gateway signals

Data bits described inTable 6-30 are set in the **Service** > **Service Parameters** area in Cisco CallManager Administration. Following are two examples of desired values based on the particular problem.

- The recommended value for normal system debugging is SdlTraceDataFlags=0x110
- The recommended value when tracking problems with SDL links is 0x13D (non-compacted trace; if a compact trace is desired, bit 0x200 must be set. It can be set in combination with any other bits).

Table 6-30 SDLTraceDataFlags Definitions

SDLTraceDataFlag	Value	Definition
TraceSdlLinkState	= 0x001	Enable trace of SDL Link Initialization
TraceSdlLowLevel	= 0x002	Enable tracing of low-level SDL events, fileOpen and socket events (for example)
TraceSdlLinkPoll	= 0x004	Enable tracing of SDL Link Poll message
TraceSdlLinkMsg	= 0x008	Enable tracing of SDL Link message
traceRawData	= 0x010	Enable raw signal data trace on all signals
TraceSdlTagMap	= 0x020	Enable tag mapping
traceCreate	= 0x100	Enable process create and stop traces
TraceNoPrettyPrint	= 0x200	Disable pretty printing of trace files

⚠ Caution **Disk Space Warning—**Be advised that information obtained from this interface could be very detailed, and therefore can consume a large amount of disk space. For this reason, we advise you to turn on the trace file for a specific amount of time, review the information, and turn off the trace.

Sniffer Trace

A *sniffer* is a software application that monitors IP traffic on a network and provides information in the form of a trace. Sniffer traces provide information about the quantity and type of network traffic on your network. TCP/IP or UDP packets are protocols utilized by Cisco CallManager and by endpoint devices, such as phones and gateways. Sniffer traces can also help you identify high levels of broadcast traffic that could result in voice audio problems or dropped calls. Common sniffer applications include Network Associates SnifferPro, Hewlett Packard Internet Advisor, and W&G Domino. Domino offers sniffing hardware and software solutions and a network analyzer. If you want to use Domino, we recommend using the analysis software to evaluate a captured sniffer file (such as from the SnifferPro application).

Sniffer Trace Applications

Use the following links to learn more about some available sniffer trace applications. Any sniffer application will work with Cisco CallManager.

- Network Associates *SnifferPro*

  http://www.sniffer.com/

- Acterna *Domino Analyzer*

  http://www.acterna.com/products/domino/index.html

Call Detail Records (CDR) and Call Management Records (CMR)

CDR is a reporting option that logs every call made (or attempted) from any Cisco IP phone. There are two kinds of CDRs—basic CDRs and Diagnostic CDRs (or CMRs). Once enabled, you can open CDRs or Diagnostic CDRs (CMRs) in the SQL Server Enterprise Manager. CDR files are saved in a SQL database that can be exported to nearly any application, including Microsoft Access or Excel.

CDR records contain information needed to generate billing records. In a distributed environment, all CDR data is collected in a central location, or a set of locations. The failure of a Cisco CallManager node does not make the CDR data associated with that node unavailable. The data is no longer stored on the Cisco CallManager disk as a flat file, but is instead stored in a central database in tables.

If the Cisco CallManager fails before any records are written, no record of the call will exist. This means that no record will be written for calls that are active on a given Cisco CallManager when it fails before the calls terminate.

Refer to the Call Detail Records section of this document for detailed information about CDRs and CMRs.

The information provided includes:

- Reading and writing records
- Known issues
- List of record types generated
- List of fields contained in each record and a description of what that field represents
- Description of the types of calls logged, and the fields logged with each of them
- List of cause codes that may appear in the CDR records

Enabling or Disabling CDRs

CDR record creation is disabled by default when the system is installed. If you wish to have CDR data, you must enable CDRs in the **Service > Service Parameters** area of Cisco CallManager Administration. CDR processing can be enabled and disabled at any time while the system is in operation. You do not need to restart Cisco CallManager for the enabling or disabling of CDRs to take effect. The system will respond to all changes within a few seconds. CMR or diagnostic data is enabled separately from CDR data. CMR data will not be generated unless both CDRs and Call Diagnostics are enabled, but CDR data may be generated and logged without CMR data.

Use the following steps to enable CDRs. Figure 6-38 illustrates the screen associated with this process.

Step 1 Open Cisco CallManager Administration.

Step 2 Select **Service** > **Service Parameters**.

Step 3 Select the IP address of your Cisco CallManager installation.

Step 4 Select Cisco CallManager from the list of configured services.

Step 5 From the list of Parameters, select **CDREnabled**.

Step 6 Define type as **boolean**.

Step 7 Select **T** for True.

Step 8 Click Update.

Step 9 Repeat Step 3 to Step 8 for each CallManager in the cluster.

Result: Call Detail Records will start logging immediately.

⚠ **Caution** Tracing voice connectivity requires that CDR logging be enabled on every Cisco CallManager installation in a cluster.

Figure 6-38 Enabling CDRs in the Service Parameters Configuration Window

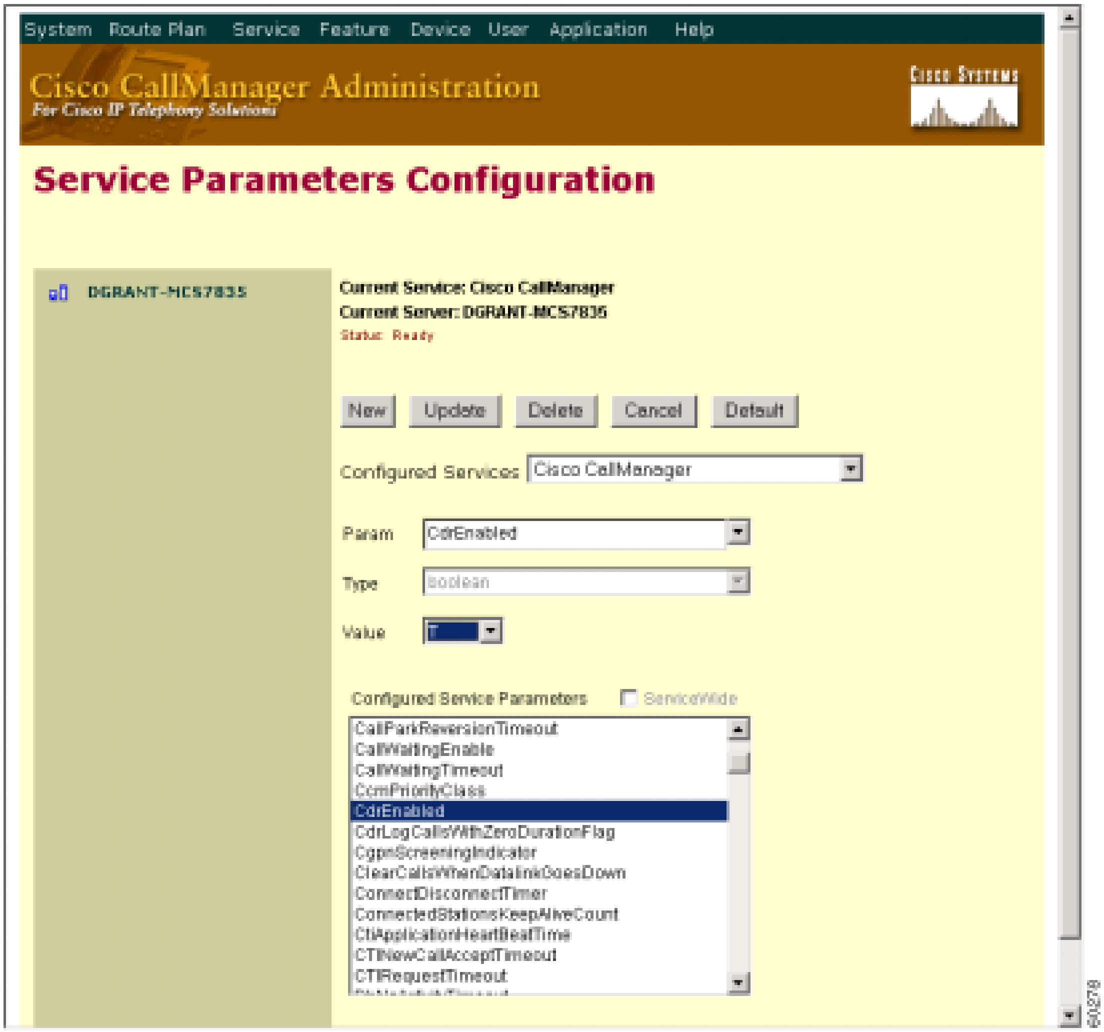

Service Parameters Configuration

DGRANT-MCS7835

Current Service: Cisco CallManager
Current Server: DGRANT-MCS7835
Status: Ready

Configured Services: Cisco CallManager

Param: CdrEnabled

Type: boolean

Value: T

Configured Service Parameters

CallParkReversionTimeout
CallWaitingEnable
CallWaitingTimeout
CcmPriorityClass
CdrEnabled
CdrLogCallsWithZeroDurationFlag
CgpnScreeningIndicator
ClearCallsWhenDatalinkGoesDown
ConnectDisconnectTimer
ConnectedStationsKeepAliveCount
CtiApplicationHeartBeatTime
CTINewCallAcceptTimeout
CTIRequestTimeout

CDRs

CDRs provide basic information that can help you understand the more detailed information contained in CallManager traces. Basic CDRs provide information such as the calling number, called number, originating IP address, destination IP address, call duration, and so on. CDRs can help you troubleshoot phone problems. For example, if a user reports a problem with a call occurring at a specific time, you can consult the CDRs that occurred around the time indicated to learn additional information about that call and others. CDRs are commonly used for billing.

Diagnostic CDRs (Also Known As CMRs)

Diagnostic CDRs provide detailed call information such as the number of packets sent, received, and lost, and the amount of jitter and latency. This level of detail can provide explanations for some problems, such as one-way audio. For example, a one-way audio problem is indicated if a packet size of 10,000 is sent, but the received size is only 10.

Q.931 Translator

The Q.931 Translator that is packaged with Cisco CallManager is a tool that is used to translate CallManager traces of H.225 setup messages into an easier to read, IOS-equivalent format. H.225 is part of the H.323 protocol stack, and is used for call control signaling, and is based on Q.931. This tool is only useful for calls that access an H.323 Gateway (Cisco IOS router) or other H.323 device (e.g. Microsoft NetMeeting).

Understanding the Translation Process

The Message Translator works by filtering incoming data from Cisco CallManager log files, then parsing and translating them into Cisco IOS-equivalent messages. The application displays the messages in the Message Translator interface, as shown below.

Using the Message Translator

Start the message translation process by using the following procedure to locate the Cisco CallManager log file in the directory structure. Figure 6-39 illustrates the associated translation screen.

Figure 6-39 Q931 Message Translator

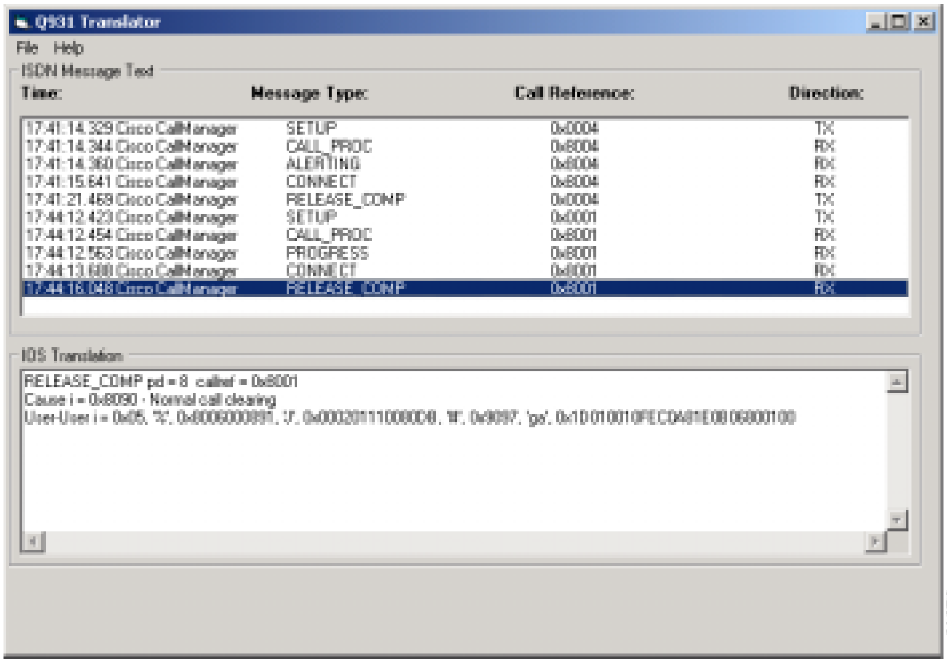

Time:	Message Type:	Call Reference:	Direction:
17:41:14.329 Cisco CallManager	SETUP	0x0004	TX
17:41:14.344 Cisco CallManager	CALL_PROC	0x8004	RX
17:41:14.360 Cisco CallManager	ALERTING	0x8004	RX
17:41:15.641 Cisco CallManager	CONNECT	0x8004	RX
17:41:21.469 Cisco CallManager	RELEASE_COMP	0x0004	TX
17:44:12.423 Cisco CallManager	SETUP	0x0001	TX
17:44:12.454 Cisco CallManager	CALL_PROC	0x8001	RX
17:44:12.563 Cisco CallManager	PROGRESS	0x8001	RX
17:44:13.688 Cisco CallManager	CONNECT	0x8001	RX
17:44:16.048 Cisco CallManager	RELEASE_COMP	0x8001	RX

```
RELEASE_COMP pd = 8  callref = 0x8001
Cause i = 0x8090 - Normal call clearing
User-User i = 0x05, 'X', 0x8006000891, '/', 0x00020111000D8, '#', 0x9097, 'gs', 0x1D010010FECDA81E0806000100
```

Step 1 Launch Q.931 Translator from C:\Program Files\Cisco\Bin\q931translator.exe.

Step 2 First pull down the **File** menu and select **Open**.

Step 3 Select a log file in the directory listing for translation.

Step 4 To save the log file you have selected from the directory listing, type in a name that will identify it as a translated Cisco IOS log file. The log file you select is loaded into the top box of the Message Translator.

Step 5 Select the message to be translated. The Cisco IOS translation of the selected message appears immediately in the ISDN translation box at the bottom of the screen.

Step 6 Keeping the translated files on hand may be of use to you later on. To do this, pull down the **File** menu and select the **Save As IOS** option, which translates and saves all ISDN messages in the log file.

Troubleshooting Cisco CallManager Devices

This section addresses some common problem categories that may occur with Cisco CallManager and related devices. Each problem category suggests troubleshooting tools you should use to help isolate the problem. This document provides general categories of potential problems and suggestions on how to troubleshoot those problems. It does not provide an exhaustive list of problems and resolutions. If you encounter a problem that cannot be resolved using the tools and utilities described in this document, consult the Cisco Technical Assistance Center (TAC) for assistance. Be sure to have the Cisco CallManager Administration Details available, plus the any diagnostic information (such as traces) you have gathered up to the point of calling the TAC.

Problem categories presented here are as follows:

- Voice Quality
- Phone Resets
- Dropped Calls
- Cisco CallManager Feature Issues
- Slow Server Response
- Reorder Tone Through Gateways
- Gateway Registration Problems
- Gatekeeper Registration Problems
- Cisco IP Phone Initialization Process
- Skinny Station Registration Process
- Cisco CallManager Initialization Process
- Self-Starting Processes
- Cisco CallManager Registration Process
- Cisco CallManager KeepAlive Process
- Cisco IP Phone-to-Cisco IP Phone Exchange of Skinny Station Messages During Call Flow
- Cisco CallManager Intra-Cluster Call Flow Traces
- Cisco IOS Gateways
- Experiencing Busy Signal After Dialing International Numbers
- Inter-Cluster IP Phone to IP Phone
- Call Flow Traces
- Failed Call Flow

Voice Quality

Voice quality issues include lost or distorted audio during phone calls. Common problems include breaks in the sound, which cause the audio to be intermittent (like broken words), or the presence of odd noises that distort the audio (echo) or effects that cause spoken words to sound watery or robotic. One-way audio, that is, a conversation between two people where only one person can hear anything, is not actually a voice quality issue. This will be discussed later in this section.

One or more of the following components may cause audio problems:

- Gateway
- Phone
- Network

To properly troubleshoot voice quality issues, you must consider the infrastructure and all the devices for drops and delays.

i Button Help

The Cisco IP Phone 7960 provides another tool for diagnosing possible audio problems. On an active call, you can press the ⓘ button twice (rapidly) and the phone displays an information screen that contains packet receive and transmit statistics, as well as average and maximum jitter counters. On this screen, note that *jitter* is the average of the last 16 packets that arrived; the maximum jitter is the high-water mark for the average jitter. Information provided includes the following:

- **RxType/TxType**—These are the codecs being used on the current call in progress.
- **RxSize/TxSize**—This is the size of the payload in each packet, measured in milliseconds.
- **RxCnt/TxCnt**—The amount of packets sent/received during the current call in progress.
- **AvgJtr**—This is the average jitter observed in the last 16 RTP packets.
- **MaxJtr**—This is the maximum jitter seen during the life of the RTP receive stream. Note: This is on a per-stream basis, **not** for the life of the call. If you put a call on hold, the stream stops, and new stream will be created when the call is taken off hold.
- **RxDisc**—The amount of inbound packets discarded.
- **RxLost**—The amount of packets lost in the path from the remote side.

The most common sources for delay and packet loss are devices where a higher speed interface feeds into a lower speed interface. For example, a router may have a 100 Mb fast Ethernet interface connected to the LAN and a slow frame-relay connected to the WAN. When the poor quality occurs only when communicating to the remote site (only the remote site may be reporting the poor voice quality while in the other direction everything appears to be fine), below are the most likely causes of the problem:

- The router has not been properly configured to give the voice traffic priority over the data traffic.
- There are too many calls active for the WAN to support (that is, there is no call admission control to restrict the number of calls that can be placed).
- There are physical port errors.
- There is congestion in the WAN itself.

On the LAN, the most common problems are physical-level errors (such as CRC errors) caused by faulty cables and interfaces, or by incorrectly configured devices (such as a port speed or duplex mismatch). Make sure that the traffic is not crossing any shared-media device, such as a hub. There could also be situations where the traffic is taking a slower path through the network than expected. If QoS has been configured correctly, it is possible that there is no call admission control. Depending on your topology, this can be accomplished through the use of Locations in Cisco CallManager Administration configuration, or by using a Cisco IOS router as a gatekeeper. In any case, you should always know how many calls could be supported across your WAN. If possible, test this by disabling silence suppression as described earlier, then place calls between the two sites. Do not place the calls on hold or on mute, since this will stop packets from being transmitted. With the maximum number of calls across the WAN, the calls should all have acceptable quality. Test to make sure that a fast busy is returned when trying to make one more call.

Lost or Distorted Audio

One of the most common problems encountered is a *breaking up* of audio (often described as garbled speech or a loss of syllables within a word or sentence). There are two common causes for this: packet loss and/or jitter. Packet loss means that audio packets do not arrive at their destination because they were dropped or arrived too late to be useful. Jitter is the variation in the arrival times of packets. In the ideal situation, all VoIP packets from one phone to another would arrive exactly at a rate of 1 every 20 msec. Notice that this does not mention how *long* it takes for a packet to get from point A to point B, simply the *variation* in the arrival times.

There are many sources of variable delay in a real network. Some of these cannot be controlled, and some can. Variable delay cannot be eliminated entirely in a packetized voice network. Digital Signal Processors (DSPs) on phones and other voice-capable devices are designed to buffer some of the audio, in anticipation of variable delay. This *dejittering* is done only when the audio packet has reached its destination and is ready to be put into a conventional audio stream (that is, played into the user's ear, or sent to the PSTN via a digital PCM stream). The Cisco IP Phone 7960 can buffer as much as one second of voice samples. The jitter buffer is adaptive, meaning if a burst of packets is received, the Cisco IP Phone 7960 can play them out in an attempt to control the jitter. The network administrator needs to minimize the variation between packet arrival times by applying quality-of-service (QoS) and other measures in advance (especially if calls cross a wide-area network).

When faced with a lost or distorted audio problem, you should first try to isolate the path of the audio. Try to identify each network device (switches and routers) in the path of the call's audio stream. Keep in mind that the audio may be between two phones, between a phone and a gateway, or it could have multiple legs (from a phone to a transcoding device and from there to another phone). Try to identify whether the problem occurs only between two sites, only through a certain gateway, on a certain subnet, and so on. This will help to narrow down which devices you need to examine more carefully. Next, it is often best to disable silence suppression (also known as Voice Activation Detection or VAD) if this hasn't been done already. This mechanism does save bandwidth by not transmitting any audio when there is silence, but may cause noticeable (and unacceptable) clipping at the beginning of words. You can disable this in Cisco CallManager Administration, under **Service** > **Service Parameters**. From there, select the server and the Cisco CallManager service as illustrated in Figure 6-40. Set SilenceSuppressionSystemWide to **F** (alternatively you can set SilenceSuppressionWithGateways to **F**, but this does not apply to H.323 gateways or MGCP gateways). When in doubt, turn both off by selecting the Value F for each.

Figure 6-40 SilenceSuppressionSystemWide Configuration Display

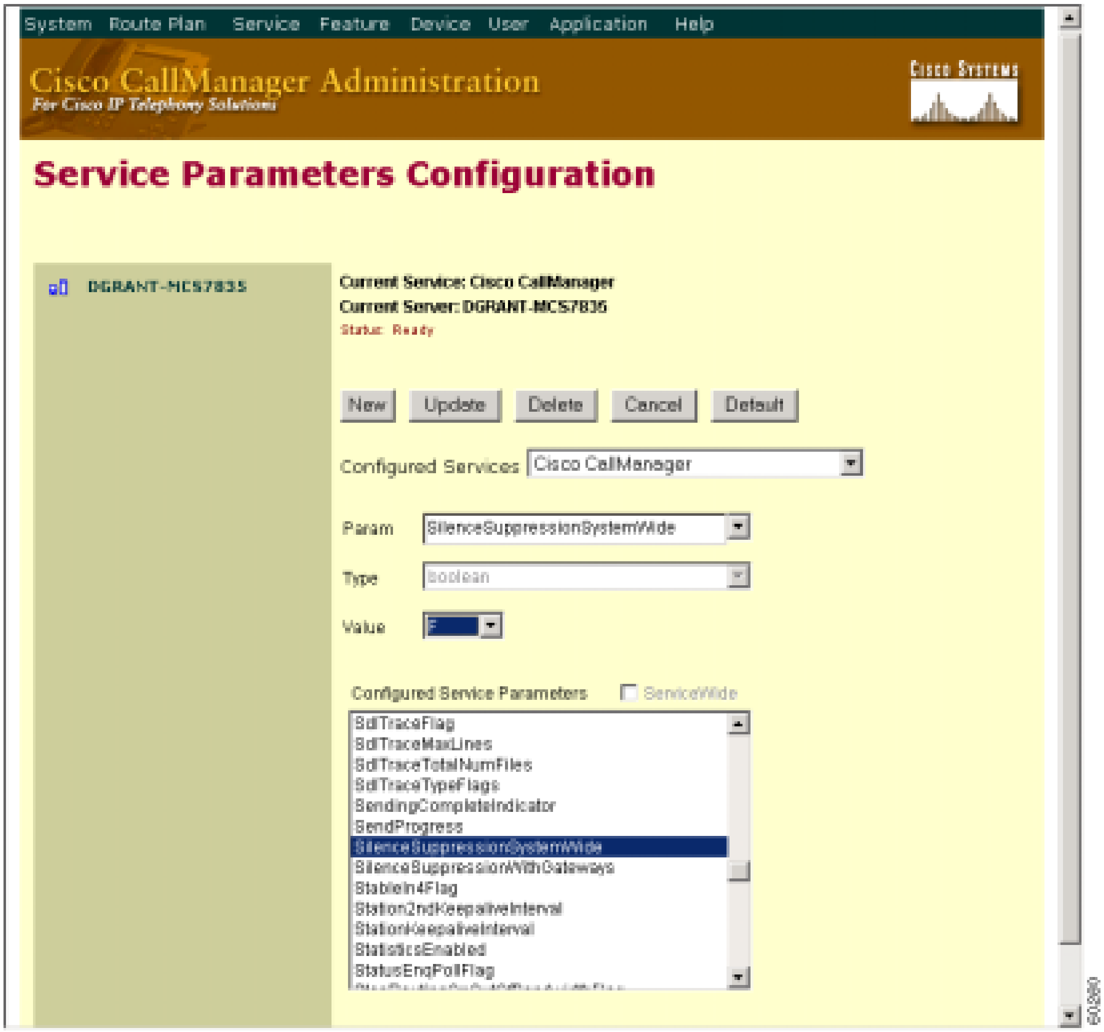

Service Parameters Configuration

Current Service: Cisco CallManager
Current Server: DGRANT-MCS7835
Status: Ready

Configured Services: Cisco CallManager

Param: SilenceSuppressionSystemWide

Type: boolean

Value: F

Configured Service Parameters □ ServiceWide

- SdlTraceFlag
- SdlTraceMaxLines
- SdlTraceTotalNumFiles
- SdlTraceTypeFlags
- SendingCompleteIndicator
- SendProgress
- SilenceSuppressionSystemWide
- SilenceSuppressionWithGateways
- StableIn4Flag
- Station2ndKeepaliveInterval
- StationKeepaliveInterval
- StatisticsEnabled
- StatusEnqPollFlag

If a network analyzer is available, a monitored call between two phones should have 50 packets per second (or 1 packet every 20 ms) when silence suppression is disabled. With proper filtering, it should be possible to identify if packets are being lost or delayed excessively.

Remember that delay by itself won't cause clipping, only variable delay will. Table 6-31 illustrates a *perfect* trace, with arrival times between the audio packets (which will have an RTP header) of 20 msec. In a poor quality call (such as a call with a lot of jitter), the arrival times would vary greatly.

Table 6-31 Example Output of a "Perfect" Trace

Packet Number	Time, Absolute (msec)	Time, Delta (msec)
1	0	
2	0.02	20
3	0.04	20
4	0.06	20
5	0.08	20

Placing the packet analyzer into various points in the network will help narrow down from where the delay is coming. If no analyzer is available, other methods will be required. It is important to examine interface statistics of each device in the path of the audio. Another tool for tracking calls with poor voice quality is the Diagnostic Call Detail Records (CDRs). See the **Tools and Utilities** section and the **Call Detail Records** section for more information about CDRs.

The values for jitter and latency can be retrieved for all calls (but only *after* the call has terminated). Figure 6-41 shows a sample Diagnostic CDR (CallDetailRecordDiagnostic is the actual table name). The number of packets sent, received, lost, jitter, and latency are all recorded. The globalCallID value can be used to find the call in the regular CDR table so that the disconnect cause and other information can be obtained. The diagram below shows both tables open. Notice that in the Diagnostic CDR, every device that can possibly report this information is included. So, if the problem is between two Cisco IP phones, we see two table entries per call. If we have a call through a Cisco IOS gateway, for example, we only see the diagnostic information from the Cisco IP phone, not the gateway because there is no mechanism for it to notify the SQL database with this information.

Figure 6-41 Call Detail Record Diagnostic Example

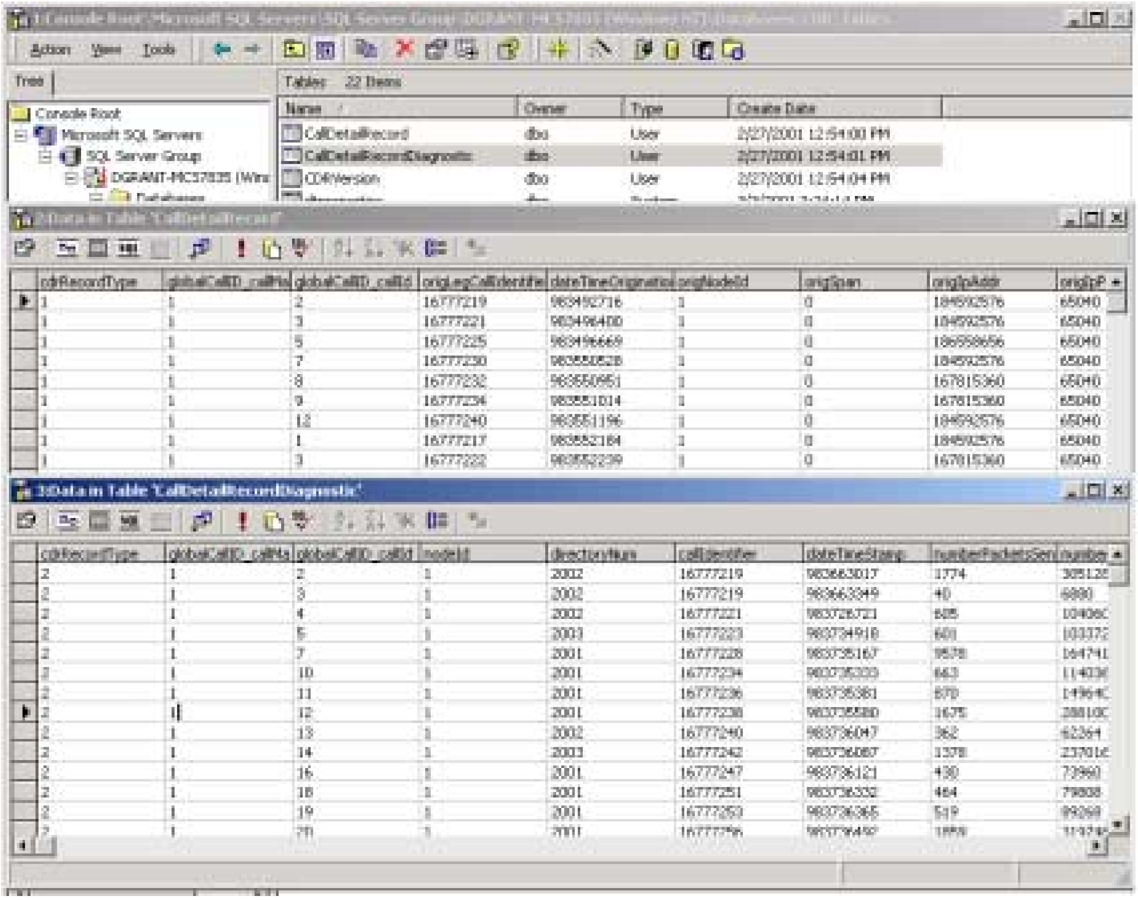
1:Console Root\Microsoft SQL Servers\SQL Server Group\DGRANT-MC37835 (Windows NT)\Databases\CDR\Tables
Action
View
Tools
Tree
Console Root
Microsoft SQL Servers
SQL Server Group
DGRANT-MC37835 (Wind
Databases
Tables 22 Items
2:Data in Table 'CallDetailRecord'
3:Data in Table 'CallDetailRecordDiagnostic'
60355

Name	Owner	Type	Create Date
CallDetailRecord	dbo	User	2/27/2001 12:54:00 PM
CallDetailRecordDiagnostic	dbo	User	2/27/2001 12:54:01 PM
CDRVersion	dbo	User	2/27/2001 12:54:04 PM

cdrRecordType	globalCallID_callMa	globalCallID_callId	origLegCallIdentifie	dateTimeOriginatio	origNodeId	origSpan	origIpAddr	origIpP
1	1	2	16777219	983492716	1	0	184592576	65040
1	1	3	16777221	983496400	1	0	184592576	65040
1	1	5	16777225	983496669	1	0	186558656	65040
1	1	7	16777230	983550528	1	0	184592576	65040
1	1	8	16777232	983550951	1	0	167815360	65040
1	1	9	16777234	983551014	1	0	167815360	65040
1	1	12	16777240	983551196	1	0	184592576	65040
1	1	1	16777217	983552184	1	0	184592576	65040
1	1	3	16777222	983552239	1	0	167815360	65040

cdrRecordType	globalCallID_callMa	globalCallID_callId	nodeId	directoryNum	callIdentifier	dateTimeStamp	numberPacketsSen	number
2	1	2	1	2002	16777219	983663017	1774	30512
2	1	3	1	2002	16777219	983663349	40	6880
2	1	4	1	2002	16777221	983726721	605	10408
2	1	5	1	2003	16777223	983734918	601	10337
2	1	7	1	2001	16777228	983735167	9578	16474
2	1	10	1	2001	16777234	983735333	663	11403
2	1	11	1	2001	16777236	983735381	870	14964
2	1	12	1	2001	16777238	983735580	1675	28810
2	1	13	1	2002	16777240	983736047	362	62264
2	1	14	1	2003	16777242	983736087	1378	23701
2	1	16	1	2001	16777247	983736121	430	73960
2	1	18	1	2001	16777251	983736332	464	79808
2	1	19	1	2001	16777253	983736365	519	89268
2	1	20	1	2001	16777256	983736492	1859	31374

Crackling

Another "poor quality" symptom may be a *crackling*, which is sometimes caused by a defective power supply or some kind of strong electrical interference close to the phone. Try swapping the power supply and moving the phone to a different location.

Echo

Echo (also known as *talker echo*) occurs when a talker's speech energy, transmitted down the primary signal path, is coupled into the receive path from the far end as illustrated in Figure 6-42. The talker then hears his or her own voice, delayed by the total echo path delay time. It is important to remember that echo may have existed in a traditional PBX network, but went unnoticed because of the low delay.

Figure 6-42 Echo Effect Illustrated

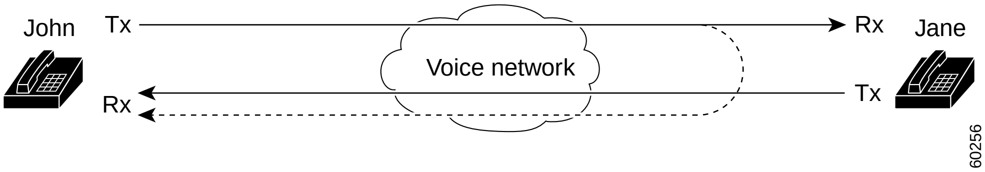

In Figure 6-42, John's voice is reflected back. This can go unnoticed in a traditional voice network because the delay is so low. To the user, it sounds more like a side-tone than an echo. In a VoIP network, this echo may be noticed since packetization and compression inserts enough delay such that echo may be heard. The important thing to remember is that the cause of the echo is always with analog components and wiring. For instance, IP packets cannot simply turn around and go back to the source at a lower audio level. The same is impossible on digital T1/E1 circuits. So on a call from one Cisco IP phone to another, there should never be any problem. The only exception may be if one party is using a speakerphone that has the volume set too high or some other situation where an audio loop is created.

When troubleshooting echo problems, make sure that the phones that are being tested or examined are not using the speakerphone and that they have the headset volume set to reasonable levels (start with 50 percent of the maximum audio level). Most of the time, the problems will occur when attaching to the PSTN by way of a digital or analog gateway. Cisco IP phone users may complain that they hear their own voice being reflected back to them. Although the true source of the problem is almost always at the far end, it is nearly always impossible to change anything in the PSTN. So, the first step is to determine which gateway is being used. If a digital gateway is in use, it may be possible to add additional padding in the transmit direction (towards the PSTN) in the hopes that the lower signal strength will yield less reflected energy. Additionally, you can adjust the receive level so that any reflected audio is reduced even further. It is very important to remember to make small adjustments at a time. Too much attenuation of the signal will make the audio impossible to hear on both sides. Alternatively, you can contact the carrier and request to have the lines checked. On a typical T1/PRI circuit in North America, the input signal should be—15 dB. If the signal level is much higher (-5 dB, for example), echo will be the likely result.

Keep a log of all calls that experience echo. The time of the problem, the source phone number, and the number called should all be recorded. Gateways can have up to 32 msec of echo cancellation. If the delay in the reflected audio is longer than this, the echo chancellor will be unable to work properly. This should not be an issue for local calls, and long distance calls should have external echo chancellors built into the network at the central office. This is one of the reasons why it is important to note the external phone number of a call that experiences echo.

So, in John's case above, the echo is being introduced somewhere on the PSTN side. One may argue that the cause of the echo can be anywhere in the network, this is true, but we must also remember this fact about echo:

> **Analog circuits cause echo**—It is *not possible* for bits on a digital circuit or in an IP network to leak from the Transmit path to the Receive path.

Echo Cancellation

Cisco IOS Gateways and Skinny Gateways have built-in echo cancellers. An echo canceller is a device that reduces the level of echo's introduced by a leak in the *tail circuit*. The following descriptions discuss the functionality and effectiveness of echo cancellers. Specific discussion include:

- Tail Circuit
- Echo Canceller Coverage
- Un-Cancelable Echo
- Troubleshooting Persistent Echo

Tail Circuit

The tail circuit is *everything* connected to the PSTN side of a packet voice gateway. That includes all the switches, multiplexors, cabling, PBXs—everything between the voice gateway and the telephone.

Figure 6-43 Echo Cancellation Effects Illustrated

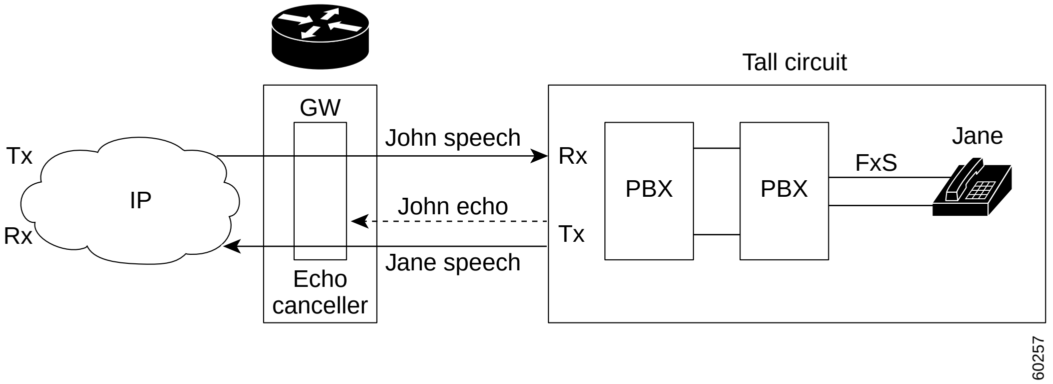

In Figure 6-43, the Rx Signal is John's voice coming from the IP network, and the Tx Signal is a mixture of Jane's voice and the echo of Bob's voice.

An echo canceller only works in the PSTN direction, that is, the echo canceller will only remove the echo generated on the tail circuit that it is connected to. The echo canceller will remove the echo portion of the signal that is coming out of the tail circuit heading for the IP network. The echo canceller does this by *learning* the electrical characteristics of the tail circuit, and forming its own model of the tail in memory. Using this model, it forms its own 'estimated echo' signal based on the current and past Rx signal (John's voice). John's voice is run through this functional model to come up with an estimate of John's echo signal. This estimated 'John echo' is then subtracted from the actual Tx signal coming out of the tail circuit.

Echo Canceller Coverage

The *echo canceller coverage* (*tail coverage* or *tail length*) is a parameter of an echo canceller that controls the length of the tail circuit in memory, and therefore determines the time window of cancelable echoes.

Figure 6-44 Echo Peak Amplitude versus Time Illustrated

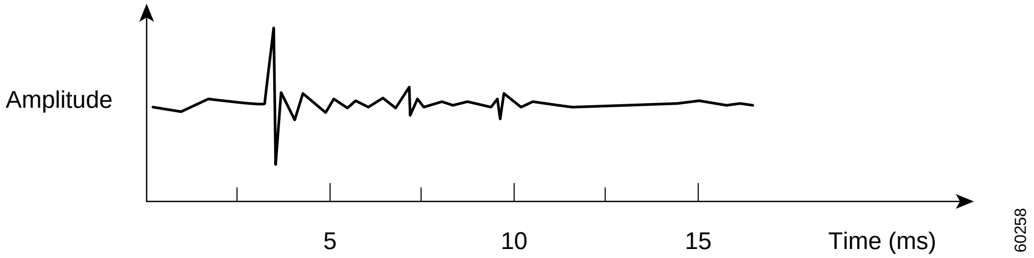

In Figure 6-44, the peaks correspond to individual echoes in the tail circuit. We see that this system has three echoes, a strong one around 3ms and two weaker ones around 7 and 9 msec. After about 12ms, there is no significant energy in the impulse response. An echo canceller facing into such a tail circuit should be provisioned for at least 12ms of tail coverage, to cancel all three echoes. An echo canceller with 5 msec of coverage would do pretty well with this circuit, as the primary echo falls within this window, but the second two echoes would remain un-cancelled, as the echo canceller would be unable to 'see' them. It is important to stress again that the echo canceller faces into a static tail circuit. The tail coverage parameter has nothing to do with the IP network, the round trip delay, or whether the network delay is changing.

Note *Quality of service* (QoS) might improve your end-to-end network delay for a given level of congestion. The shorter the delay, the less annoying a given echo becomes. However, you will never be able to reduce the delay below the danger zone for echo perception with any form of QoS. QoS will help in other ways, but there is no magic QoS switch for echo, nor could there ever be.

One way to ensure that you have a working echo canceller in the circuit is to make a call, and immediately begin to say something repeatedly. The person on the other end of the line should be silent, so if you are calling a voice-mail system, wait for the announcer to stop talking before starting the experiment. If the echo canceller is enabled, and the echoes in the system are cancelable, you will hear echo for the first few utterances and will notice the echo dying away. After a few seconds of speech, the echo should be gone, or at least very quiet compared to the echo level at the beginning of the call. This is the signature of a working echo canceller. Recall that an echo canceller starts out with no knowledge of the tail circuit that it is looking into. It needs to observe a certain amount of speech and echo flowing through the tail circuit to form the virtual tail circuit model. This training period is known as the convergence time of the echo canceller. You should expect convergence within the first few seconds of active speech. If you try this experiment, and do not obtain echo reduction with time, there are two possibilities:

1. The echo canceller is disabled or broken
2. The echo source is un-cancelable.

Try making calls to other destinations and looking for the standard echo-die-away behavior.

Un-Cancelable Echo

An un-cancelable echo is an echo that is either 1) too loud to render inaudible or 2) delayed beyond the time window of the echo canceller's coverage. Too loud an echo can make the echo look like Jane's voice (from the perspective of the echo canceller) or require more attenuation than an echo canceller can give to be made imperceptible. Tail circuits, which involve multiple PSTN hops, some long distance trunks, and alternating series of digital and analog links, can have echoes that exceed the tail coverage window.

Troubleshooting Persistent Echo

If echo is still persistent (lasts throughout the call) and the echo canceller has been demonstrated to be working, this implies that the echo is beyond the ability of the echo canceller to fix. The echo is thus either too loud (more likely) or too delayed (much less likely). There are two problems. One is to identify which is the case, and the second problem is to find out where to fix it.

1. Ensure that the echo canceller is provisioned ON, with coverage set to the max.
2. Match output impedances and levels with analog telecom equipment attached to the gateway's analog voice ports.
3. Identify which tail circuit is the problem
4. If the destination telephone is a speakerphone or headset, you can probably stop right now. Try replacing the speakerphone or headset with a decent handset, and see if the echo dies away normally.
5. Ensure the voice gateway echo canceller is provisioned ON (Skinny Gateways are always set to ON) and coverage is set to maximum. Test for normal echo canceller behavior, i.e. see if echo dies away within a few seconds of speech. Follow the procedure above to verify that echo canceller is working correctly.

In order for an echo canceller to detect echo, the difference between the transmitted voice and the received echo must be at least 10dB for Cisco IOS gateways and 15dB for Skinny gateways. A difference of 15dB or more is optimal. If you are using an IOS Gateway, you can observe the Tx and Rx dB levels by issuing the command "**show call active voice**". Issue this command several times while talking loud and soft, and record the results. If the difference between the Tx and Rx dB level is less than or close to10dB, this is the cause of the echo canceller not converging. In this case, we can try to alleviate the problem by attenuating or padding the signal coming in from the PSTN. If the difference in signal level between Tx and Rx are 15dB or greater, this indicates that the echo is occurring outside of the echo canceller coverage period.

To adjust the amount of padding in or out of a Skinny Gateway (6608, or DT-24/DE-30), Click on **Device -> Gateway**. Find and select the gateway that you would like to adjust, and scroll to the bottom of the configuration. Figure 6-45 illustrates this configuration screen.

Figure 6-45 Example Skinny Gateway Configuration Display

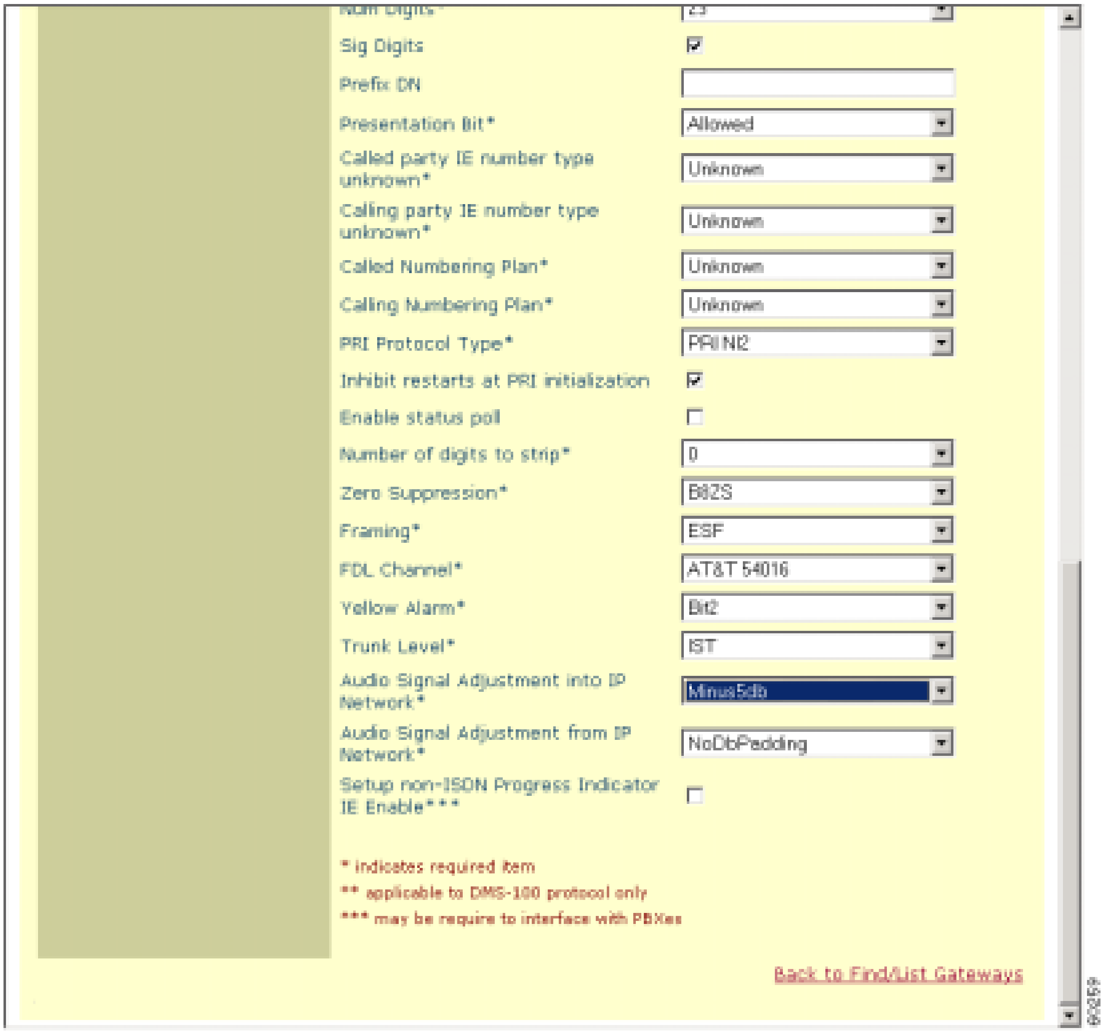

Num Digits*	23
Sig Digits	☑
Prefix DN	
Presentation Bit*	Allowed
Called party IE number type unknown*	Unknown
Calling party IE number type unknown*	Unknown
Called Numbering Plan*	Unknown
Calling Numbering Plan*	Unknown
PRI Protocol Type*	PRI NI2
Inhibit restarts at PRI initialization	☑
Enable status poll	□
Number of digits to strip*	0
Zero Suppression*	B8ZS
Framing*	ESF
FDL Channel*	AT&T 54016
Yellow Alarm*	Bit2
Trunk Level*	0T
Audio Signal Adjustment into IP Network*	Minus5db
Audio Signal Adjustment from IP Network*	NoDbPadding
Setup non-ISDN Progress Indicator IE Enable***	□

* indicates required item
** applicable to DMS-100 protocol only
*** may be require to interface with PBXes

Back to Find/List Gateways

To adjust padding on Cisco IOS gateways, adopt the following configuration settings:

```
voice-port 3/0:23
    input gain -6
    output attenuation 5
    echo-cancel coverage 32
```

If adjusting the gain and attenuation on the gateway does not help, work with your provider to adjust the signal level coming into your gateway.

One-Way Audio or No Audio

One-way audio occurs when one person cannot hear another person during a call. This can be caused by an improperly configured Cisco IOS gateway, a firewall, or a routing or default gateway problem, among other things.

There are a number of causes for one-way audio or no audio during a call. The most common cause is an improperly-configured device. For instance, Cisco CallManager handles the call setup for a Cisco IP phone. The actual audio stream occurs between the two Cisco IP phones (or between the Cisco IP phone and a gateway). So, it is entirely possible that the Cisco CallManager is able to signal to a destination phone (making it ring) when the phone originating the call does not have an IP route to the destination phone. This commonly occurs when the default gateway in the phone is improperly configured (manually or on the Dynamic Host Configuration Protocol (DHCP) server).

If a call consistently has one-way audio, try to ping the destination Cisco IP phone using a PC that is on the same subnet as the phone and has the same default gateway. Take a PC that is on the same subnet as the destination phone (with the same default gateway as the destination phone) and ping the source phone. Both of those tests should work. Audio traffic can also be affected by a firewall or packet filter (such as access lists on a router) that may be blocking the audio in one or both directions. If the one-way audio occurs only through a voice-enabled Cisco IOS gateway, check the configuration carefully. IP routing must be enabled (examine the configuration to make sure that **no ip routing** command is not found near the beginning of the configuration). Also, if you're using RTP header compression to save bandwidth across the WAN, make sure that it is enabled on each router carrying voice traffic that attaches to the WAN circuit. There should not be a situation where the RTP header is compressed on one end but cannot be de-compressed on the other side of the WAN. A sniffer is a very useful tool when troubleshooting one-way audio problems because you can verify that the phone or gateway is actually sending or receiving packets. Diagnostic CDRs are useful in determining if a call is experiencing one-way audio because they log transmitted and received packets (refer to the "Lost or Distorted Audio" section on page 6-117). You can also press the ⓘ button twice (quickly) on a Cisco IP Phone 7960 during an active call to view details about transmitted and received packets.

Note When a call is muted (mute button pressed on a phone), no packets will be transmitted. The Hold button stops the audio stream, so no packets are sent in either direction. When the Hold button is released, all the packet counters are reset. Remember that Silence Suppression must be disabled on both devices for the TX and RX counters to stay equal. Disabling Silence Suppression system-wide will not affect Cisco IOS Gateways.

MTP and One-Way Audio

If you are using Media Termination Point (MTP) in a call (to support supplementary services such as hold and transfer with H.323 devices that do not support H.323 version 2), check to see if the MTP allocated is working correctly. Cisco IOS routers support H.323 version 2 beginning in release 11.3(9)NA and 12.0(3)T. Starting with Cisco IOS release 12.0(7)T, the optional H.323 Open/Close LogicalChannel is supported, so that software-based MTP is no longer required for supplementary services.

The MTP device, as well as Conference Bridge and Transcoder, will bridge two or more audio streams. If the MTP, Conference Bridge, or Transcoder is not working properly, one-way audio or audio loss might be experienced. Shut down MTP to find out if MTP is causing the problem.

Phone Resets

Phones will power cycle or reset for one of the following two reasons:

- TCP failure connecting to Cisco CallManager, or
- Failure to receive an acknowledgement to the phone's KeepAlive messages.

Below are steps for troubleshooting phone resets:

Step 1 Check the Cisco CallManager release notes on CCO (Cisco Connection Online at http://www.cisco.com/) for release notes relating to the problem.

Step 2 Check the Event Viewer for instances of phone(s) resetting. Phone resets are considered Information events, as shown in Figure 6-46.

Figure 6-46 Event Viewer Information Window

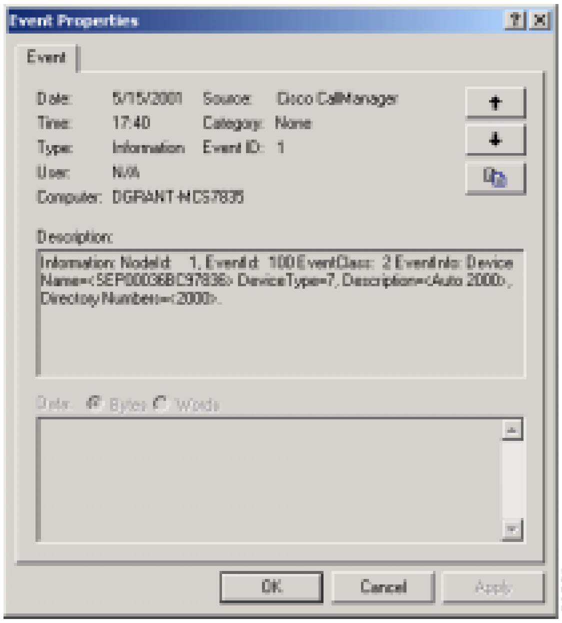

Date: 5/15/2001 Source: Cisco CallManager
Time: 17:40 Category: None
Type: Information Event ID: 1
User: N/A
Computer: DGRANT-MCS7835

Description:
Information: NodeId: 1, EventId: 100 EventClass: 2 EventInfo: Device Name=<SEP00036BC97836> DeviceType=7, Description=<Auto 2000>, Directory Number=<2000>.

Data: ● Bytes ○ Words

Step 3 Look for these and any errors that may have occurred around the time that the phone(s) reset.

Step 4 Start a CallManager trace and try to isolate the problem by identifying any common characteristics in the phones that are resetting. For example, check whether they are all located on the same subnet, same VLAN, and so on. Look at the trace and determine if:

a. The resets occur during a call or happen intermittently, or

b. There are any similarities of phone model – Cisco IP Phone 7960, Cisco IP Phone 30VIP, and so on.

Step 5 Start a Sniffer trace on a phone that frequently resets. After it has reset, look at the trace to determine if there are any TCP retries occurring. If so, this indicates a network problem. The trace may show some consistencies in the resets, such as the phone resetting every seven days. This might indicate a DHCP lease expiration every seven days (this value is user-configurable and could be every two minutes, and so on).

Dropped Calls

Dropped calls occur when a call is prematurely terminated. You can use CDRs to determine the possible cause of dropped calls, particularly if the problem is intermittent. Dropped calls can be the result of a phone or gateway resetting (see above section) or a circuit problem, such as incorrect PRI configuration or error.

The first step is to determine if this problem is isolated to one phone or a group of phones. Perhaps the affected phones are all on a particular subnet or location. The next step is to check the Event Viewer for phone or gateway resets, as illustrated in Figure 6-47.

Figure 6-47 Event Viewer Screen Accessed when Diagnosing Dropped Call Causes

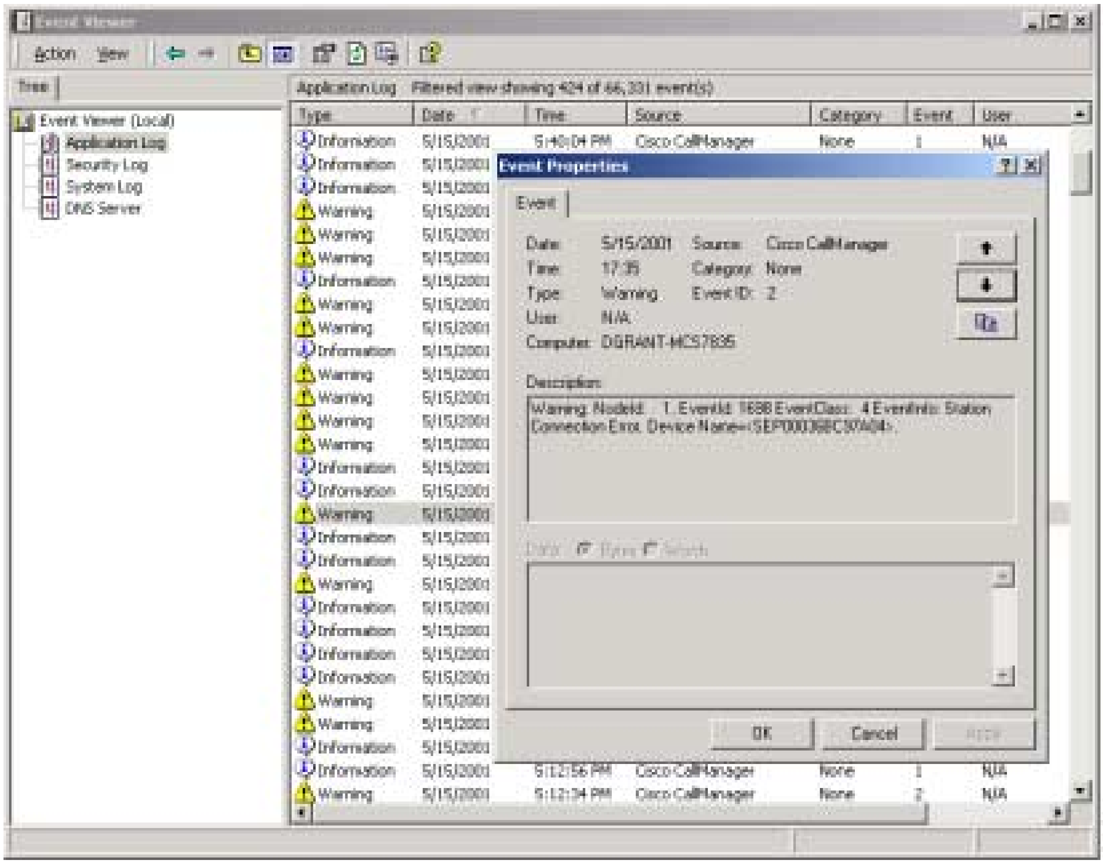

```
Application Log   Filtered view showing 424 of 66,331 event(s)

Type          Date        Time         Source             Category   Event   User
Information   5/15/2001   5:40:04 PM   Cisco CallManager  None       1       N/A
Information   5/15/2001
Information   5/15/2001
Warning       5/15/2001
Warning       5/15/2001
Warning       5/15/2001
Information   5/15/2001
Warning       5/15/2001
Warning       5/15/2001
Information   5/15/2001
Warning       5/15/2001
Warning       5/15/2001
Warning       5/15/2001
Warning       5/15/2001
Information   5/15/2001
Information   5/15/2001
Warning       5/15/2001
Information   5/15/2001
Information   5/15/2001
Warning       5/15/2001
Information   5/15/2001
Information   5/15/2001
Information   5/15/2001
Information   5/15/2001
Warning       5/15/2001
Warning       5/15/2001
Information   5/15/2001
Information   5/15/2001   5:12:56 PM   Cisco CallManager  None       1       N/A
Warning       5/15/2001   5:12:34 PM   Cisco CallManager  None       2       N/A
```

```
Date:      5/15/2001      Source:    Cisco CallManager
Time:      17:35          Category:  None
Type:      Warning        Event ID:  2
User:      N/A
Computer:  DGRANT-MCS7835

Description:
Warning: NodeId:   1, EventId: 1698 EventClass:  4 EventInfo: Station
Connection Error. Device Name=<SEP003068C97A04>.

Data:  ● Bytes  ○ Words
```

There should be one Warning and one Error message for each phone that resets. In this case, the problem is often that the phone cannot keep its TCP connection to the Cisco CallManager alive, so the Cisco CallManager resets the connection. This may be because a phone was turned off or there may be a problem in the network. If this is an intermittent problem, it may be useful to use Microsoft Performance to record phone registrations, as illustrated in Figure 6-48.

Figure 6-48 Microsoft Performance Screen Used for Verifying Phone Registration

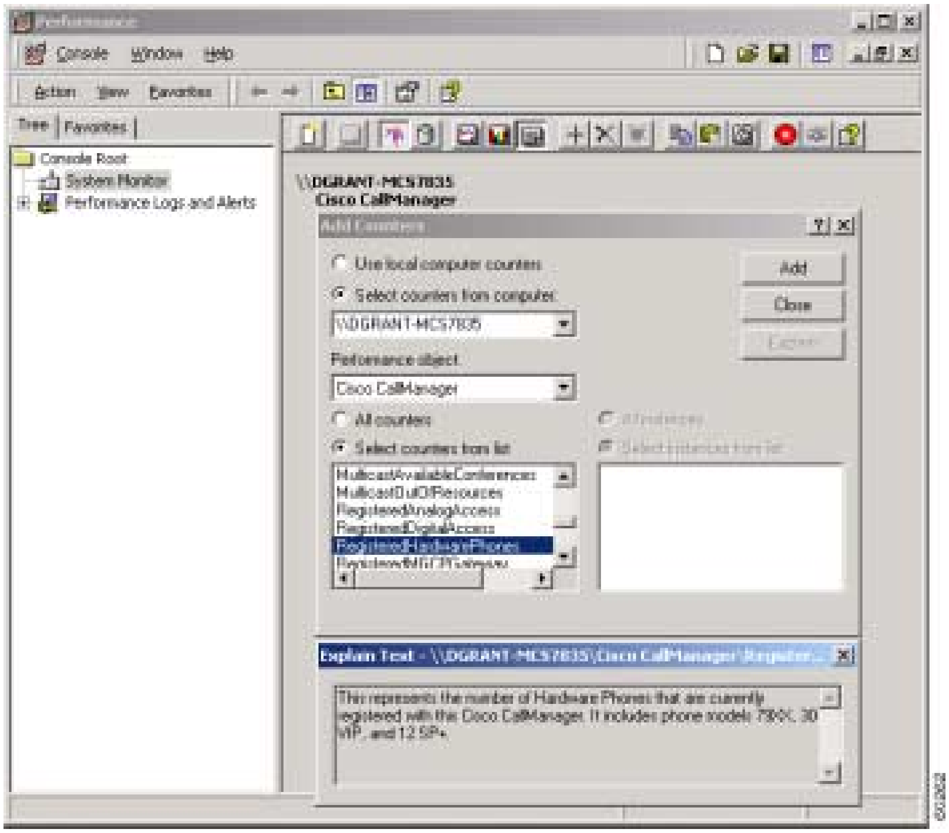

○ Use local computer counters
● Select counters from computer:
\\DGRANT-MCS7835

Performance object:
Cisco CallManager

○ All counters
● Select counters from list
MulticastAvailableConferences
MulticastOutOfResources
RegisteredAnalogAccess
RegisteredDigitalAccess
RegisteredHardwarePhones
RegisteredMGCPGateway

Explain Text - \\DGRANT-MCS7835\Cisco CallManager\Register...

This represents the number of Hardware Phones that are currently registered with the Cisco CallManager. It includes phone models 79XX, 30 VIP, and 12 SP+.

If the problem seems to be occurring only through a certain gateway, such as a Cisco Access DT-24+, the best course of action is to enable tracing and/or view the CDRs. The CDR files will give a Cause Of Termination (COT) that may help determine the cause of the problem. Figure 6-49 illustrates viewing the CDR files in search of the COT.

Figure 6-49 CDR File View Example for COT Assessment

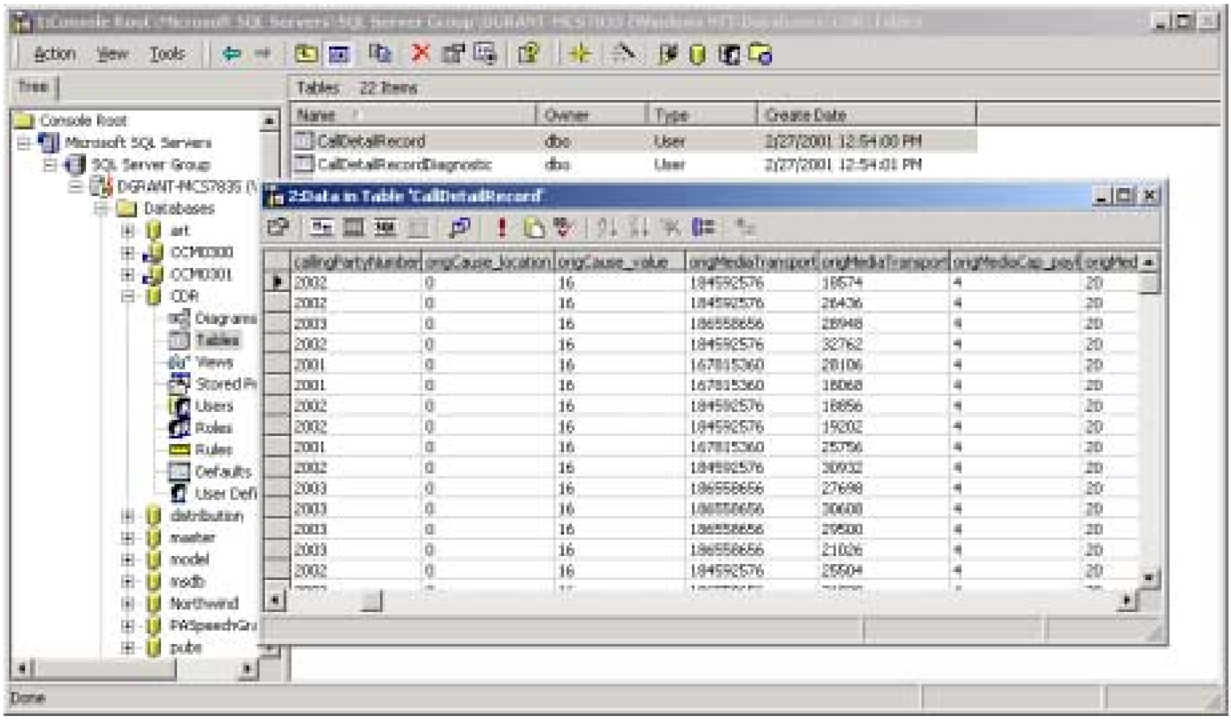

Tables 22 Items

Name	Owner	Type	Create Date
CallDetailRecord	dbo	User	2/27/2001 12:54:00 PM
CallDetailRecordDiagnostic	dbo	User	2/27/2001 12:54:01 PM

callingPartyNumber	origCause_location	origCause_value	origMediaTransport	origMediaTransport	origMediaCap_payl	origMed
2002	0	16	184592576	18574	4	20
2002	0	16	184592576	26436	4	20
2003	0	16	186558656	28948	4	20
2002	0	16	184592576	32762	4	20
2001	0	16	167815360	28106	4	20
2001	0	16	167815360	18068	4	20
2002	0	16	184592576	16856	4	20
2002	0	16	184592576	19202	4	20
2001	0	16	167815360	25756	4	20
2002	0	16	184592576	30932	4	20
2003	0	16	186558656	27698	4	20
2003	0	16	186558656	30608	4	20
2003	0	16	186558656	29500	4	20
2003	0	16	186558656	21026	4	20
2002	0	16	184592576	25504	4	20

The disconnect cause values (origCause_value and destCause_value — depending on which side hung up the call), map to Q.931 disconnect cause codes (in decimal) that can be found at http://www.cisco.com/univercd/cc/td/doc/product/software/ios113ed/dbook/disdn.htm. In Figure 6-49, cause 16 refers to a normal call clearing. If the call is going out a gateway to the PSTN, the CDR can be used to determine which side is hanging up the call. Much of the same information can be obtained by enabling tracing on the Cisco CallManager. Use the trace tool only as a last resort or if the network is not yet in production.

Check Your Loads

As with any problem, check the phone and gateway loads and CCO (Cisco Connection Online at http://www.cisco.com/) for the latest software loads, new patches, or release notes relating to the problem.

Cisco CallManager Feature Issues

Problems may occur with features, such as Conference Bridge or Media Termination Point, which are used in conjunction with Cisco CallManager. Some of these feature problems are caused by configuration errors or a lack of resources. For example, users may not be able to conference calls if the specified number of Ad Hoc conference resources has been exceeded. The result would be a dropped call when the user attempted to initiate the conference feature. This could appear to be a Cisco CallManager feature issue, when in fact it is a problem with the number of available conference resources. The number of times a conference resource was required, but not available, is one of the counters logged in Microsoft Performance. The same behavior occurs if there are conference resources available, but the conferencing service had stopped.

Codec/Regions: Codec Mismatch

If a user gets a reorder tone when going off-hook, it could be the result of codec disagreement between regions. Verify that both call ends support at least one common codec (for example, G.711). If not, you will need to use transcoders.

A region specifies the range of supported codecs that can be used with each of the other regions. Every device belongs to a region.

Note Codec negotiation with a Cisco IOS router is not supported.

Region1<->Region2 = G.711 means that a call between a device in Region1 and a device in Region2 can use G.711 or any other supported codec that requires the same or less bandwidth as G.711 (any supported codecs within G.711, G.729, G.723, and so on).

The following codecs are supported for each device:

- Cisco IP Phone 79xx — G.711A-law/µ-law, G.729
- Cisco IP Phone SP12 series and VIP 30 — G.711A-law/µ-law, G.723.1
- Cisco Access Gateway DE30 and DT-24+ — G.711A-law/µ-law, G.723.1
- Cisco Catalyst WS-X6608-T1/E1 — G.711A-law//µ-law, G.729

Locations

If a user receives a reorder tone after dialing a number, it could be because the Cisco CallManager bandwidth allocation for the location of one of the call end devices has been exceeded (less than 24k). Cisco CallManager checks for 24k available bandwidth for each device before making a call. If less than 24k bandwidth is available, Cisco CallManager will not setup the call and the user will hear a reorder tone.

```
12:42:09.017 Cisco CallManager|Locations: Orig=1 BW=12 Dest=0 BW=-1 (-1 implies infinite
bw available)
12:42:09.017 Cisco CallManager|StationD - stationOutputCallState tcpHandle=0x4f1ad98
12:42:09.017 Cisco CallManager|StationD - stationOutputCallInfo CallingPartyName=,
CallingParty=5003, CalledPartyName=, CalledParty=5005, tcpHandle=0x4f1ad98
12:42:09.017 Cisco CallManager|StationD - stationOutputStartTone: 37=ReorderTone
tcpHandle=0x4f1ad98
```

Once the call is established, the Cisco CallManager subtracts bandwidth from the locations depending on the codec used in that call. If the call is using G.711, Cisco CallManager subtracts 80k; if the call is using G.723, Cisco CallManager subtracts 24k; if the call is using G729, Cisco CallManager subtracts 24k.

Conference Bridge

Use the following information to help troubleshooting a *No Conference Bridge Available* problem. This could be either a software or a hardware problem.

First, check to see if you have any available Conference Bridge resources registered with Cisco CallManager (either software or hardware). To do so, you can use Microsoft Performance to check the number of *Unicast AvailableConferences*.

The Cisco IP Voice Media Streaming application performs the conference bridge function. One software installation of Cisco IP Voice Media Streaming will support 16 Unicast Available Conferences (3 people/conference), as shown in the following trace.

```
10:59:29.951 Cisco CallManager|UnicastBridgeControl - wait_capabilities_StationCapRes -
Device= CFB_kirribilli - Registered - ConfBridges= 16, Streams= 48, tcpHandle=4f12738
10:59:29.951 Cisco CallManager|UnicastBridgeManager - UnicastBridgeRegistrationReq -
Device Registration Complete for Name= Xoð_ô%ð_ - DeviceType= 50, ResourcesAvailable= 16,
deviceTblIndex= 0
```

One E1 port (WS-X6608-E1 card contains 8x E1 ports) provides five Unicast Available Conferences (max conference size = 6), as shown in the following trace.

```
11:14:05.390 Cisco CallManager|UnicastBridgeControl - wait_capabilities_StationCapRes -
Device= CFB00107B000FB0 - Registered - ConfBridges= 5, Streams= 16, tcpHandle=4f19d64
11:14:05.480 Cisco CallManager|UnicastBridgeManager - UnicastBridgeRegistrationReq -
Device Registration Complete for Name= Xoð_ô%ð_ - DeviceType= 51, ResourcesAvailable= 5,
deviceTblIndex= 0
```

The following hardware trace on the Cisco Catalyst 6000 8 Port Voice T1/E1 and Services Module indicates that the E1 port 4/1 in the card has registered as a Conference Bridge with Cisco CallManager.

```
greece-sup (enable) sh port 4/1
Port Name Status Vlan Duplex Speed Type
----- ------------------ ---------- ---------- ------ ----- ------------
4/1 enabled 1 full - Conf Bridge
Port DHCP MAC-Address IP-Address Subnet-Mask
-------- ------- ----------------- --------------- ---------------
4/1 disable 00-10-7b-00-0f-b0 10.200.72.31 255.255.255.0
Port Call-Manager(s) DHCP-Server TFTP-Server Gateway
-------- ----------------- --------------- --------------- ---------------
4/1 10.200.72.25 - 10.200.72.25 -
Port DNS-Server(s) Domain
-------- ----------------- -------------------------------------------------
4/1 - 0.0.0.0
Port CallManagerState DSP-Type
-------- ---------------- --------
4/1 registered C549
Port NoiseRegen NonLinearProcessing
----- ---------- -------------------
4/1 disabled disabled
```

Second, check the maximum number of users configured in the conference (Ad Hoc or Meet-Me) to determine if the problem occurred because this number was exceeded.

Transcoding Problems

If you have installed a hardware transcoder in the Cisco Catalyst 6000 8 Port Voice T1/E1 and Services Module, and it doesn't work as expected (meaning you cannot make calls between two users with no common codec), check to see if you have any available Transcoder resources registered with Cisco CallManager (this must be hardware). Use Microsoft Performance to check the number of `MediaTermPointsAvailable` available.

One E1 port (WS-X6608-E1 card contains 8x E1 ports) provides Transcoder/MTP resources for 16 calls, as shown in the following trace.

```
11:51:09.939 Cisco CallManager|MediaTerminationPointControl - Capabilities Received -
Device= MTP00107B000FB1 - Registered - Supports 16 calls
```

The following hardware trace on the Cisco Catalyst 6000 8 Port Voice T1/E1 and Services Module indicates that the E1 port 4/2 in the card has registered as an MTP/transcoder with Cisco CallManager.

```
greece-sup (enable) sh port 4/2
Port Name Status Vlan Duplex Speed Type
----- ------------------ ---------- ---------- ------ ----- ------------
4/2 enabled 1 full - MTP
Port DHCP MAC-Address IP-Address Subnet-Mask
```

```
-------- ------- ---------------- -------------- ---------------
4/2 disable 00-10-7b-00-0f-b1 10.200.72.32 255.255.255.0
Port Call-Manager(s) DHCP-Server TFTP-Server Gateway
-------- ---------------- -------------- -------------- --------------
4/2 10.200.72.25 - 10.200.72.25 -
Port DNS-Server(s) Domain
-------- ---------------- -------------------------------------------------
4/2 - 0.0.0.0
Port CallManagerState DSP-Type
-------- --------------- --------
4/2 registered C549
Port NoiseRegen NonLinearProcessing
----- ---------- ------------------
4/2 disabled disabled
```

Note The same T1/E1 port cannot be configured for both Conference Bridge and Transcoder/MTP

In order to make a call between two devices using a low bit rate code (such as G.729 and G.723) that do not support the same codec, a transcoder resource is required. Figure 6-50 provides a basis discussing transcoder implementation.

Figure 6-50 Transcoder Implementation Environment

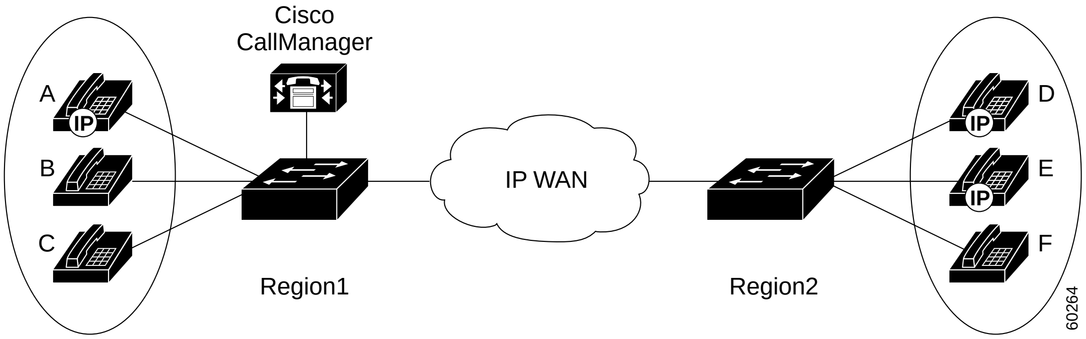

Assume Cisco CallManager has been configured such that the codec between Region 1 and Region 2 in Figure 6-50 is G.729. The following scenarios are possible:

- If caller on Phone A initiates a call, Cisco CallManager realizes that it is a Cisco IP Phone 7960, which happens to support G.729. After the digits have been collected, the Cisco CallManager determines that the call is destined for User D who is in Region 2. Since the destination device also supports G.729, the call is set up and the audio flows directly between Phone A and Phone D.
- If a caller on Phone B, who has a Cisco IP Phone 12SP+, were to initiate a call to Phone D, this time the Cisco CallManager would realize that the originating phone only supports G.723 or G.711. Cisco CallManager would need to allocate a transcoding resource so that audio would flow as G.711 between Phone B and the transcoder, but as G.729 between the transcoder and Phone D. If no transcoder were available, Phone D's phone would ring, but as soon as the call was answered, the call would disconnect.
- If a user on Phone B were to call Phone F (Cisco IP Phone 12SP+), the two phones would actually use G.723, even though G.729 is configured as the codec to use between the regions. G.723 is used because both endpoints support it and it uses less bandwidth than G.729.
- If a Cisco uOne voice mail system is added (which only supports G.711) or a Cisco IOS router configured for G.711 to Region 1, then a transcoding device must be used if calling from Region 2. If none is available, then the call will fail.

MTP Resource Problems

An MTP resource problem could be the culprit if a call is established and supplementary services are not available on an H.323 device that does not support H323v2. First, determine whether you have any available MTP resources (either software or hardware) registered with Cisco CallManager. You can do so by using Microsoft Performance to check the number of `MediaTermPointsAvailable`.

One MTP software application supports 24 calls (using MTP to support supplementary services with H.323 devices that not support H.323v2), as shown in the following trace.

```
10:12:19.161 Cisco CallManager|MediaTerminationPointControl - Capabilities Received -
Device= MTP_kirribilli. - Registered - Supports 24 calls
```

One E1 port (WS-X6608-E1 card contains 8x E1 ports) provides MTP resources for 16 calls, as shown in the following trace.

```
11:51:09.939 Cisco CallManager|MediaTerminationPointControl - Capabilities Received -
Device= MTP00107B000FB1 - Registered - Supports 16 calls
```

The following hardware trace, from the Cisco Catalyst 6000 8 Port Voice T1/E1 and Services Module, indicates that the E1 port 4/2 in the card has registered as an MTP/transcoder with Cisco CallManager.

```
greece-sup (enable) sh port 4/2
Port Name Status Vlan Duplex Speed Type
----- ------------------ ---------- ---------- ------ ----- ------------
4/2 enabled 1 full - MTP
Port DHCP MAC-Address IP-Address Subnet-Mask
-------- ------- ----------------- --------------- ---------------
4/2 disable 00-10-7b-00-0f-b1 10.200.72.32 255.255.255.0
Port Call-Manager(s) DHCP-Server TFTP-Server Gateway
-------- ----------------- --------------- --------------- ---------------
4/2 10.200.72.25 - 10.200.72.25 -
Port DNS-Server(s) Domain
-------- ----------------- -------------------------------------------------
4/2 - 0.0.0.0
Port CallManagerState DSP-Type
-------- ---------------- --------
4/2 registered C549
Port NoiseRegen NonLinearProcessing
----- ---------- ------------------
4/2 disabled disabled
```

Second, see if the **Media Termination Point Required** check box is selected in the Gateway Configuration screen of Cisco CallManager Administration.

Third, verify that Cisco CallManager has allocated the required number of MTP devices.

From the CCM file:

```
15:22:23.848 Cisco CallManager|MediaManager(40) started
15:22:23.848 Cisco CallManager|MediaManager - wait_AuConnectRequest
15:22:23.848 Cisco CallManager|MediaManager - wait_AuConnectRequest - Transcoder Enabled
15:22:23.848 Cisco CallManager|MediaManager - wait_AuConnectRequest - party1(16777357),
party2(16777358), proxies=1, connections=2, current proxies=0
15:22:23.848 Cisco CallManager|MediaManager - wait_AuConnectRequest - proxy connections
15:22:23.848 Cisco CallManager|MediaManager - wait_AuConnectRequest - allocating
MTP(ci=16777359)
15:22:23.848 Cisco CallManager|MediaManager - wait_AllocateMtpResourceRes
15:22:23.848 Cisco CallManager|MediaManager - wait_AllocateMtpResourceRes - start 2
connections
15:22:23.848 Cisco CallManager|MediaManager - wait_AllocateMtpResourceRes - creating
connection between party1(16777357) and party2(16777359)
15:22:23.848 Cisco CallManager|MediaManager - wait_AllocateMtpResourceRes - creating
connection between party1(16777358) and party2(16777359)
```

```
15:22:23.848 Cisco CallManager|MediaCoordinator - wait_MediaCoordinatorAddResource -
CI=16777359 count=1
15:22:23.848 Cisco CallManager|MediaCoordinator - wait_MediaCoordinatorAddResource -
CI=16777359 count=2
```

Dial Plans

A Dial Plan is a list of numbers (and groups of numbers) that tells the Cisco CallManager to which devices (phones, gateways, and so on) to send calls when a certain string of digits is collected. It is analogous to a static routing table in a router. Please be certain that your dial plan concepts, basic call routing, and planning have been carefully considered and properly configured before trying to troubleshoot a potential dial plan issue. Very often, the problem lies with planning and configuration.

Consider the following questions when troubleshooting dial plans problems:

- What is the Directory Number (DN) originating the call?
- What is the Calling Search Space of this DN?
- If applicable, what is the Calling Search Space of the device (such as a Cisco IP phone) with which the DN is associated? Make sure that you identify the correct device; multiple line appearances are supported, and it's possible to have a DN on multiple devices. Note the device's Calling Search Space. If the call is originated by a Cisco IP phone, remember that the particular line (DN) and the device to which that line is associated will each have Calling Search Spaces. They will be combined when making a call. As an example, assume that line instance 1000 has a Calling Search Space of AccessLevelX and the Cisco IP phone that has extension 1000 configured on it has AccessLevelY as its Calling Search Space. Therefore, when making a call from that line appearance, Cisco CallManager will search through partitions contained in Calling Search Space AccessLevelX and AccessLevelY.
- What partitions are associated with the Calling Search Space(s)?
- What is the partition of the device to which the call should (or should not) go?
- What is the number that is being dialed? Note if and when the callers are getting a secondary dial tone, at any stage. Also, what do callers hear after all the digits have been entered (re-order, fast-busy)? Do they get the progress tones before they expect to hear anything? Make sure callers wait at least 10 seconds after entering the last digit, since they may have to wait for the inter-digit timer to expire.
- Generate a Route Plan Report in Cisco CallManager Administration. Use it to examine all the route patterns for the partitions that are in the Calling Search Space for the call.
- If necessary, add or modify the Route Patterns or Route Filters.
- If you can find the Route Pattern to which the call is being sent, note the Route List or Gateway to which the pattern points.
- For a Route List, check which Route Groups are part of the list and which Gateway(s) are part of the Route Groups.
- Verify that the applicable devices are registered with Cisco CallManager.
- If there is no access to Cisco CallManager, use the **show tech** command to capture this information and verify.
- Watch out for the @ sign. This is a macro that can expand to include many different things. It is often used in combination with filtering options.
- If a device isn't part of a partition, it is said to be part of the Null or default partition. **Every** user should be able to call that device. The Null partition is always searched **last**.

- If you dial an outside number that is matching a 9.@ pattern and it takes 10 seconds before the call goes through, check the filtering options. A 9.@ pattern, when dialing a 7-digit number, **will** (by default) wait 10 seconds. You need to apply a Route Filter to the pattern that says LOCAL-AREA-CODE DOES-NOT- EXIST and END-OF-DIALING DOES-NOT-EXIST.

Partitions

Route partitions inherit the error handling capabilities of the Cisco CallManager software. That is, a console and CCM file trace are provided for logging information and error messages. These messages will be part of the digit analysis component of the traces. Even with the traces below, an understanding of the Partitions and Calling Search Spaces configurations and which devices are in each partition, along with its associated calling search space, is vital in determining the problem.

The trace below is an example of a number dialed that is in the Calling Search Space of the device. For more detailed explanations about CallManager traces, please review the case studies in this document.

```
08:38:54.968 Cisco CallManager|StationInit - InboundStim - OffHookMessageID
tcpHandle=0x6b88028
08:38:54.968 Cisco CallManager|StationD - stationOutputDisplayText tcpHandle=0x6b88028,
Display= 5000
08:38:54.968 Cisco CallManager|StationD - stationOutputSetLamp stim: 9=Line instance=1
lampMode=LampOn tcpHandle=0x6b88028
08:38:54.968 Cisco CallManager|StationD - stationOutputCallState tcpHandle=0x6b88028
08:38:54.968 Cisco CallManager|StationD - stationOutputDisplayPromptStatus
tcpHandle=0x6b88028
08:38:54.968 Cisco CallManager|StationD - stationOutputSelectSoftKeys tcpHandle=0x6b88028
08:38:54.968 Cisco CallManager|StationD - stationOutputActivateCallPlane
tcpHandle=0x6b88028|
08:38:54.968 Cisco CallManager|Digit analysis: match(fqcn="5000", cn="5000",
pss="RTP_NC_Hardwood:RTP_NC_Woodland:Local RTP", dd="")
```

In the Digit Analysis component of the trace (above), the `pss` (Partition Search Space, also known as Calling Search Space) is listed for the device placing the call. Below, you can see that `RTP_NC_Hardwood;RTP_NC_Woodland;Local_RTP` are the partitions this device is allowed to call.

```
08:38:54.968 Cisco CallManager|Digit analysis: potentialMatches=PotentialMatchesExist
08:38:54.968 Cisco CallManager|StationD - stationOutputStartTone: 33=InsideDialTone
tcpHandle=0x6b88028
08:38:55.671 Cisco CallManager|StationInit - InboundStim - KeypadButtonMessageID kpButton:
5 tcpHandle=0x6b88028
08:38:55.671 Cisco CallManager|StationD - stationOutputStopTone tcpHandle=0x6b88028
08:38:55.671 Cisco CallManager|StationD - stationOutputSelectSoftKeys tcpHandle=0x6b88028
08:38:55.671 Cisco CallManager|Digit analysis: match(fqcn="5000", cn="5000",
pss=]RTP_NC_Hardwood:RTP_NC_Woodland:Local RTP", dd="5")
08:38:55.671 Cisco CallManager|Digit analysis: potentialMatches=PotentialMatchesExist
08:38:56.015 Cisco CallManager|StationInit - InboundStim - KeypadButtonMessageID kpButton:
0 tcpHandle=0x6b88028
08:38:56.015 Cisco CallManager|Digit analysis: match(fqcn="5000", cn="5000",
pss="RTP_NC_Hardwood:RTP_NC_Woodland:Local RTP", dd="50")
08:38:56.015 Cisco CallManager|Digit analysis: potentialMatches=PotentialMatchesExist
08:38:56.187 Cisco CallManager|StationInit - InboundStim - KeypadButtonMessageID kpButton:
0 tcpHandle=0x6b88028
08:38:56.187 Cisco CallManager|Digit analysis: match(fqcn="5000", cn="5000",
pss="RTP_NC_Hardwood:RTP_NC_Woodland:Local RTP", dd="500")
08:38:56.187 Cisco CallManager|Digit analysis: potentialMatches=PotentialMatchesExist
08:38:56.515 Cisco CallManager|StationInit - InboundStim - KeypadButtonMessageID kpButton:
3 tcpHandle=0x6b88028
08:38:56.515 Cisco CallManager|Digit analysis: match(fqcn="5000", cn="5000",
pss="RTP_NC_Hardwood:RTP_NC_Woodland:Local RTP", dd="5003")
08:38:56.515 Cisco CallManager|Digit analysis: analysis results
08:38:56.515 Cisco CallManager||PretransformCallingPartyNumber=5000
```

From the example above, it is key that you notice that `PotentialMatchesExist` is the result of a Digit Analysis of the numbers that were dialed until the exact match is found and the call is routed accordingly.

Below is a trace where the number that was attempted to be dialed (1001) is not in the Calling Search Space of the device. Again, it is key that you note that the digit analysis routine had potential matches until only the first digit was dialed. The route pattern associated with the digit `1` is in a partition that is not in the device's calling search space, “RTP_NC_Hardwood;RTP_NC_Woodland;Local_RTP”. Therefore the phone was sent Reorder Tone.

```
08:38:58.734 Cisco CallManager|StationInit - InboundStim - OffHookMessageID
tcpHandle=0x6b88028
08:38:58.734 Cisco CallManager|StationD - stationOutputDisplayText tcpHandle=0x6b88028,
Display= 5000
08:38:58.734 Cisco CallManager|StationD - stationOutputSetLamp stim: 9=Line instance=1
lampMode=LampOn tcpHandle=0x6b88028
08:38:58.734 Cisco CallManager|StationD - stationOutputCallState tcpHandle=0x6b88028
08:38:58.734 Cisco CallManager|StationD - stationOutputDisplayPromptStatus
tcpHandle=0x6b88028
08:38:58.734 Cisco CallManager|StationD - stationOutputSelectSoftKeys tcpHandle=0x6b88028
08:38:58.734 Cisco CallManager|StationD - stationOutputActivateCallPlane
tcpHandle=0x6b88028
08:38:58.734 Cisco CallManager|Digit analysis: match(fqcn="5000", cn="5000",
pss="RTP_NC_Hardwood:RTP_NC_Woodland:Local RTP", dd="")
08:38:58.734 Cisco CallManager|Digit analysis: potentialMatches=PotentialMatchesExist
08:38:58.734 Cisco CallManager|StationD - stationOutputStartTone: 33=InsideDialTone
tcpHandle=0x6b88028
08:38:59.703 Cisco CallManager|StationInit - InboundStim - KeypadButtonMessageID kpButton:
1 tcpHandle=0x6b88028
08:38:59.703 Cisco CallManager|StationD - stationOutputStopTone tcpHandle=0x6b88028
08:38:59.703 Cisco CallManager|StationD - stationOutputSelectSoftKeys tcpHandle=0x6b88028
08:38:59.703 Cisco CallManager|Digit analysis: match(fqcn="5000", cn="5000",
pss="RTP_NC_Hardwood:RTP_NC_Woodland:Local RTP", dd="1")
08:38:59.703 Cisco CallManager|Digit analysis: potentialMatches=NoPotentialMatchesExist
08:38:59.703 Cisco CallManager|StationD - stationOutputStartTone: 37=ReorderTone
tcpHandle=0x6b88028
```

Route partitions work by associating a partition name with every directory number in the system. The directory number can be called only if the calling device contains the partition within a list of partitions to which it is permitted to place calls — its partition search space. This provides extremely powerful control over routing.

When a call is being placed, Digit Analysis attempts to resolve the dialed address only in those partitions that the partition search space specifies. Each partition name comprises a discrete subset of the global dialable address space. From each listed partition, Digit Analysis retrieves the pattern that best matches the sequence of dialed digits. Then, from among the matching patterns, Digit Analysis chooses the best match. If two patterns equally match the sequence of dialed digits, Digit Analysis selects the pattern associated with the partition listed first in the partition search space (for more information, review the documentation about Closest-Match Routing).

Security

Cisco CallManager can be configured to create a secure dialing plan for users. This can be done through the use of partitions and calling search spaces, in addition to more common filtering based on sections of the `@` macro (which stands for the North American Numbering Plan) in a route pattern, such as the Area Code. Partitions and Calling Search Spaces are an integral part of security and are especially useful for multi-tenant environments and creating an individual user level. Filtering is a subset of the Calling Search Space/Partition concept that can add additional granularity to the security plan.

This is an extension to the Dial Plan section, above. Be advised, it is not advisable to run an CallManager trace when trying to fix a filtering problem. There is not enough information and the potential for causing additional harm is too great.

Run the **show tech** command on Cisco CallManager. Table 6-32 lists information that appears in the Route Filter section.

Table 6-32 Route Filter Section Output from Show Tech Command

Name	dialPlanWizardG	Clause
CiscoDallasInte	1	(INTERNATIONAL-
CiscoRTPTollByP	1	(AREA-CODE == 9
CiscoRTPLongDis	1	(AREA-CODE EXIS
CiscoDallasToll	1	(AREA-CODE == 9
CiscoDallas911R	1	(SERVICE == 911
CiscoRTPLocal7D	1	(AREA-CODE DOES
CiscoDallasLong	1	(AREA-CODE EXIS
CiscoRTP911RF	1	(SERVICE == 911
CiscoRTPInterna	1	(INTERNATIONAL-
CiscoDallasLoca	1	(LOCAL-AREA-COD

Unfortunately, this display is incomplete. It does, however, give a listing of all the Route Filters in the system. The **show** command does not allow you to see which filters are associated with which Route Pattern. Another method to better understand dial plan is to go to the Route Plan Report window. Figure 6-51 illustrates an option on the far right-hand side to `View In File`. The output shown in Table 6-33 is a comma-separated file that can be viewed in Microsoft Excel or a similar application.

Figure 6-51 Route Plan Report Window—View In File Option

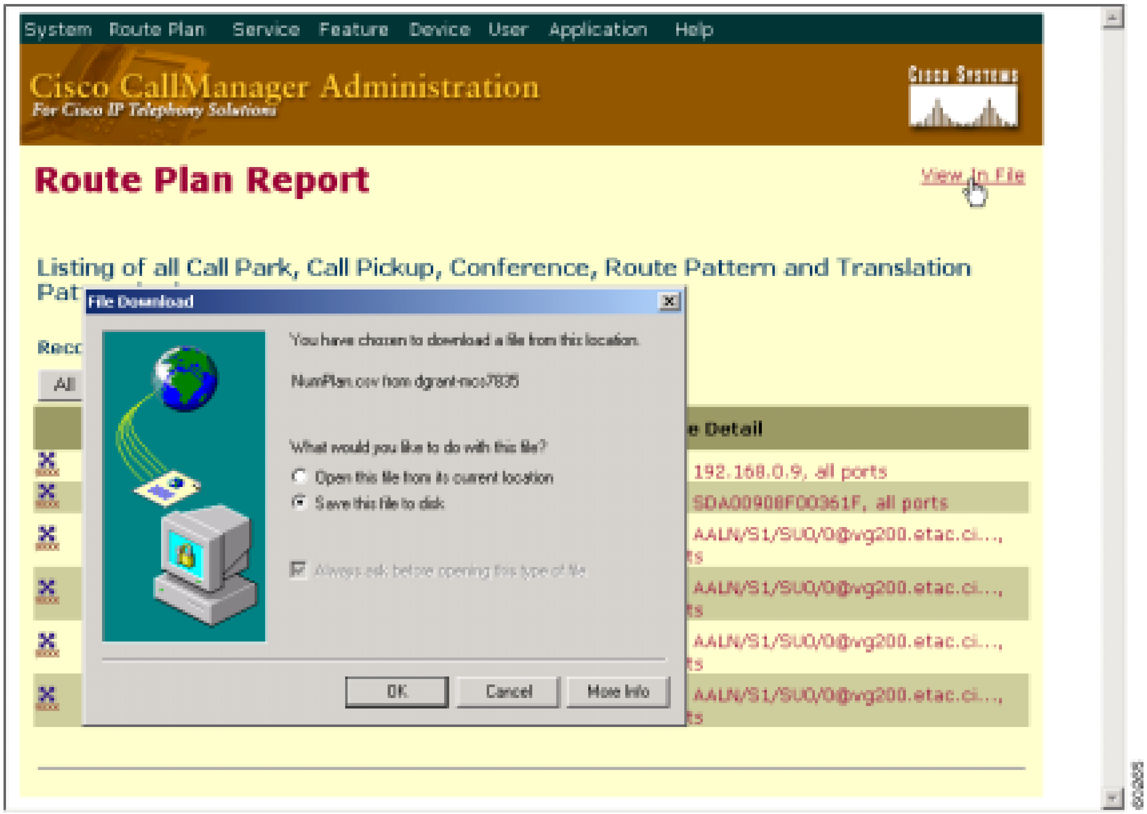

Route Plan Report

Listing of all Call Park, Call Pickup, Conference, Route Pattern and Translation Pat

Reco

File Download

You have chosen to download a file from this location.

NumPlan.csv from dgrant-mcs7835

What would you like to do with this file?

○ Open this file from its current location

● Save this file to disk

☑ Always ask before opening this type of file

192.168.0.9, all ports

SDA00908F00361F, all ports

AALN/S1/SU0/0@vg200.etac.ci...,
ts

AALN/S1/SU0/0@vg200.etac.ci...,
ts

AALN/S1/SU0/0@vg200.etac.ci...,
ts

AALN/S1/SU0/0@vg200.etac.ci...,
ts

Table 6-33 Route Plan Report

Pattern/DN	Partition	Pattern Usage	Device Name	Device Description
1000		Device	SEP003094C2635E	Telecaster
1010		Device	SEP003094C2635E	Telecaster
1111		Device	SEP00308062CDF1	SEP00308062CDF1
1211		Device	SEP00308062CDF1	SEP00308062CDF1
2999		Device	SAA0010EB007FFE	SAA0010EB007FFE
4444		Device	SEP003094C26302	Guest
4500		Conference		
9.@	CiscoRTPLocalPT	Route	CiscoRTPLocalRL	
9.@	CiscoDallasLocalPT	Route	CiscoDallasLocalRL	
9.@	CiscoRTPIntlPT	Route	CiscoRTPIntlRL	
9.@	CiscoDallasLongDistPT	Route	CiscoDallasLongDistRL	
9.@	CiscoRTP911PT	Route	CiscoRTP911RL	
9.@	CiscoRTPLongDistPT	Route	CiscoRTPLongDistRL	
9.@	CiscoTollByPassToDallasPT	Route	CiscoTollByPassToDallasRL	
9.@	CiscoDallasIntlPT	Route	CiscoDallasIntlRL	

Table 6-33 Route Plan Report (continued)

Pattern/DN	Partition	Pattern Usage	Device Name	Device Description
9.@	CiscoDallas911PT	Route	CiscoDallas911RL	
9.@	CiscoTollByPassToRTPPT	Route	CiscoTollByPassToRTPRL	

The output in Table 6-33 shows the route patterns and their corresponding partitions. It does not show the route filters or the calling search spaces of the directory numbers. More information is available on the actual Route Plan Report. If you need to contact the Cisco TAC, you should send this page via email (if the Cisco CallManager is inaccessible).

Slow Server Response

Slow server response can result from several possible problems. These include:

- **Mismatched Duplex—**Slow response from the server could result if the duplex of the switch does not match the duplex of the Cisco CallManager server. For optimal performance, set both the switch and the server to **100/Full**. We do not recommend using *Auto* on either the switch or the server. You must restart the MCS server for the change to take effect.
- **Screen Saver—**Some screen savers, particularly OpenGL screen savers, will consume all of the CPU when active. Cisco recommends disabling screen savers on the CallManager.
- **Third Party Software**—Third party software is ***not*** supported, and is strongly discouraged. Software such as Virus Scanners will consume Cisco CallManager resources, and can cause other problems also. The ***only*** software that should ever be installed on a Cisco CallManager should come from the Cisco CallManager Software Center on CCO. All of the software available on this page has been tested, and is supported by Cisco. Service packs and HotFixes for Windows 2000 should always be downloaded from this page also. Refer to the following location for detailed CallManager information: http://www.cisco.com/cgi-bin/tablebuild.pl/callmgr.

Reorder Tone Through Gateways

Users placing a call through the gateway might get a reorder tone if they are attempting to make a restricted call or call a number that has been blocked. A reorder tone may occur if the dialed number is out of service or if the PSTN has an equipment or service problem. Be sure the device giving the reorder tone has registered. Also, check your dial plan configuration to ensure that the call can be successfully routed.

Steps for troubleshooting reorder tones through gateways:

1. Check the gateways to ensure that you are using the latest software loads.
2. Check CCO (Cisco Connection Online at http://www.cisco.com/) for the latest software loads, new patches, or release notes relating to the problem.
3. Start a CallManager trace and recreate the problem. Reorder tones could be the result of a configuration issue with location-based admission control or gatekeeper-based admission control where the Cisco CallManager might limit the number of allowable calls. In the CallManager trace, locate the call to determine if it was blocked intentionally by a route pattern or the calling search space, or by any other configuration setting.
4. Reorder tones can also occur when calling through the PSTN. Check the CallManager traces for Q.931 disconnect messages. If a Q.931 disconnect message is present, it means the other party caused the disconnect, and we cannot correct that.

Gateway Registration Problems

One of the most common issues encountered with gateways on a Cisco CallManager is a registration problem. Registration can fail for a variety of reasons.

This section deals with two similar, but different, categories of gateways. The Analog Access AS-X, AT-X and Digital Access DT-24+ and DE-30+ belong to one category. These gateways are stand-alone units that are not directly connected to a Network Management Processor (NMP). The second category includes the Analog Access WS-X6624 and Digital Access WS-X6608. These gateways are blades installed in a Catalyst 6000 chassis with direct connectivity to the NMP for control and statusing.

In the examples below, the messages being explained are identified using bold text. This is to make it easier for you to see. In the actual display output, text is not bold. The examples are from an WS-X6624.

The first thing to check is that the gateway is up and running. All of the gateways have a heartbeat LED that blinks 1-second on, 1-second off when the gateway software is running normally. If this LED is not blinking at all, or blinking very rapidly, then the gateway software is not running. Normally, this will result in an automatic reset of the gateway. Also, it is normal for the gateway to reset itself if it cannot complete the registration process after about 2 to 3 minutes. So, you may happen to look at the heartbeat LED while the device is resetting. If the normal blinking pattern does not appear in 10 to 15 seconds, then the gateway has suffered a serious failure. On the AS-X or AT-X gateway, the heartbeat LED is the far right green LED showing on the front panel. On the DT-24+ or DE-30+ gateway, it is the far left red LED on the top edge of the card. On the Analog Access WS-X6624, it is a green LED inside the blade (not visible from the front panel) on the far right card edge near the front. Finally, on the Digital Access WS-X6608 there is a separate heartbeat LED for each of the 8 spans on the blade. There are 8 red LEDs across the card (not visible from the front panel) about 2/3 of the way towards the back.

The second thing to check is that the gateway has received its IP address. A standalone gateway **must** receive its IP address via DHCP or BOOTP. A Catalyst gateway may receive its IP address by DHCP, BOOTP, or by manual configuration through the NMP. If you have access to the DHCP server, the best way to check a standalone gateway is to verify that the device has an outstanding lease on an IP address. If the gateway shows up on your server, this is a good indication, but not definitive. Delete the lease at the DHCP server, and then reset the gateway. If the gateway re-appears on the server with a lease within a couple of minutes, then everything is working fine in this area. If not, then either the gateway cannot contact the DHCP server (Is a router improperly configured and not forwarding DHCP broadcasts? Is the server running?), or cannot get a positive response (Is the IP address pool depleted?). If checking these suggestions does not yield the answer, use a sniffer trace to determine the specific problem.

For a Catalyst 6000 gateway, you should make sure that the NMP can communicate with the gateway. You can check this by trying to ping its internal IP address from the NMP. The IP address is in the format:

127.1.*module.port*

So, in our example, we would do:

```
Console (enable) ping 127.1.7.1
127.1.7.1 is alive
```

If pinging works, then the ‘sh-port’ command will show IP address information. Make sure the IP address information and the TFTP IP address is correct as well. If the gateway is failing to obtain valid DHCP information, the tracy utility (which can be supplied by Cisco TAC) can be used to determine the problem. Issue the command from the Cat6000 CLI:

tracy_start *mod port*

In this example, the WS-X6624 is module 7 and it only has a single 860 processor, so it is port 1. The command issued is as follows:

tracy_start 7 1

The following output is actually from the 860-console port on the gateway board itself. However, the output of the tracy command is nothing more than a remote copy of the 860-console port.

```
      |               |
      |               |
    | | |           | | |
  | | | | |       | | | | |
| | | | | | |:.:| | | | | | |:..
C i s c o S y s t e m s
CAT6K Analog Gateway (ELVIS)
APP Version : A0020300, DSP Version : A0030300, Built Jun 1 2000 16:33:01
ELVIS>> 00:00:00.020 (XA) MAC Addr : 00-10-7B-00-13-DE
00:00:00.050 NMPTask:got message from XA Task
00:00:00.050 (NMP) Open TCP Connection ip:7f010101
00:00:00.050 NMPTask:Send Module Slot Info
00:00:00.060 NMPTask:get DIAGCMD
00:00:00.160 (DSP) Test Begin -> Mask<0x00FFFFFF>
00:00:01.260 (DSP) Test Complete -> Results<0x00FFFFFF/0x00FFFFFF>
00:00:01.260 NMPTask:get VLANCONFIG
00:00:02.870 (CFG) Starting DHCP
00:00:02.870 (CFG) Booting DHCP for dynamic configuration.
00:00:06.570 (CFG) DHCP Request or Discovery Sent, DHCPState = INIT_REBOOT
00:00:06.570 (CFG) DHCP Server Response Processed, DHCPState = INIT_REBOOT
00:00:06.780 (CFG) IP Configuration Change! Restarting now...
00:00:10.480 (CFG) DHCP Request or Discovery Sent, DHCPState = INIT
00:00:14:480 (CFG) DHCP Timeout Waiting on Server, DHCPState = INIT
00:00:22:480 (CFG) DHCP Timeout Waiting on Server, DHCPState = INIT
00:00:38:480 (CFG) DHCP Timeout Waiting on Server, DHCPState = INIT
```

If the above timeout message continues to scroll by, then there is a problem contacting the DHCP server. Check that the Catalyst 6000 gateway port is in the correct VLAN.

This information is in the **show port** command from before. If the DHCP server is not on the same VLAN as the Catalyst 6000 gateway, then make sure the appropriate IP Helper addresses have been configured to forward the DHCP requests to the DHCP server. It is possible for the gateway to get stuck in the INIT state after a VLAN number change until the gateway resets. When in this state, you can try resetting the gateway. Every time the 860 is reset, your tracy session will be lost. Therefore, you must close your existing session and re-establish a new one by issuing the following commands:

tracy_close *mod port*

tracy_start *mod port*

If all this checks out and you're still seeing the **DHCPState = INIT** messages, then check to see if the DHCP server is functioning correctly. If so, start a sniffer trace to see if the requests are being sent and if the server is responding.

Once DHCP is working correctly, the gateway will have an IP address that will allow the use of the tracy debugging utility. This utility is a built-in feature of the NMP command set for the Catalyst gateways and is available as a helper application that runs on Windows 98/NT/2000 for the standalone gateways. To use the helper application tracy utility, you need to connect to the gateway by using the IP address to which it is assigned. This tracy application works on all the gateways, provides a separate trace window for each gateway (up to eight may be traced at once), and allows traces to be logged directly to a file you specify.

The next step is to verify that the TFTP server IP address was correctly provided to the gateway. This is normally provided by DHCP in either Option 66 (by name or IP address), Option 150 (IP address only), or si_addr (IP address only). If your server has multiple Options configured, si_addr will take precedence over Option 150, which will take precedence over Option 66. If Option 66 provides the DNS_NAME of the TFTP server, then the DNS server(s) IP address(es) must have been specified by

DHCP, **and** the name entered in Option 66 must resolve to the correct TFTP server IP address. A Catalyst gateway could be configured by the NMP to disable DHCP, and the NMP operator must then enter all configuration parameters by hand at the console, including the TFTP server address.

Additionally, the gateways will always attempt to resolve the name *CiscoCMI* via DNS. If successful, the CiscoCM1 IP address will take precedence over anything the DHCP server or NMP tells it for the TFTP server address, even if the NMP has DHCP disabled.

You can check the current TFTP server IP address in a gateway by using the **tracy** utility. Enter the following command to get the configuration task number:

```
TaskID: 0
Cmd: show tl
```

Look for a line with **config** or **CFG** and use the corresponding number as the taskID for the next line. For example, for the Digital Access WS-X6624 gateway, the command to dump the DHCP information is:

```
TaskID: 6
Cmd: show dhcp
```

The TFTP server IP address is then clearly shown. If it is not correct, verify that your DHCP options and other information it provides are correct.

Once the TFTP address is correct, the next step is to ensure that the gateway is getting its configuration file from the TFTP server. If you see the following in the tracy output, your TFTP service may not be working correctly or the gateway might not be configured on the Cisco CallManager:

```
00:09:05.620 (CFG) Requesting SAA00107B0013DE.cnf File From TFTP Server
00:09:18.620 (CFG) TFTP
Error: Timeout Awaiting Server Response for .cnf File!
```

The gateway will attempt to connect to the same IP address as the TFTP server if it does not receive a configuration file. This is fine unless you are in a clustered environment in which the gateway needs to receive its list of redundant Cisco CallManagers. If the card is not receiving its TFTP information correctly, check the TFTP service on the Cisco CallManager and make sure it is running. Also, check the TFTP trace on the Cisco CallManager as well.

Another common problem is that the gateway is not configured correctly on the Cisco CallManager. A typical error is entering an incorrect MAC address for the gateway. If this is the case, for a Catalyst 6000 gateway, you will probably get the following messages on the NMP console every two minutes:

```
2000 Apr 14 19:24:08 %SYS-4-MODHPRESET:Host process (860) 7/1 got reset asynchronously
2000 Apr 14 19:26:05 %SYS-4-MODHPRESET:Host process (860) 7/1 got reset asynchronously
2000 Apr 14 19:28:02 %SYS-4-MODHPRESET:Host process (860) 7/1 got reset asynchronously
```

This is what the tracy output would look like if the gateway is not in the Cisco CallManager database:

```
00:00:01.670 (CFG) Booting DHCP for dynamic configuration.
00:00:05.370 (CFG) DHCP Request or Discovery Sent, DHCPState = INIT_REBOOT
00:00:05.370 (CFG) DHCP Server Response Processed, DHCPState = BOUND
00:00:05.370 (CFG) Requesting DNS Resolution of CiscoCM1
00:00:05.370 (CFG) DNS Error on Resolving TFTP Server Name.
00:00:05.370 (CFG) TFTP Server IP Set by DHCP Option 150 = 10.123.9.2
00:00:05.370 (CFG) Requesting SAA00107B0013DE.cnf File From TFTP Server
00:00:05.370 (CFG) TFTP Error: .cnf File Not Found!
00:00:05.370 (CFG) Requesting SAADefault.cnf File From TFTP Server
00:00:05.380 (CFG) .cnf File Received and Parsed Successfully.
00:00:05.380 (CFG) Updating Configuration ROM...
00:00:05.610 GMSG: GWEvent = CFG_DONE --> GWState = SrchActive
00:00:05.610 GMSG: CCM#0 CPEvent = CONNECT_REQ --> CPState = AttemptingSocket
00:00:05.610 GMSG: Attempting TCP socket with CCM 10.123.9.2
00:00:05.610 GMSG: CCM#0 CPEvent = SOCKET_ACK --> CPState = BackupCCM
```

```
00:00:05.610 GMSG: GWEvent = SOCKET_ACK --> GWState = RegActive
00:00:05.610 GMSG: CCM#0 CPEvent = REGISTER_REQ --> CPState = SentRegister
00:00:05.680 GMSG: CCM#0 CPEvent = CLOSED --> CPState = NoTCPSocket
00:00:05.680 GMSG: GWEvent = DISCONNECT --> GWState = Rollover
00:00:20.600 GMSG: GWEvent = TIMEOUT --> GWState = SrchActive
00:00:20.600 GMSG: CCM#0 CPEvent = CONNECT_REQ --> CPState = AttemptingSocket
00:00:20.600 GMSG: Attempting TCP socket with CCM 10.123.9.2
00:00:20.600 GMSG: CCM#0 CPEvent = SOCKET_ACK --> CPState = BackupCCM
```

Another possible registration problem could be if the load information is incorrect or the load file is corrupt. The problem could also occur if the TFTP server is not working. In this case, tracy clearly shows that the TFTP server reported that the file is not found:

```
00:00:07.390 GMSG: CCM#0 CPEvent = REGISTER_REQ --> CPState = SentRegister
00:00:08.010 GMSG: TFTP Request for application load A0021300
00:00:08.010 GMSG: CCM#0 CPEvent = LOADID --> CPState = AppLoadRequest
00:00:08.010 GMSG: ***TFTP Error: File Not Found***
00:00:08.010 GMSG: CCM#0 CPEvent = LOAD_UPDATE --> CPState = LoadResponse
```

In this case, you can see that the gateway is requesting application load A0021300, although the correct load name would be A0020300. For a Catalyst 6000 gateway, the same problem can occur when a new application load needs to get its corresponding DSP load as well. If the new DSP load is not found, a similar message will appear.

The following shows the output when an Analog Access WS-X6224 has been configured to retrieve an incorrect application load. The output looks similar to that of a gateway that has not been configured on the Cisco CallManager:

```
      |               |
      |               |
    | | |           | | |
  | | | | |       | | | | |
| | | | | | |:.:| | | | | | |:..
C i s c o S y s t e m s
CAT6K Analog Gateway (ELVIS)
APP Version : A0020300, DSP Version : A0030300, Built Jun 1 2000 16:33:01
ELVIS>> 00:00:00.020 (XA) MAC Addr : 00-10-7B-00-13-DE
00:00:00.050 NMPTask:got message from XA Task
00:00:00.050 (NMP) Open TCP Connection ip:7f010101
00:00:00.050 NMPTask:Send Module Slot Info
00:00:00.060 NMPTask:get DIAGCMD
00:00:00.160 (DSP) Test Begin -> Mask<0x00FFFFFF>
00:00:01.260 (DSP) Test Complete -> Results<0x00FFFFFF/0x00FFFFFF>
00:00:01.260 NMPTask:get VLANCONFIG
00:00:02.030 (CFG) Starting DHCP
00:00:02.030 (CFG) Booting DHCP for dynamic configuration.
00:00:05.730 (CFG) DHCP Request or Discovery Sent, DHCPState = INIT_REBOOT
00:00:05.730 (CFG) DHCP Server Response Processed, DHCPState = BOUND
00:00:05.730 (CFG) Requesting DNS Resolution of CiscoCM1
00:00:05.730 (CFG) DNS Error on Resolving TFTP Server Name.
00:00:05.730 (CFG) TFTP Server IP Set by DHCP Option 150 = 10.123.9.2
00:00:05.730 (CFG) Requesting SAA00107B0013DE.cnf File From TFTP Server
00:00:05.730 (CFG) .cnf File Received and Parsed Successfully.
00:00:05.730 GMSG: GWEvent = CFG_DONE --> GWState = SrchActive
00:00:05.730 GMSG: CCM#0 CPEvent = CONNECT_REQ --> CPState = AttemptingSocket
00:00:05.730 GMSG: Attempting TCP socket with CCM 10.123.9.2
00:00:05.730 GMSG: CCM#0 CPEvent = SOCKET_ACK --> CPState = BackupCCM
00:00:05.730 GMSG: GWEvent = SOCKET_ACK --> GWState = RegActive
00:00:05.730 GMSG: CCM#0 CPEvent = REGISTER_REQ --> CPState = SentRegister
00:00:06.320 GMSG: CCM#0 CPEvent = LOADID --> CPState = LoadResponse
00:01:36.300 GMSG: CCM#0 CPEvent = TIMEOUT --> CPState = BadCCM
00:01:36.300 GMSG: GWEvent = DISCONNECT --> GWState = Rollover
00:01:46.870 GMSG: CCM#0 CPEvent = CLOSED --> CPState = NoTCPSocket
```

```
00:01:51.300 GMSG: GWEvent = TIMEOUT --> GWState = SrchActive
00:01:51.300 GMSG: CCM#0 CPEvent = CONNECT_REQ --> CPState = AttemptingSocket
00:01:51.300 GMSG: Attempting TCP socket with CCM 10.123.9.2
00:01:51.300 GMSG: CCM#0 CPEvent = SOCKET_ACK --> CPState = BackupCCM
00:01:51.300 GMSG: GWEvent = SOCKET_ACK --> GWState = RegActive
00:01:51.300 GMSG: CCM#0 CPEvent = REGISTER_REQ --> CPState = SentRegister
00:01:51.890 GMSG: CCM#0 CPEvent = LOADID --> CPState = LoadResponse
```

The difference here is that the gateway gets stuck in the **LoadResponse** stage and eventually times out. This problem can be resolved by correcting the load file name in the Device Defaults area of Cisco CallManager Administration.

Gatekeeper Registration Problems

Before starting any gateway-to-gatekeeper troubleshooting, verify that there is IP connectivity within the network. Assuming that there is IP connectivity, use the information in this section to troubleshoot your gateway.

Inter-Cluster Trunks Only

Note that gatekeeper control for Cisco CallManager Release 3.0(1) and later is only available for inter-cluster trunks. Gatekeeper control is configurable for other devices, but the configuration is not supported.

Admission Rejects

Admission Rejects (ARJs) are issued when Cisco CallManager has registered with the Gatekeeper, but can’t send a phone call. Configuration issues on the gatekeeper should be the primary focus when the gatekeeper is issuing a ARJ. However, below are the general guidelines for troubleshooting:

1. Verify IP connectivity from the gateway to the gatekeeper.
2. Show gatekeeper status – verify the gatekeeper state is up.
3. Is there a zone subnet defined on the gatekeeper? If so, verify that the subnet of the gateway is in the allowed subnets.

Registration Rejects

Registration Rejects (RRJ) are issued when Cisco CallManager cannot register with the Gatekeeper. Configuration issues on the gatekeeper should be the primary focus when the gatekeeper is issuing an RRJ.

However, here are the general guidelines for troubleshooting:

1. Verify IP connectivity from the gateway to the gatekeeper.
2. Show gatekeeper status – verify the gatekeeper state is up.
3. Is there a zone subnet defined on the gatekeeper? If so, verify that the subnet of the gateway is in the allowed subnets.

Cisco IP Phone Initialization Process

The Cisco IP phone initialization (or *boot up*) process is explained in detail below.

1. At initialization, the Cisco IP phone sends a request to the DHCP server to get an IP address, DNS server address, and TFTP server name or address, if appropriate. Options are set in DHCP server (Option 066, Option 150, and so on). It also gets a default gateway address if set in DHCP server (Option 003).
2. If a DNS name of the TFTP sever is sent by DHCP, then a DNS sever IP address is required to map the name to an IP address. This step is bypassed if the DHCP server sends the IP address of the TFTP server. In this case study, the DHCP server sent the IP address of TFTP because DNS was not configured.
3. If a TFTP server name is not included in the DHCP reply, then the Cisco IP phone uses the default server name.
4. The configuration file (.cnf) file is retrieved from the TFTP server. All .cnf files have the name SEP<mac_address>.cnf, where “SEP” is an acronym for Selsius Ethernet Phone. If this is the first time the phone is registering with the Cisco CallManager, then a default file, SEPdefault.cnf, is downloaded to the Cisco IP phone. In this case study, the first Cisco IP phone uses the IP address 172.16.70.230 (its MAC address is SEP0010EB001720), and the second Cisco IP phone uses the IP address 172.16.70.231 (its MAC address is SEP003094C26105).
5. All .cnf files include the IP address(es) of the CallManager(s) defined in the CallManager group for that device.
6. Once the Cisco IP phone has connected and registered with Cisco CallManager, the Cisco CallManager tells the Cisco IP phone which executable version (called a load ID) to run. If the specified version does not match the executing version on the Cisco IP phone, the Cisco IP phone will request the new executable from the TFTP server and reset automatically.

Skinny Station Registration Process

Cisco IP phones communicate with Cisco CallManager using the Cisco Skinny Station Protocol. The registration process allows a Skinny Station, such as a Cisco IP phone, to inform Cisco CallManager of its existence and to make calling possible. Figure 6-52 shows the different messages that are exchanged between the Cisco IP phone (the *station*) and the Cisco CallManager. The primary messages in the Skinny Station registration process are described in Table 6-34.

Figure 6-52 Example Cisco Skinny Station Protocol Message Exchange

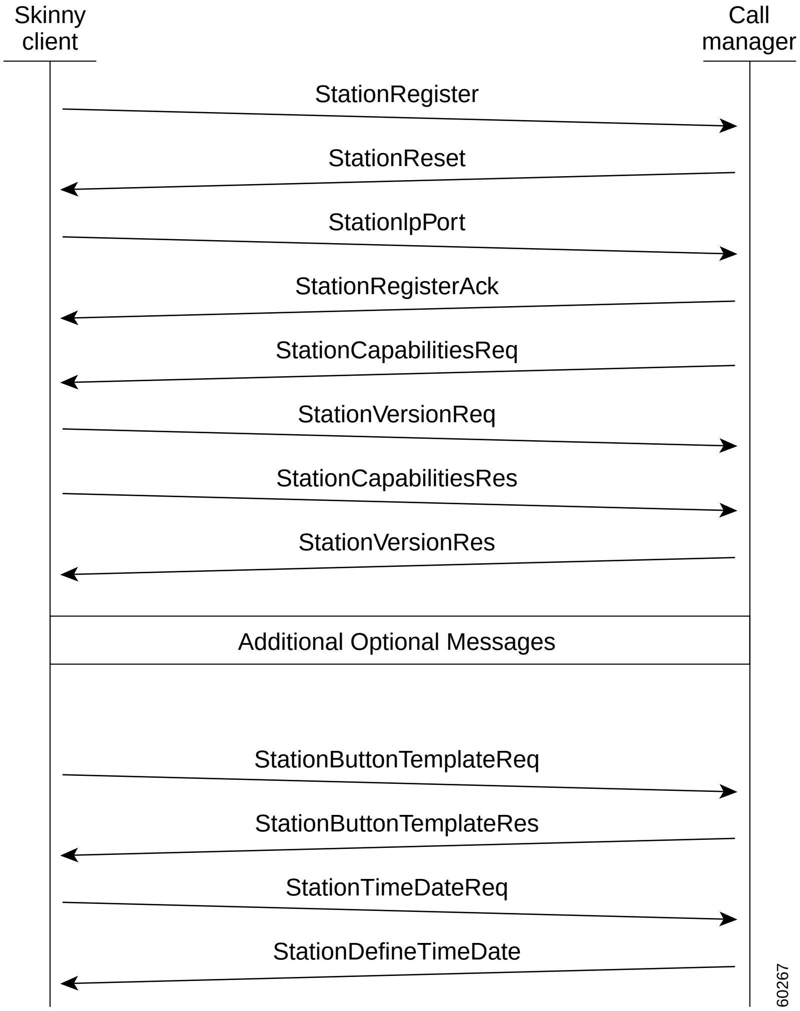

Table 6-34 Skinny Station Registration Process Descriptions

Message	Description
Station Register	The station sends this message to announce its existence to the controlling Cisco CallManager.
Station Reset	Cisco CallManager sends this message to command the station to reset its processes.
Station IP Port	The station sends this message to provide Cisco CallManager with the User Datagram Protocol (UDP) port to be used with the RTP stream.
Station Register Acknowledge	Cisco CallManager sends this message to acknowledge the registration of a station.

Table 6-34 Skinny Station Registration Process Descriptions (continued)

Message	Description

Station Register Reject	Cisco CallManager sends this message to reject a registration attempt from the indicated phone.
	`char text[StationMaxDisplayTextSize];`
	`};`
	Where:
	`text` is a character string, maximum length of 33 bytes, containing a textual description of the reason that registration is rejected.

Station Capabilities Request	Cisco CallManager sends this message to request the current capabilities of the station. Station capabilities may include compression standard and other H.323 capabilities.
Station Version Request	The station sends this message to request the version number of the software load for the station.
Station Version Response	Cisco CallManager sends this message to inform the station of the appropriate software version number.
Station Capabilities Response	The station sends this message to Cisco CallManager in response to a Station Capabilities Request. The station’s capabilities are cached in the Cisco CallManager and used to negotiate terminal capabilities with an H.323 compliant Terminal.
Station Button Template Request	The station sends this message to request the button template definition for that specific terminal or Cisco IP phone.
Station Button Template Response	Cisco CallManager sends this message to update the button template information contained in the station.
Station Time Date Request	The station sends this message to request the current date and time for internal usage and for displaying as a text string.
Station Define Time and Date	Cisco CallManager uses this message to provide the date and time information to the station. It provides time synchronization for the stations.

Cisco CallManager Initialization Process

In this section the initialization process of Cisco CallManager will be explained with the help of traces that are captured from CCM1 (identified by the IP address 172.16.70.228). As described previously, CallManager traces are a very effective troubleshooting tool because they detail every packet sent between endpoints. This section will describe the events that occur when Cisco CallManager is initialized. Understanding how to read the trace helps you to properly troubleshoot the various Cisco CallManager processes, and the effect of those processes on services such as conferencing, call forwarding, and so on.

- The following messages from the Cisco CallManager CCM trace utility show the initialization process on one of the Cisco CallManagers, in this case, CCM1. Review the descriptions of each message below.
- The first message indicates that Cisco CallManager started its initialization process.
- The second message indicates that Cisco CallManager read the default database values, which would be the primary or publisher database (for this case).

- The third message indicates that CallManager established TCP port 8002 for other CallManagers in the cluster to communicate with the server.
- The fourth message shows that, after listening to these messages, Cisco CallManager added a second Cisco CallManager to its list: CCM2 (172.16.70.229).
- The fifth message indicates that Cisco CallManager has started and is running Cisco CallManager version 3.0.20.

```
16:02:47.765 CCM|CMProcMon - CallManagerState Changed - Initialization Started.
16:02:47.796 CCM|NodeId: 0, EventId: 107 EventClass: 3 EventInfo: Cisco CM Database
Defaults Read
16:02:49.937 CCM| SDL Info - NodeId: [1], Listen IP/Hostname: [172.16.70.228], Listen
Port: [8002]
16:02:49.984 CCM|dBProcs - Adding SdlLink to NodeId: [2], IP/Hostname: [172.16.70.229]
16:02:51.031 CCM|NodeId: 1, EventId: 1 EventClass: 3 EventInfo: Cisco CallManager
Version=<3.0(0.20)> started
```

Self-Starting Processes

Once Cisco CallManager is up and running, it starts several other processes within itself. Some of these processes are shown below, including MulticastPoint Manager, UnicastBridge Manager, digit analysis, and route list. The messages described during these processes can be very useful when troubleshooting a problem related to the features in Cisco CallManager.

For example, assume that the route lists are not functioning and are unusable. To troubleshoot this problem, you would monitor these traces to determine whether the Cisco CallManager has started RoutePlanManager and if it is trying to load the route lists. In the sample configuration below, RouteListName="ipwan" and RouteGroupName="ipwan" are loading and starting.

```
16:02:51.031 CCM|MulicastPointManager - Started
16:02:51.031 CCM|UnicastBridgeManager - Started
16:02:51.031 CCM|MediaTerminationPointManager - Started
16:02:51.125 CCM|MediaCoordinator(1) - started
16:02:51.125 CCM|NodeId: 1, EventId: 1543 EventClass: 2 EventInfo: Database manager
started
16:02:51.234 CCM|NodeId: 1, EventId: 1542 EventClass: 2 EventInfo: Link manager started
16:02:51.390 CCM|NodeId: 1, EventId: 1541 EventClass: 2 EventInfo: Digit analysis started
16:02:51.406 CCM|RoutePlanManager - Started, loading RouteLists
16:02:51.562 CCM|RoutePlanManager - finished loading RouteLists
16:02:51.671 CCM|RoutePlanManager - finished loading RouteGroups
16:02:51.671 CCM|RoutePlanManager - Displaying Resulting RoutePlan
16:02:51.671 CCM|RoutePlanServer - RouteList Info, by RouteList and RouteGroup Selection
Order
16:02:51.671 CCM|RouteList - RouteListName=''ipwan''
16:02:51.671 CCM|RouteList - RouteGroupName=''ipwan''
16:02:51.671 CCM|RoutePlanServer - RouteGroup Info, by RouteGroup and Device Selection
Order
16:02:51.671 CCM|RouteGroup - RouteGroupName=''ipwan''
```

The following trace shows the RouteGroup adding the device 172.16.70.245, which is a CallManager located in a different cluster, and is considered an H.323 device. In this case, the RouteGroup is created to route calls to 172.16.70.245 with Cisco IOS gatekeeper permission. If there is a problem routing the call to a Cisco IP phone located in that cluster, then the following messages would help you find the cause of the problem.

```
16:02:51.671 CCM|RouteGroup - DeviceName=''172.16.70.245''
16:02:51.671 CCM|RouteGroup -AllPorts
```

Part of the initialization process shows that Cisco CallManager is adding DNs. By reviewing these messages, you can determine whether the Cisco CallManager has read the route pattern from the database.

```
16:02:51.671 CCM|NodeId: 1, EventId: 1540 EventClass: 2 EventInfo: Call control started
16:02:51.843 CCM|ProcessDb - Dn = 2XXX, Line = 0, Display = , RouteThisPattern,
NetworkLocation = OffNet, DigitDiscardingInstruction = 1, WhereClause =
16:02:51.859 CCM|Digit analysis: Add local pattern 2XXX , PID: 1,80,1
16:02:51.859 CCM|ForwardManager - Started
16:02:51.984 CCM|CallParkManager - Started
16:02:52.046 CCM|ConferenceManager - Started
```

In the following traces the Device Manager in Cisco CallManager is statically initializing two devices. The device with IP address 172.17.70.226 is a gatekeeper and the device with IP address 172.17.70.245 is another Cisco CallManager in a different cluster. That Cisco CallManager is registered as an H.323 Gateway with this Cisco CallManager.

```
16:02:52.250 CCM|DeviceManager: Statically Initializing Device; DeviceName=172.16.70.226
16:02:52.250 CCM|DeviceManager: Statically Initializing Device; DeviceName=172.16.70.245
```

Cisco CallManager Registration Process

Another important part of the CallManager trace is the registration process. When a device is powered up, it gets information via DHCP, connects to the TFTP server for its .cnf file, and then connects to the Cisco CallManager specified in the .cnf file. The device could be an MGCP Gateway, a Skinny Gateway, or a Cisco IP phone. Therefore, it is important to be able to discover whether or not devices have successfully registered on the Cisco IP Telephony network.

In the following trace, Cisco CallManager has received new connections for registration. The registering devices are *MTP_nsa-cm1* (MTP services on CCM1), and *CFB_nsa-cm1* (Conference Bridge service on CCM1). These are software services running on Cisco CallManager but are treated internally as different external services and are therefore assigned a TCPHandle, socket number, and port number as well as a device name.

```
16:02:52.750 CCM|StationInit - New connection accepted. DeviceName=, TCPHandle=0x4fbaa00,
Socket=0x594, IPAddr=172.16.70.228, Port=3279, StationD=[0,0,0]
16:02:52.750 CCM|StationInit - New connection accepted. DeviceName=, TCPHandle=0x4fe05e8,
Socket=0x59c, IPAddr=172.16.70.228, Port=3280, StationD=[0,0,0]
16:02:52.781 CCM|StationInit - Processing StationReg. regCount: 1 DeviceName=MTP_nsa-cm1,
TCPHandle=0x4fbaa00, Socket=0x594, IPAddr=172.16.70.228, Port=3279, StationD=[1,45,2]
16:02:52.781 CCM|StationInit - Processing StationReg. regCount: 1 DeviceName=CFB_nsa-cm1,
TCPHandle=0x4fe05e8, Socket=0x59c, IPAddr=172.16.70.228, Port=3280, StationD=[1,96,2]
```

In the following trace, Skinny Station messages are sent between a Cisco IP phone and Cisco CallManager. The Cisco IP phone (172.16.70.231) is registering with Cisco CallManager. Refer to the descriptions of Skinny Station messages earlier in this section for more information. As soon as Cisco CallManager receives the registration request from a Cisco IP phone, it assigns a TCPHandle number to this device. This number remains the same until the device or Cisco CallManager is restarted. Therefore, you can follow all the events related to a particular device by searching for or keeping track of the device’s TCPHandle number, which appears in hex. Also, notice that Cisco CallManager provides the load ID to the Cisco IP phone. Based on this load ID, the Cisco IP phone runs the executable file (acquired from the TFTP server) that corresponds to the device.

```
16:02:57.000 CCM|StationInit - New connection accepted. DeviceName=, TCPHandle=0x4fbbc30,
Socket=0x5a4, IPAddr=172.16.70.231, Port=52095, StationD=[0,0,0]
16:02:57.046 CCM|NodeId: 1, EventId: 1703 EventClass: 2 EventInfo: Station Alarm, TCP
Handle: 4fbbc30, Text: Name=SEP003094C26105 Load=AJ.30 Parms=Status/IPaddr LastTime=A P1:
2304(900) P2: -414838612(e74610ac)
```

```
16:02:57.046 CCM|StationInit - ***** InboundStim - AlarmMessageID tcpHandle=0x4fbbc30
Message="Name=SEP003094C26105 Load=AJ.30 Parms=Status/IPaddr LastTime=A" Parm1=2304 (900)
Parm2=-414838612 (e74610ac)
16:02:57.093 CCM|StationInit - Processing StationReg. regCount: 1
DeviceName=SEP003094C26105, TCPHandle=0x4fbbc30, Socket=0x5a4, IPAddr=172.16.70.231,
Port=52095, StationD=[1,85,1]
16:02:57.093 CCM|StationInit - InboundStim - IpPortMessageID: 32715(0x7fcb)
tcpHandle=0x4fbbc30
```

Cisco CallManager KeepAlive Process

The station, device, or service and the Cisco CallManager use the following messages to maintain knowledge of the communications channel between them. The messages are used to begin the KeepAlive sequence that ensures that the communications link between the Cisco CallManager and the station remains active. The following messages can originate from either the Cisco CallManager or the station.

```
16:03:02.328 CCM|StationInit - InboundStim - KeepAliveMessage - Forward KeepAlive to
StationD. DeviceName=MTP_nsa-cm2, TCPHandle=0x4fa7dc0, Socket=0x568, IPAddr=172.16.70.229,
Port=1556, StationD=[1,45,1]
16:03:02.328 CCM|StationInit - InboundStim - KeepAliveMessage - Forward KeepAlive to
StationD. DeviceName=CFB_nsa-cm2, TCPHandle=0x4bf8a70, Socket=0x57c, IPAddr=172.16.70.229,
Port=1557, StationD=[1,96,1]
16:03:06.640 CCM|StationInit - InboundStim - KeepAliveMessage - Forward KeepAlive to
StationD. DeviceName=SEP0010EB001720, TCPHandle=0x4fbb150, Socket=0x600,
IPAddr=172.16.70.230, Port=49211, StationD=[1,85,2]
16:03:06.703 CCM|StationInit - InboundStim - KeepAliveMessage - Forward KeepAlive to
StationD. DeviceName=SEP003094C26105, TCPHandle=0x4fbbc30, Socket=0x5a4,
IPAddr=172.16.70.231, Port=52095, StationD=[1,85,1]
```

The messages in the following trace depict the KeepAlive sequence, which indicates that the communications link between the Cisco CallManager and the station is active. Again, these messages can originate either by the Cisco CallManager or the station.

```
16:03:02.328 CCM|MediaTerminationPointControl - stationOutputKeepAliveAck
tcpHandle=4fa7dc0
16:03:02.328 CCM|UnicastBridgeControl - stationOutputKeepAliveAck tcpHandle=4bf8a70
16:03:06.703 CCM|StationInit - InboundStim - IpPortMessageID: 32715(0x7fcb)
tcpHandle=0x4fbbc30
16:03:06.703 CCM|StationD - stationOutputKeepAliveAck tcpHandle=0x4fbbc30
```

This section describes a Cisco IP phone (directory number 1000) calling another Cisco IP phone (directory number 1001) within the same cluster. A cluster is a group of Cisco CallManagers that have one server that is the *publisher* SQL database and one or more servers with *subscriber* SQL databases. In a centralized call processing model, all phones must be registered to one CallManager.

In our sample topology, CCM1 is the publisher and CCM2 is a subscriber. The two Cisco IP phones (1000 and 1001) are registered to CCM2. The call flow is shown in the Figure 6-53. When a Cisco IP phone goes off-hook, it initiates a Skinny Station message to the CallManager server over the TCP handle established during initialization. After call control signaling is established between the two Cisco IP phones and the Cisco CallManager, the RTP stream starts flowing directly between the two phones, as shown in Figure 6-53. The Skinny Station call flow messages for this intra-cluster call are explained in next section.

Figure 6-53 Example Call Flow Within the Same Cluster

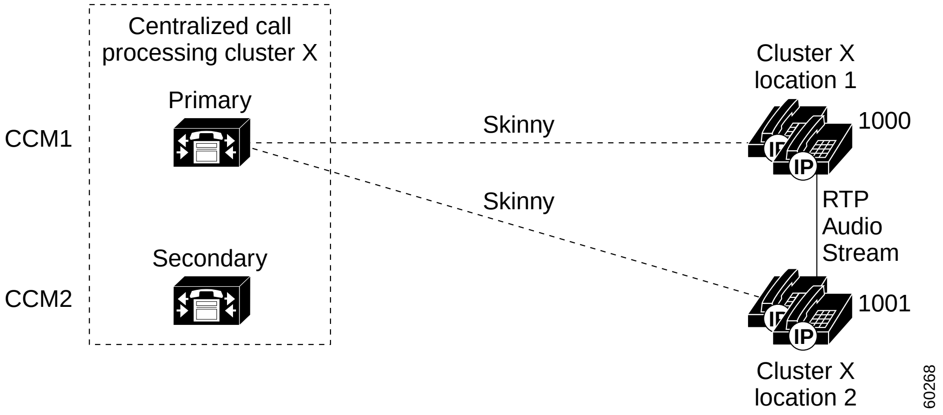

Cisco IP Phone-to-Cisco IP Phone Exchange of Skinny Station Messages During Call Flow

Figure 6-54 illustrates an example exchange of messages between two Skinny Stations. The Skinny Station, or Cisco IP phone, initiates a connection to the Cisco CallManager, and then Cisco CallManager performs digit analysis before opening a control session with the destination Skinny Station. As the following diagram indicates, the Skinny Station messages are written using simple English so they can be readily understood by end-users. Because of this, these messages are not explained in this section. However, these call flow Skinny Station messages are explained in more detail in later sections when traces are being examined.

Figure 6-54 Example Message Exchange Between Skinny Stations

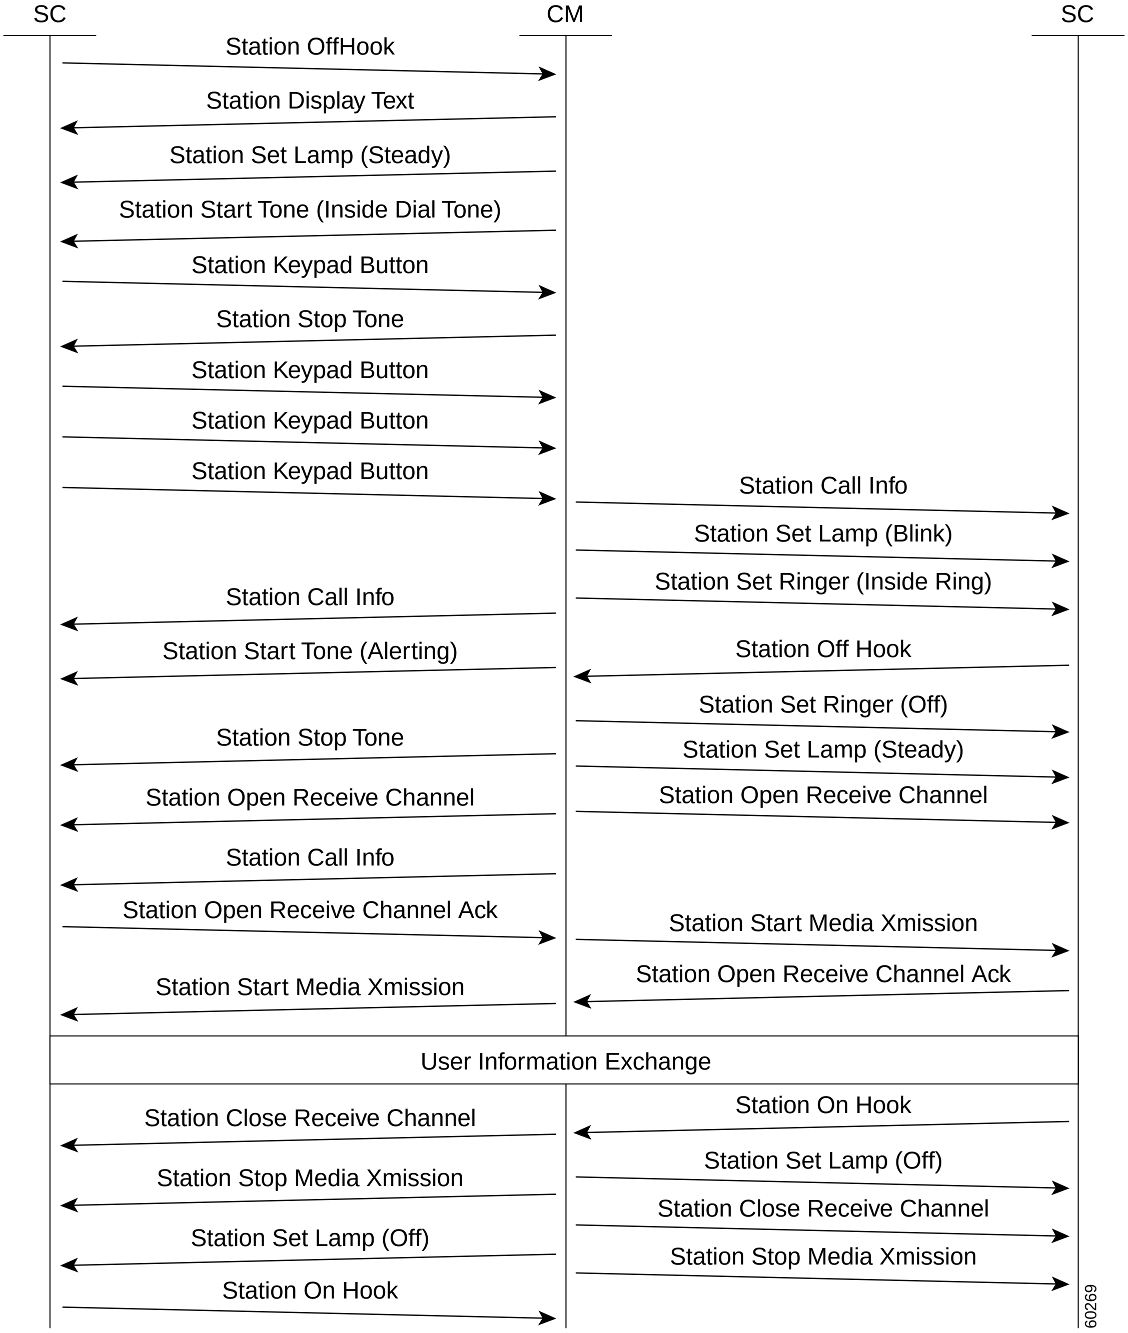
SC
CM
SC
Station OffHook
Station Display Text
Station Set Lamp (Steady)
Station Start Tone (Inside Dial Tone)
Station Keypad Button
Station Stop Tone
Station Keypad Button
Station Keypad Button
Station Keypad Button
Station Call Info
Station Set Lamp (Blink)
Station Set Ringer (Inside Ring)
Station Call Info
Station Start Tone (Alerting)
Station Off Hook
Station Set Ringer (Off)
Station Stop Tone
Station Set Lamp (Steady)
Station Open Receive Channel
Station Open Receive Channel
Station Call Info
Station Open Receive Channel Ack
Station Start Media Xmission
Station Start Media Xmission
Station Open Receive Channel Ack
User Information Exchange
Station On Hook
Station Close Receive Channel
Station Set Lamp (Off)
Station Stop Media Xmission
Station Close Receive Channel
Station Set Lamp (Off)
Station Stop Media Xmission
Station On Hook
60269

Cisco CallManager Intra-Cluster Call Flow Traces

The following CallManager traces explore in detail the intra-cluster call flow. The Cisco IP phones in the call flow can be identified by the DN, tcpHandle, and IP address. A Cisco IP phone (DN=1001, tcpHandle=0x4fbbc30, IP address=172.16.70.231) is calling another Cisco IP phone in the same Cluster (DN=1000, tcpHandle= 0x4fbb150, IP address= 172.16.70.230). Remember that you can follow a device through the trace by looking at the TCP handle value, time stamp, or name of the device. The TCP handle value for the device remains the same until the device is rebooted or goes offline.

The following traces show that the Cisco IP phone (1001) has gone off-hook. The trace below shows the unique messages, TCP handle, and the DN, which is displayed on the Cisco IP phone. There is no called number at this point, because the user has not tried to dial any digits. The information below is in the form of Skinny Station messages between the Cisco IP phones and the Cisco CallManager.

```
16:05:41.625 CCM|StationInit - InboundStim - OffHookMessageID tcpHandle=0x4fbbc30
16:05:41.625 CCM|StationD - stationOutputDisplayText tcpHandle=0x4fbbc30, Display= 1001\
```

The next trace shows Skinny Station messages going from Cisco CallManager to a Cisco IP phone. The first message is to turn on the lamp on the calling party's Cisco IP phone.

```
16:05:41.625 CCM|StationD - stationOutputSetLamp stim: 9=Line instance=1 lampMode=LampOn
tcpHandle=0x4fbbc30
```

The stationOutputCallState message is used by Cisco CallManager to notify the station of certain call-related information.

```
16:05:41.625 CCM|StationD - stationOutputCallState tcpHandle=0x4fbbc30
```

The stationOutputDisplayPromptStatus message is used by Cisco CallManager to cause a call-related prompt message to be displayed on the Cisco IP phone.

```
16:05:41.625 CCM|StationD - stationOutputDisplayPromptStatus tcpHandle=0x4fbbc30
```

The stationOutputSelectSoftKey message is used by Cisco CallManager to cause the Skinny Station to select a specific set of soft keys to be displayed on the phone.

```
16:05:41.625 CCM|StationD - stationOutputSelectSoftKeys tcpHandle=0x4fbbc30
```

The next message is used by Cisco CallManager to instruct the Skinny Station as to the correct line context for the display.

```
16:05:41.625 CCM|StationD - stationOutputActivateCallPlane tcpHandle=0x4fbbc30
```

In the following message, the digit analysis process is ready to identify incoming digits and check them for potential routing matches in the database. The entry, cn=1001, represents the calling party number. `dd=""` represents the dialed digit, which would show the called part number. Note that StationInit messages are sent by the phone, StationD messages are sent by Cisco CallManager, and digit analysis is performed by Cisco CallManager.

```
16:05:41.625 CCM|Digit analysis: match(fqcn="", cn="1001", pss="", dd="")
16:05:41.625 CCM|Digit analysis: potentialMatches=PotentialMatchesExist
```

The following debug message shows that the Cisco CallManager is providing inside dial tone to the calling party Cisco IP phone.

```
16:05:41.625 CCM|StationD - stationOutputStartTone: 33=InsideDialTone tcpHandle=0x4fbbc30
```

Once Cisco CallManager detects an incoming message and recognizes the keypad button 1 has been pressed on the Cisco IP phone, it immediately stops the output tone.

```
16:05:42.890 CCM|StationInit - InboundStim - KeypadButtonMessageID kpButton: 1
tcpHandle=0x4fbbc30
16:05:45.140 CCM|StationInit - Offhook
16:05:42.890 CCM|StationD - stationOutputStopTone tcpHandle=0x4fbbc30
16:05:42.890 CCM|StationD - stationOutputSelectSoftKeys tcpHandle=0x4fbbc30
16:05:42.890 CCM|Digit analysis: match(fqcn="", cn="1001", pss="", dd="1")
16:05:42.890 CCM|Digit analysis: potentialMatches=PotentialMatchesExist
16:05:43.203 CCM|StationInit - InboundStim - KeypadButtonMessageID kpButton: 0
tcpHandle=0x4fbbc30
16:05:43.203 CCM|Digit analysis: match(fqcn="", cn="1001", pss="", dd="10")
16:05:43.203 CCM|Digit analysis: potentialMatches=PotentialMatchesExist
16:05:43.406 CCM|StationInit - InboundStim - KeypadButtonMessageID kpButton: 0
tcpHandle=0x4fbbc30
```

```
16:05:43.406 CCM|Digit analysis: match(fqcn="", cn="1001", pss="", dd="100")
16:05:43.406 CCM|Digit analysis: potentialMatches=PotentialMatchesExist|
16:05:43.562 CCM|StationInit - InboundStim - KeypadButtonMessageID kpButton: 0
tcpHandle=0x4fbbc30
16:05:43.562 CCM|Digit analysis: match(fqcn="", cn="1001", pss="", dd="1000")
```

Once the Cisco CallManager has received enough digits to match, it provides the digit analysis results in a table format. Cisco CallManager will ignore any extra digits pressed on the phone after this point, since a match has already been found.

```
16:05:43.562 CCM|Digit analysis: analysis results
16:05:43.562 CCM||PretransformCallingPartyNumber=1001
|CallingPartyNumber=1001
|DialingPattern=1000
|DialingRoutePatternRegularExpression=(1000)
|PotentialMatches=PotentialMatchesExist
|DialingSdlProcessId=(1,38,2)
|PretransformDigitString=1000
|PretransformPositionalMatchList=1000
|CollectedDigits=1000
|PositionalMatchList=1000
|RouteBlockFlag=RouteThisPattern
```

The next trace shows that Cisco CallManager is sending out this information to a called party phone (the phone is identified by the tcpHandle number).

```
16:05:43.578 CCM|StationD - stationOutputCallInfo CallingPartyName=1001,
CallingParty=1001, CalledPartyName=1000, CalledParty=1000, tcpHandle=0x4fbb150
```

The next trace indicates that Cisco CallManager is ordering the lamp to blink for incoming call indication on the called party's Cisco IP phone.

```
16:05:43.578 CCM|StationD - stationOutputSetLamp stim: 9=Line instance=1
lampMode=LampBlink tcpHandle=0x4fbb150
```

In the following traces, Cisco CallManager is providing ringer, display notification, and other call-related information to the called party's Cisco IP phone. Again, you can see that all messages are directed to the same Cisco IP phone because the same tcpHandle is used throughout the traces.

```
16:05:43.578 CCM|StationD - stationOutputSetRinger: 2=InsideRing tcpHandle=0x4fbb150
16:05:43.578 CCM|StationD - stationOutputDisplayNotify tcpHandle=0x4fbb150
16:05:43.578 CCM|StationD - stationOutputDisplayPromptStatus tcpHandle=0x4fbb150
16:05:43.578 CCM|StationD - stationOutputSelectSoftKeys tcpHandle=0x4fbb150
```

Notice that Cisco CallManager is also providing similar information to the calling party's Cisco IP phone. Again, the tcpHandle is used to differentiate between Cisco IP phones.

```
16:05:43.578 CCM|StationD - stationOutputCallInfo CallingPartyName=1001,
CallingParty=1001, CalledPartyName=, CalledParty=1000, tcpHandle=0x4fbbc30
16:05:43.578 CCM|StationD - stationOutputCallInfo CallingPartyName=1001,
CallingParty=1001, CalledPartyName=1000, CalledParty=1000, tcpHandle=0x4fbbc30
```

In the next trace, Cisco CallManager provides an alerting or ringing tone to the calling party's Cisco IP phone, notifying that the connection has been established.

```
16:05:43.578 CCM|StationD - stationOutputStartTone: 36=AlertingTone tcpHandle=0x4fbbc30
16:05:43.578 CCM|StationD - stationOutputCallState tcpHandle=0x4fbbc30
16:05:43.578 CCM|StationD - stationOutputSelectSoftKeys tcpHandle=0x4fbbc30
16:05:43.578 CCM|StationD - stationOutputDisplayPromptStatus tcpHandle=0x4fbbc30
```

At this point, the called party's Cisco IP phone goes off-hook. Therefore, Cisco CallManager stops generating the ringer tone to the calling party.

```
16:05:45.140 CCM|StationD - stationOutputStopTone tcpHandle=0x4fbbc30
```

In the following messages, Cisco CallManager requrest the Skinny Station to open an RTP port to receive the unicast audio stream from the other device. To do so, Cisco CallManager provides the IP address of the called party as well as codec information, and packet size in msec (milliseconds). PacketSize is an integer containing the sampling time in milliseconds used to create the RTP packets.

```
16:05:45.140 CCM|StationD - stationOutputOpenReceiveChannel tcpHandle=0x4fbbc30 myIP:
e74610ac (172.16.70.231)
16:05:45.140 CCM|StationD - ConferenceID: 0 msecPacketSize: 20
compressionType:(4)Media_Payload_G711Ulaw64k
```

Similarly, Cisco CallManager provides information to the called party (1000).

```
16:05:45.140 CCM|StationD - stationOutputOpenReceiveChannel tcpHandle=0x4fbb150 myIP:
e64610ac (172.16.70.230)
16:05:45.140 CCM|StationD - ConferenceID: 0 msecPacketSize: 20
compressionType:(4)Media_Payload_G711Ulaw64k
```

Cisco CallManager has received the acknowledgment message from called party for establishing the open channel for RTP stream, as well as the IP address of the called party. This message is to inform the Cisco CallManager of two pieces of information about the Skinny Station. First, it contains the status of the open request. Secondly, it contains the receive port number that was opened. The IPAddr and Port information will be used in the next command for the remote end to use when transmitting the RTP audio stream.

```
16:05:45.265 CCM|StationInit - InboundStim - StationOpenReceiveChannelAckID
tcpHandle=0x4fbb150, Status=0, IpAddr=0xe64610ac, Port=17054, PartyID=2
```

The following messages are used by Cisco CallManager to order the station to begin transmitting the audio stream to the indicated remote Cisco IP phone’s IP address and port number.

```
16:05:45.265 CCM|StationD - stationOutputStartMediaTransmission tcpHandle=0x4fbbc30 myIP:
e74610ac (172.16.70.231)
16:05:45.265 CCM|StationD - RemoteIpAddr: e64610ac (172.16.70.230) RemoteRtpPortNumber:
17054 msecPacketSize: 20 compressionType:(4)Media_Payload_G711Ulaw64k
```

In the following traces, the previously explained messages are sent to the called party. These messages are followed by the messages indicating the RTP media stream has begun between the called and calling party.

```
16:05:45.312 CCM|StationD - stationOutputStartMediaTransmission tcpHandle=0x4fbb150 myIP:
e64610ac (172.16.70.230)
16:05:45.328 CCM|StationD - RemoteIpAddr: e74610ac (172.16.70.231) RemoteRtpPortNumber:
18448 msecPacketSize: 20 compressionType:(4)Media_Payload_G711Ulaw64k
16:05:46.203 CCM|StationInit - InboundStim - OnHookMessageID tcpHandle=0x4fbbc30
```

The calling party’s Cisco IP phone finally goes on-hook, which terminates all the control messages between the Skinny Station and Cisco CallManager, as well as the RTP stream between Skinny Stations.

```
16:05:46.203 CCM|StationInit - InboundStim - OnHookMessageID tcpHandle=0x4fbbc30
```

Cisco IOS Gateways

This section provides diagnostic information supporting Cisco CallManager features for monitoring Cisco IOS Gateways. Specific sections include:

- Call Flow Traces
- Debug Messages and Show Commands on the Cisco IOS Gatekeeper
- Debug Messages and Show Commands on the Cisco IOS Gateway
- Cisco IOS Gateway with T1/PRI Interface

- Cisco IOS Gateway with T1/CAS Interface

Call Flow Traces

This section discusses call flow through examples from the Cisco CallManager trace file CCM000000000 (refer to the previous section for the location of the file). The traces in this case study focus only on the call flow itself, as the more detailed trace information has already been explained in the previous case study (initialization, registration, KeepAlive mechanism, and so on).

In this call flow, a Cisco IP phone (directory number 1001) is calling a phone (directory number 3333) located somewhere on the PSTN. Remember that you can follow a device through the trace by looking at the TCP handle value, time stamp, or name of the device. The TCP handle value for the device remains the same until the device is rebooted or goes offline.

In the following traces, the Cisco IP phone (1001) has gone off-hook. The trace shows the unique messages, TCP handle, and the calling number, which is displayed on the Cisco IP phone. There is no called number at this point because the user has not tried to dial any digits.

```
16:05:46.37515:20:18.390 CCM|StationInit - InboundStim – OffHookMessageID
tcpHandle=0x5138d98
15:20:18.390 CCM|StationD - stationOutputDisplayText tcpHandle=0x5138d98, Display=1001
```

In the following traces, the user is dialing the 3333, one digit at a time. The number 3333 is the destination number of the phone, which is located somewhere on the PSTN network. The digit analysis process of the Cisco CallManager is currently active and is analyzing the digits to discover where the call needs to be routed. A more detailed explanation of the digit analysis was provided in the previous case study.

```
15:20:18.390 CCM|Digit analysis: match(fqcn="", cn="1001", pss="", dd="")
15:20:19.703 CCM|Digit analysis: match(fqcn="", cn="1001", pss="", dd="3")
15:20:20.078 CCM|Digit analysis: match(fqcn="", cn="1001", pss="", dd="33")
15:20:20.718 CCM|Digit analysis: match(fqcn="", cn="1001", pss="", dd="333")
15:20:21.421 CCM|Digit analysis: match(fqcn="", cn="1001", pss="", dd="3333")
15:20:21.421 CCM|Digit analysis: analysis results
```

In the following traces, the digit analysis has been completed, calling and called party has been matched, and the information has been parsed.

```
|CallingPartyNumber=1001
|DialingPattern=3333
|DialingRoutePatternRegularExpression=(3333)
|PretransformDigitString=3333
|PretransformPositionalMatchList=3333
|CollectedDigits=3333
|PositionalMatchList=3333
```

In the following traces, the number 0 indicates the originating location, and the number 1 indicates the destination location. The bandwidth of the originating location is determined by BW = –1. The value –1 implies that the bandwidth is infinite. The bandwidth is infinite because the call was originated from a Cisco IP phone located in a LAN environment. The bandwidth of the destination location is determined by BW = 64. The call destination is to a phone located in a PSTN, and the codec type is used is G.711 (64Kbps).

```
15:20:21.421 CCM|Locations:Orig=0 BW=-1 Dest=1 BW=64 (-1 implies infinite bw available)
```

The following traces show the calling and called party information. In this example, the calling party name and number is the same because the administrator has not configured a display name, such as *John Smith*.

```
15:20:21.421 CCM|StationD - stationOutputCallInfo CallingPartyName=1001,
CallingParty=1001, CalledPartyName=, CalledParty=3333, tcpHandle=0x5138d98
```

Before reviewing the following traces, it is important to understand the meaning of the term H.323. By way of a brief explanation, there are several protocols that are used when establishing an H.323 session. One protocol is H.225, which is primarily used for call signaling and is a subset of Q.931. Another protocol is H.245, which is used for capability exchange. One of the more important functions of H.245 is the Compressor/Decompressor (codec) type negotiation, such as G.711, G.729, and so on, between the calling and called side. Once the capability exchange is complete, the next important function of H.245 is performing a UDP port negotiation between the calling and called sides.

The following trace shows that the H.323 code has been initialized and is sending an H.225 setup message. You can also see the traditional HDLC SAPI messages, the IP address of the called side in hex, and the port numbers.

```
15:20:21.421 CCM|Out Message -- H225SetupMsg -- Protocol= H225Protocol
15:20:21.421 CCM|MMan_Id= 1. (iep= 0 dsl= 0 sapi= 0 ces= 0 IpAddr=e24610ac IpPort=47110)
```

The following trace shows the calling and called party information, as well as the H.225 alerting message. Also shown is the mapping of a Cisco IP phone’s hex value to the IP address. 172.16.70.231 is the IP address of the Cisco IP phone (1001).

```
15:20:21.437 CCM|StationD - stationOutputCallInfo CallingPartyName=1001,
CallingParty=1001, CalledPartyName=, CalledParty=3333, tcpHandle=0x5138d98
15:20:21.453 CCM|In Message -- H225AlertMsg -- Protocol= H225Protocol
15:20:21.953 CCM|StationD - stationOutputOpenReceiveChannel tcpHandle=0x5138d98 myIP:
e74610ac (172.16.70.231)
```

The following trace shows the compression type used for this call (G.711 µ-law).

```
15:20:21.953 CCM|StationD - ConferenceID: 0 msecPacketSize: 20
compressionType:(4)Media_Payload_G711Ulaw64k
```

Once the H.225 alert message has been sent, the next part of H.323 is to initialize H.245. The following trace shows the calling and called party information and the H.245 messages. Notice the TCP handle value is the same as before, indicating this is the continuation of the same call.

```
15:20:22.062 CCM|H245Interface(3) paths established ip = e74610ac, port = 23752
15:20:22.062 CCM|H245Interface(3) OLC outgoing confirm ip = e24610ac, port = 16758
15:20:22.062 CCM|MediaManager - wait_AuConnectInfo - received response, forwarding
```

The following trace shows the H.225 connect message, as well as other information that was explained earlier. When the H.225 connect message is received, the call has been connected.

```
15:20:22.968 CCM|In Message -- H225ConnectMsg -- Protocol= H225Protocol
15:20:22.968 CCM|StationD - stationOutputCallInfo CallingPartyName=1001,
CallingParty=1001, CalledPartyName=, CalledParty=3333, tcpHandle=0x5138d98
15:20:22.062 CCM|MediaCoordinator - wait_AuConnectInfoInd
15:20:22.062 CCM|StationD - stationOutputStartMediaTransmission tcpHandle=0x5138d98 myIP:
e74610ac (172.16.70.231)
15:20:22.062 CCM|StationD - RemoteIpAddr: e24610ac (172.16.70.226) RemoteRtpPortNumber:
16758 msecPacketSize: 20 compressionType:(4)Media_Payload_G711Ulaw64k
15:20:22.062 CCM|Locations:Orig=0 BW=-1Dest=1 BW=6(-1 implies infinite bw available)
```

The following message shows that an on-hook message from the Cisco IP phone (1001) is being received. As soon as an on-hook message is received, the H.225 and Skinny disconnect messages are sent and the entire H.225 message is seen. This final message indicates the call has been terminated.

```
15:20:27.296 CCM|StationInit - InboundStim – OnHookMessageID tcpHandle=0x5138d98
15:20:27.296 CCM|ConnectionManager -wait_AuDisconnectRequest (16777247,16777248): STOP
SESSION
15:20:27.296 CCM|MediaManager - wait_AuDisconnectRequest - StopSession sending disconnect
to (64,5) and remove connection from list
15:20:27.296 CCM| Device SEP003094C26105 , UnRegisters with SDL Link to monitor NodeID= 1
```

```
15:20:27.296 CCM|StationD - stationOutputCloseReceiveChannel tcpHandle=0x5138d98 myIP: 
e74610ac (172.16.70.231)
15:20:27.296 CCM|StationD - stationOutputStopMediaTransmission tcpHandle=0x5138d98 myIP: 
e74610ac (172.16.70.231)
15:20:28.328 CCM|In Message -- H225ReleaseCompleteMsg -- Protocol= H225Protocol
```

Debug Messages and Show Commands on the Cisco IOS Gatekeeper

In the previous section, the Cisco CallManager trace was discussed in detail. In the topology for this case study, **debug ras** has been turned on in the Cisco IOS gatekeeper.

The following debug messages show that the Cisco IOS gatekeeper is receiving the admission request (ARQ) for the Cisco CallManager (172.16.70.228), followed by other successful RAS messages. Finally, the Cisco IOS gatekeeper sends an admission confirmed (ACF) message to the Cisco CallManager.

```
*Mar 12 04:03:57.181: RASLibRASRecvData ARQ (seq# 3365) rcvd from [172.16.70.228883] on 
sock [0x60AF038C]
*Mar 12 04:03:57.181: RASLibRAS_WK_TInit ipsock [0x60A7A68C] setup successful
*Mar 12 04:03:57.181: RASlibras_sendto msg length 16 from 172.16.70.2251719 to 
172.16.70.228883
*Mar 12 04:03:57.181: RASLibRASSendACF ACF (seq# 3365) sent to 172.16.70.228
```

The following debug messages show that the call is in progress.

```
*Mar 12 04:03:57.181: RASLibRASRecvData successfully rcvd message of length 55 from 
172.16.70.228883
```

The following debug messages show that the Cisco IOS gatekeeper has received a disengaged request (DRQ) from the Cisco CallManager (172.16.70.228), and the Cisco IOS gatekeeper has sent a disengage confirmed (DCF) to the Cisco CallManager.

```
*Mar 12 04:03:57.181: RASLibRASRecvData DRQ (seq# 3366) rcvd from [172.16.70.228883] on 
sock [0x60AF038C]
*Mar 12 04:03:57.181: RASlibras_sendto msg length 3 from 172.16.70.2251719 to 
172.16.70.228883
*Mar 12 04:03:57.181: RASLibRASSendDCF DCF (seq# 3366) sent to 172.16.70.228
*Mar 12 04:03:57.181: RASLibRASRecvData successfully rcvd message of length 124 from 
172.16.70.228883
```

The **show gatekeeper endpoints** command on the Cisco IOS gatekeeper shows that all four Cisco CallManagers are registered with the Cisco IOS gatekeeper. Remember that in the topology for this case study, there are four Cisco CallManagers, two in each cluster. This Cisco IOS gatekeeper has two zones and each zone has two Cisco CallManagers.

```
R2514-1#show gatekeeper endpoints
GATEKEEPER ENDPOINT REGISTRATION
===================================
CallSignalAddr Port RASSignalAddr Port Zone Name Type
--------------- ----- --------------- ----- --------- ---- --
172.16.70.228 2 172.16.70.228 1493 gka.cisco.com VOIP-GW H323-ID: ac1046e4->ac1046f5
172.16.70.229 2 172.16.70.229 3923 gka.cisco.com VOIP-GW H323-ID: ac1046e5->ac1046f5
172.16.70.245 1 172.16.70.245 1041 gkb.cisco.com VOIP-GW H323-ID: ac1046f5->ac1046e4
172.16.70.243 1 172.16.70.243 2043 gkb.cisco.com VOIP-GW H32
3-ID: ac1046f5->ac1046e4
Total number of active registrations = 4
```

Debug Messages and Show Commands on the Cisco IOS Gateway

In the previous section, the Cisco IOS gatekeeper show commands and debug outputs were discussed in detail. This section focuses on the debug output and show commands on the Cisco IOS gateway. In the topology for this case study, calls are going through the Cisco IOS gateways. The Cisco IOS gateway interfaces to the PSTN or PBX with either T1/CAS or T1/PRI interfaces. Debug output of commands such as **debug voip ccapi inout**, **debug h225 events**, and **debug h225 asn1** are shown below.

In the following debug output, the Cisco IOS gateway is accepting the TCP connection request from Cisco CallManager (172.16.70.228) on port 2328 for H.225.

```
*Mar 12 04:03:57.169: H225Lib::h225TAccept: TCP connection accepted from
172.16.70.228:2328 on socket [1]
*Mar 12 04:03:57.169: H225Lib::h225TAccept: Q.931 Call State is initialized to be [Null].
*Mar 12 04:03:57.177: Hex representation of the received TPKT03000065080000100
```

The following debug output shows that the H.225 data is coming from the Cisco CallManager on this TCP session. In this debug output it is important to notice the protocolIdentifier, which indicates the H.323 version being used. The following debug shows that H.323 version 2 is being used. The called and calling party numbers are also shown.

```
- Source Address H323-ID
- Destination Address e164
*Mar 12 04:03:57.177: H225Lib::h225RecvData: Q.931 SETUP received from socket [1]value
H323-UserInformation ::=
*Mar 12 04:03:57.181: {
*Mar 12 04:03:57.181: h323-uu-pdu
*Mar 12 04:03:57.181: {
*Mar 12 04:03:57.181: h323-message-body setup :
*Mar 12 04:03:57.181: {
*Mar 12 04:03:57.181: protocolIdentifier { 0 0 8 2250 0 2 },
*Mar 12 04:03:57.181: sourceAddress
*Mar 12 04:03:57.181: {
*Mar 12 04:03:57.181: h323-ID : "1001"
*Mar 12 04:03:57.181: },
*Mar 12 04:03:57.185: destinationAddress
*Mar 12 04:03:57.185: {
*Mar 12 04:03:57.185: e164 : "3333"
*Mar 12 04:03:57.185: },
*Mar 12 04:03:57.189: H225Lib::h225RecvData: State changed to [Call Present].
```

The following is debug output for Call Control Application Programming interface (CCAPi). CCAPi is indicating an incoming call. Called and calling party information can also be seen in the following output. CCAPi matches the dial peer 0, which is the default dial peer. It is matching dial peer 0 because the CCAPi could not find any other dial peer for the calling number, so it is using the default dial peer.

```
*Mar 12 04:03:57.189: cc_api_call_setup_ind (vdbPtr=0x616C9F54, callInfo={called=3333,
calling=1001, fdest=1 peer_tag=0}, callID=0x616C4838)
*Mar 12 04:03:57.193: cc_process_call_setup_ind (event=0x617A2B18) handed call to app
"SESSION"
*Mar 12 04:03:57.193: sess_appl: ev(19=CC_EV_CALL_SETUP_IND), cid(17), disp(0)
*Mar 12 04:03:57.193: ccCallSetContext (callID=0x11, context=0x61782BBC)
Mar 12 04:03:57.193: ssaCallSetupInd finalDest cllng(1001), clled(3333)
*Mar 12 04:03:57.193: ssaSetupPeer cid(17) peer list: tag(1)
*Mar 12 04:03:57.193: ssaSetupPeer cid(17), destPat(3333), matched(4), prefix(),
peer(6179E63C)
*Mar 12 04:03:57.193: ccCallSetupRequest (peer=0x6179E63C, dest=, params=0x61782BD0
mode=0, *callID=0x617A87C0)
*Mar 12 04:03:57.193: callingNumber=1001, calledNumber=3333, redirectNumber=
*Mar 12 04:03:57.193: accountNumber=,finalDestFlag=1,
guid=0098.89c8.9233.511d.0300.cddd.ac10.46e6
```

The CCAPi matches the dial-peer 1 with the destination pattern, which is the called number 3333. Remember that peer_tag means dial peer. Notice the calling and called party number in the request packet.

```
*Mar 12 04:03:57.193: peer_tag=1
*Mar 12 04:03:57.197: ccIFCallSetupRequest: (vdbPtr=0x617BE064, dest=,
callParams={called=3333, calling=1001, fdest=1, voice_peer_tag=1}, mode=0x0)
```

The following debug output shows that the H.225 Alerting messages are returning to the Cisco CallManager.

```
*Mar 12 04:03:57.197: ccCallSetContext (callID=0x12, context=0x61466B30)
*Mar 12 04:03:57.197: ccCallProceeding (callID=0x11, prog_ind=0x0)
*Mar 12 04:03:57.197: cc_api_call_proceeding(vdbPtr=0x617BE064, callID=0x12, prog_ind=0x0)
*Mar 12 04:03:57.197: cc_api_call_alert(vdbPtr=0x617BE064, callID=0x12, prog_ind=0x8,
sig_ind=0x1)
*Mar 12 04:03:57.201: sess_appl: ev(17=CC_EV_CALL_PROCEEDING), cid(18), disp(0)
*Mar 12 04:03:57.201: ssa:
cid(18)st(1)oldst(0)cfid(-1)csize(0)in(0)fDest(0)-cid2(17)st2(1)oldst2(0)
*Mar 12 04:03:57.201: ssaIgnore cid(18), st(1),oldst(1), ev(17)
*Mar 12 04:03:57.201: sess_appl: ev(7=CC_EV_CALL_ALERT), cid(18), disp(0)
*Mar 12 04:03:57.201: ssa:
cid(18)st(1)oldst(1)cfid(-1)csize(0)in(0)fDest(0)-cid2(17)st2(1)oldst2(0)
*Mar 12 04:03:57.201: ssaFlushPeerTagQueue cid(17) peer list: (empty)
*Mar 12 04:03:57.201: ccCallAlert (callID=0x11, prog_ind=0x8, sig_ind=0x1)
*Mar 12 04:03:57.201: ccConferenceCreate (confID=0x617A8808, callID1=0x11, callID2=0x12,
tag=0x0)
*Mar 12 04:03:57.201: cc_api_bridge_done (confID=0x7, srcIF=0x616C9F54, srcCallID=0x11,
dstCallID=0x12, disposition=0, tag=0x0)value H323-UserInformation
*Mar 12 04:03:57.201: {
*Mar 12 04:03:57.201: h323-uu-pdu
*Mar 12 04:03:57.201: {
*Mar 12 04:03:57.201: h323-message-body alerting :
*Mar 12 04:03:57.201: {
*Mar 12 04:03:57.201: protocolIdentifier { 0 0 8 2250 0 2 },
*Mar 12 04:03:57.205: destinationInfo
*Mar 12 04:03:57.205: {
*Mar 12 04:03:57.205: mc FALSE,
*Mar 12 04:03:57.205: undefinedNode FALSE
*Mar 12 04:03:57.205: },
```

Notice in this packet that Cisco IOS is also sending the H.245 address and port number to Cisco CallManager. Sometimes the Cisco IOS gateway will send the unreachable address, which could cause either no audio or one-way audio.

```
*Mar 12 04:03:57.205: h245Address ipAddress :
*Mar 12 04:03:57.205: {
*Mar 12 04:03:57.205: ip 'AC1046E2'H,
*Mar 12 04:03:57.205: port 011008
*Mar 12 04:03:57.205: },
*Mar 12 04:03:57.213: Hex representation of the ALERTING TPKT to send.0300003D0100
*Mar 12 04:03:57.213:
*Mar 12 04:03:57.213: H225Lib::h225AlertRequest: Q.931 ALERTING sent from socket [1]. Call
state changed to [Call Received].
*Mar 12 04:03:57.213: cc_api_bridge_done (confID=0x7, srcIF=0x617BE064, srcCallID=0x12,
dstCallID=0x11, disposition=0, tag=0x0)
```

The following debug output shows that the H.245 session is coming up. You can see the capability indication for codec negotiation, as well as how many bytes will be present in each voice packet.

```
*Mar 12 04:03:57.217: cc_api_caps_ind (dstVdbPtr=0x616C9F54, dstCallId=0x11,
srcCallId=0x12, caps={codec=0xEBFB, fax_rate=0x7F, vad=0x3, modem=0x617C5720
codec_bytes=0, signal_type=3})
*Mar 12 04:03:57.217: sess_appl: ev(23=CC_EV_CONF_CREATE_DONE), cid(17), disp(0)
```

```
*Mar 12 04:03:57.217: ssa:
cid(17)st(3)oldst(0)cfid(7)csize(0)in(1)fDest(1)-cid2(18)st2(3)oldst2(1)
*Mar 12 04:03:57.653: cc_api_caps_ind (dstVdbPtr=0x617BE064, dstCallId=0x12,
srcCallId=0x11, caps={codec=0x1, fax_rate=0x2, vad=0x2, modem=0x1, codec_bytes=160,
signal_type=0})
```

The following debug output shows that both parties negotiated correctly and agreed on G.711 codec with 160 bytes of data.

```
*Mar 12 04:03:57.653: cc_api_caps_ack (dstVdbPtr=0x617BE064, dstCallId=0x12,
srcCallId=0x11, caps={codec=0x1, fax_rate=0x2, vad=0x2, modem=0x1, codec_bytes=160,
signal_type=0})
*Mar 12 04:03:57.653: cc_api_caps_ind (dstVdbPtr=0x617BE064, dstCallId=0x12,
srcCallId=0x11, caps={codec=0x1, fax_rate=0x2, vad=0x2, modem=0x, codec_bytes=160,
signal_type=0})
*Mar 12 04:03:57.653: cc_api_caps_ack (dstVdbPtr=0x617BE064, dstCallId=0x12,
srcCallId=0x11, caps={codec=0x1, fax_rate=0x2, vad=0x2, modem=0x1, codec_bytes=160,
signal_type=0})
*Mar 12 04:03:57.657: cc_api_caps_ack (dstVdbPtr=0x616C9F54, dstCallId=0x11,
srcCallId=0x12, caps={codec=0x1, fax_rate=0x2, vad=0x2, modem=0x1, codec_bytes=160,
signal_type=0})
*Mar 12 04:03:57.657: cc_api_caps_ack (dstVdbPtr=0x616C9F54, dstCallId=0x11,
srcCallId=0x12, caps={codec=0x1, fax_rate=0x2, vad=0x2, modem=0x1, codec_bytes=160,
signal_type=0})
```

The H.323 connect and disconnect messages can be seen below.

```
*Mar 12 04:03:59.373: cc_api_call_connected(vdbPtr=0x617BE064, callID=0x12)
*Mar 12 04:03:59.373: sess_appl: ev(8=CC_EV_CALL_CONNECTED), cid(18), disp(0)
*Mar 12 04:03:59.373: ssa:
cid(18)st(4)oldst(1)cfid(7)csize(0)in(0)fDest(0)-cid2(17)st2(4)oldst2(3)
*Mar 12 04:03:59.373: ccCallConnect (callID=0x11)
*Mar 12 04:03:59.373: {
*Mar 12 04:03:59.373: h323-uu-pdu
*Mar 12 04:03:59.373: {
*Mar 12 04:03:59.373: h323-message-body connect :
*Mar 12 04:03:59.373: {
*Mar 12 04:03:59.373: protocolIdentifier { 0 0 8 2250 0 2 },
*Mar 12 04:03:59.373: h245Address ipAddress :
*Mar 12 04:03:59.373: {
*Mar 12 04:03:59.377: ip 'AC1046E2'H,
*Mar 12 04:03:59.377: port 011008
*Mar 12 04:03:59.377: },
*Mar 12 04:03:59.389: Hex representation of the CONNECT TPKT to send.03000052080
*Mar 12 04:03:59.393: H225Lib::h225SetupResponse: Q.931 CONNECT sent from socket [1]
*Mar 12 04:03:59.393: H225Lib::h225SetupResponse: Q.931 Call State changed to [Active].
*Mar 12 04:04:08.769: cc_api_call_disconnected(vdbPtr=0x617BE064, callID=0x12, cause=0x10)
*Mar 12 04:04:08.769: sess_appl: ev(12=CC_EV_CALL_DISCONNECTED), cid(18), disp(0)
```

Cisco IOS Gateway with T1/PRI Interface

As explained earlier, there are two types of calls going through the Cisco IOS gateways: the Cisco IOS gateway interfaces to the PSTN or PBX, with either T1/CAS or T1/PRI interfaces. The following are the debug outputs when the Cisco IOS gateways use T1/PRI interface.

The **debug isdn q931** command on the Cisco IOS gateway has been turned on. This enables Q.931, a Layer Three signaling protocol for D-channel in the ISDN environment. Each time a call is placed out of the T1/PRI interface, a setup packet must be sent. The setup packet always has (protocol descriptor) pd = 8, and it generates a random hex value for the callref. The callref is used to track the call. For example, if two calls are placed, the callref value can determine the call for which the RX (received) message is intended. Bearer capability 0x8890 means a 64kb/s data call. If it were a 0x8890218F, it would be a 56kb/s data call. It would be 0x8090A3 if it were a voice call. In the debug output below, the bearer capability is 0x8090A3, which is for voice. Called and calling party numbers are also shown.

The callref uses the high-order bit of the first byte to differentiate between TX and RX. The bit set to 0 inidcates *sender* (originator) and a value of 1 indicates *receiver*. The router is completely dependent upon the PSTN or PBX to assign a Bearer channel (B-channel). If the PSTN or PBX doesn't assign a channel to the router, the call won't be routed. In our case, a CONNECT message is received from the switch with the same reference number as was received for ALERTING (0x800B). Finally, you can see the exchange of the DISCONNECT message followed by the RELEASE and RELEASE _COMP messages as the call is being disconnected. RELEASE_COMP messages are followed by a cause ID for the call rejection. The cause ID is a hex value. The meaning of the cause can be found by decoding the hex value and following up with your provider.

```
*Mar 1 225209.694 ISDN Se115 TX -> SETUP pd = 8 callref = 0x000B
*Mar 1 225209.694 Bearer Capability i = 0x8090A3
*Mar 1 225209.694 Channel ID i = 0xA98381
*Mar 1 225209.694 Calling Party Number i = 0x2183, ’1001'
*Mar 1 225209.694 Called Party Number i = 0x80, ’3333'
*Mar 1 225209.982 ISDN Se115 RX <- ALERTING pd = 8 callref = 0x800B
*Mar 1 225209.982 Channel ID i = 0xA98381
*Mar 1 225210.674 ISDN Se115 RX <- CONNECT pd = 8 callref = 0x800B
*Mar 1 225210.678 ISDN Se115 TX -> CONNECT_ACK pd = 8 callref = 0x000B
*Mar 1 225215.058 ISDN Se115 RX <- DISCONNECT pd = 8 callref = 0x800B
*Mar 1 225215.058 Cause i = 0x8090 - Normal call clearing 225217 %ISDN-6
DISCONNECT Int S10 disconnected from unknown , call lasted 4 sec
*Mar 1 225215.058 ISDN Se115 TX -> RELEASE pd = 8 callref = 0x000B
*Mar 1 225215.082 ISDN Se115 RX <- RELEASE_COMP pd = 8 callref = 0x800B
*Mar 1 225215.082 Cause i = 0x829F - Normal, unspecified or Special intercept, call
blocked group restriction
```

Cisco IOS Gateway with T1/CAS Interface

As explained earlier, there are two types of calls going through the Cisco IOS gateways: the Cisco IOS gateway interface to the PSTN or PBX, with either T1/CAS or T1/PRI interfaces. The following are the debug outputs when the Cisco IOS gateways has T1/CAS interface. The **debug cas** on the Cisco IOS gateway has been turned on.

The following debug message shows that the Cisco IOS gateway is sending an off-hook signal to the switch.

```
Apr 5 17:58:21.727: from NEAT(0): (0/15): Tx LOOP_CLOSURE (ABCD=1111)
```

The following debug message indicates the switch is sending wink after receiving the loop closure signal from the Cisco IOS gateway.

```
Apr 5 17:58:21.859: from NEAT(0): (0/15): Rx LOOP_CLOSURE (ABCD=1111)
Apr 5 17:58:22.083: from NEAT(0): (0/15): Rx LOOP_OPEN (ABCD=0000)
```

The following debug message indicates the Cisco IOS gateway is going off-hook.

```
Apr 5 17:58:23.499: from NEAT(0): (0/15): Rx LOOP_CLOSURE (ABCD=1111)
```

The following is the output of the **show call active voice brief** on the Cisco IOS gateway when the call is in progress. The called and calling party number and other useful information are also shown.

```
R5300-5#show call active voice brief
<ID>: <start>hs.<index> +<connect> pid:<peer_id> <dir> <addr> <state> tx:<packets>/<bytes>
rx:<packets>/<bytes> <state>
IP <ip>:<udp> rtt:<time>ms pl:<play>/<gap>ms lost:<lost>/<early>/<late>
delay:<last>/<min>/<max>ms <codec>
FR <protocol> [int dlci cid] vad:<y/n> dtmf:<y/n> seq:<y/n> sig:<on/off> <codec> (payload
size)
```

```
Tele <int>: tx:<tot>/<v>/<fax>ms <codec> noise:<l> acom:<l> i/o:<l>/<l> dBm
511D : 156043737hs.1 +645 pid:0 Answer 1001 active
tx:1752/280320 rx:988/158080
IP172.16.70.228:18888 rtt:0ms pl:15750/80ms lost:0/0/0 delay:25/25/65ms g711ulaw
511D : 156043738hs.1 +644 pid:1 Originate 3333 active
tx:988/136972 rx:1759/302548
Tele 1/0/0 (30): tx:39090/35195/0ms g711ulaw noise:-43 acom:0 i/0:-36/-42 dBm
```

Experiencing Busy Signal After Dialing International Numbers

Problem

User has CM 3.0 with a proper route pattern assigned to a DT-24+ or 6608 gateway. He is able to dial local and US long distance numbers with no problems. However when he punches the digits for an international number, he hears a pause, and then a busy signal.

Solution

This has been seen in the past where the CO switch does not know how to deal properly with the call IE (information element). This can be fixed by setting the Cisco CallManager Calling party IE type to `unknown` under the Gateway configuration as illustrated in Figure 6-55.

Figure 6-55 Setting Calling Party IE Type

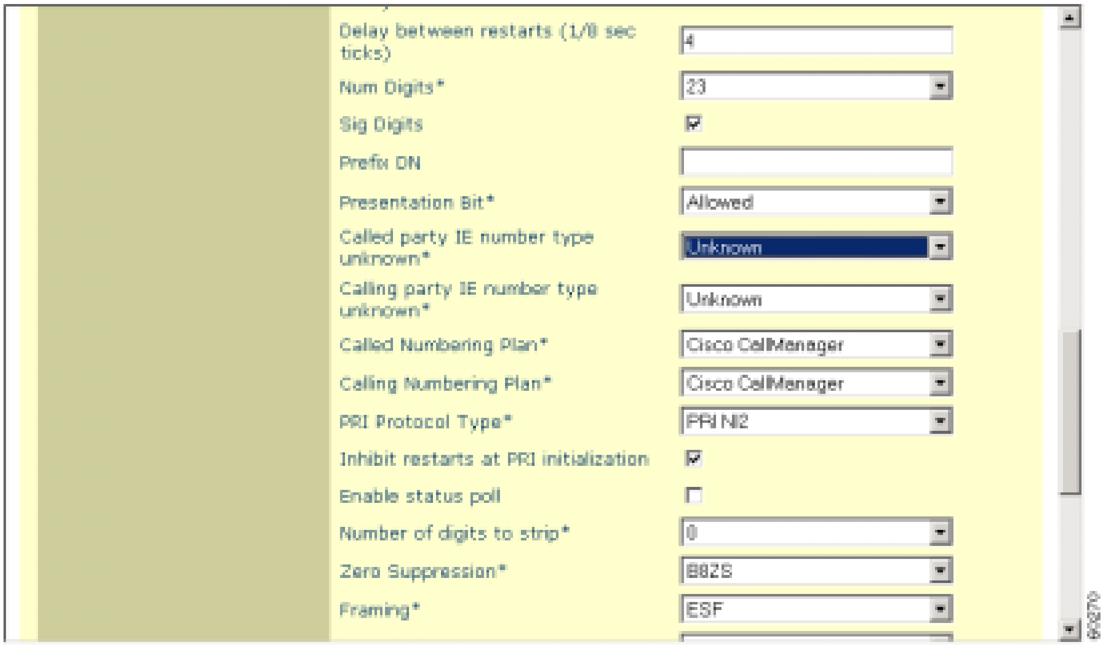

Delay between restarts (1/8 sec ticks)	4
Num Digits*	23
Sig Digits	☑
Prefix DN	
Presentation Bit*	Allowed
Called party IE number type unknown*	Unknown
Calling party IE number type unknown*	Unknown
Called Numbering Plan*	Cisco CallManager
Calling Numbering Plan*	Cisco CallManager
PRI Protocol Type*	PRI NI2
Inhibit restarts at PRI initialization	☑
Enable status poll	☐
Number of digits to strip*	0
Zero Suppression*	B8ZS
Framing*	ESF

Inter-Cluster IP Phone to IP Phone

In the previous case studies, the call flow and troubleshooting techniques of an intra-cluster call and a Cisco IP phone call through a Cisco IOS gateway to a phone hanging off of a local PBX or somewhere on the PSTN have been discussed in detail. This case study examines a Cisco IP phone calling another Cisco IP phone located in a different cluster. This type of call is also known as an inter-cluster Cisco IP phone call.

Sample Topology

Figure 6-56 illustrates and example topology used in this case study. This topology has two clusters in a Campus, each having two Cisco CallManagers. There are also Cisco IOS gateways and a Cisco IOS gatekeeper in place.

Figure 6-56 Example Inter-Cluster Topology

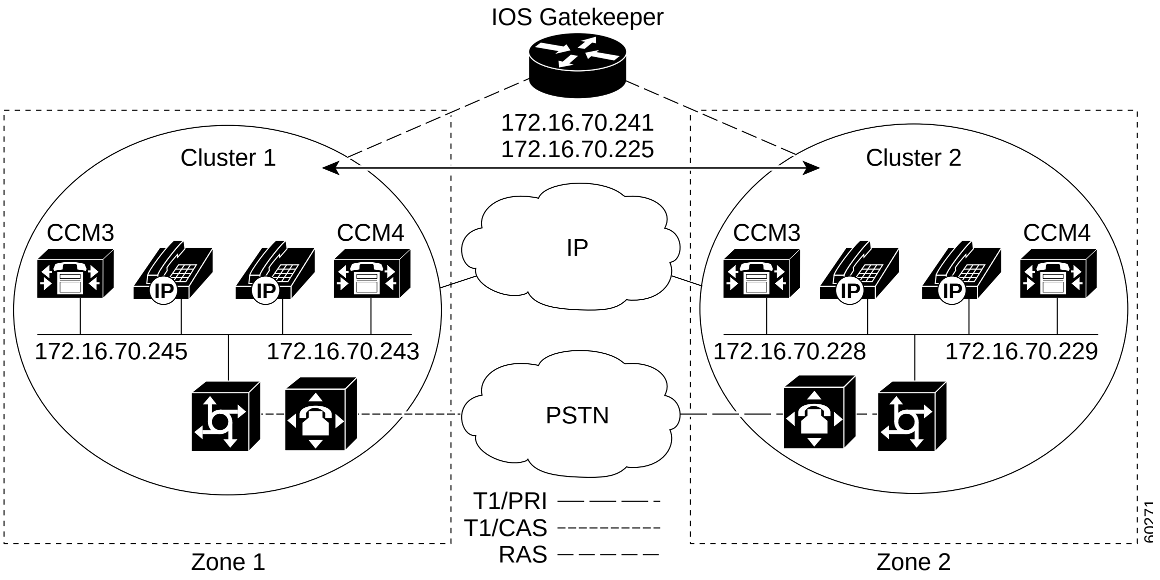

Inter-Cluster H.323 Communication

As you can see in the topology, the Cisco IP phone in Cluster 1 in Figure 6-56 is making a call to the Cisco IP phone in Cluster 2. Inter-cluster Cisco CallManager communication takes place using the H.323 Version 2 protocol. There is also a Cisco IOS gatekeeper for admission control. The detailed explanation of the debug output and show commands, and the interaction between the Cisco IOS gatekeeper and Cisco IOS gateway and Cisco CallManager devices, can be reviewed in the previous sections.

The call flow process is shown in Figure 6-57. The Cisco IP phone can talk to the Cisco CallManager via Skinny Station protocol, and the Cisco CallManager can talk with the Cisco IOS gatekeeper using the H.323 RAS protocol. The Admission Request message (ARQ) is sent to the Cisco IOS gatekeeper, which sends the Admission Confirmed message (ACF) after making sure that the inter-cluster call can be made using H.323 version 2 protocol. Once done, the audio path is made using the RTP protocol between Cisco IP phones in different clusters.

Figure 6-57 Call Flow Process for Inter-Cluster Topology

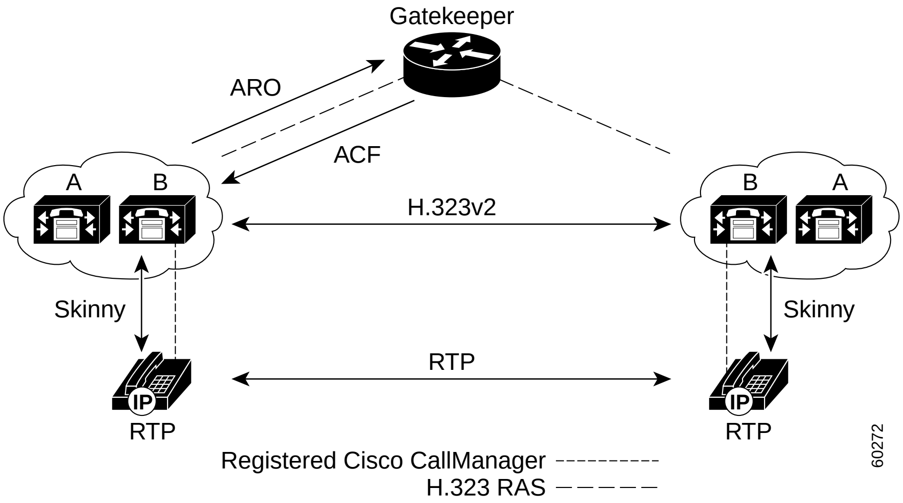

Call Flow Traces

This section discusses the call flow using CallManager trace examples captured in the CCM000000000 file. The location of this file can be found in the previous section. The traces discussed in this case study focus only on the call flow itself, because the more detailed trace information has already been explained in the previous case study (initialization, registration, KeepAlive mechanism, and so on.)

In this call flow, a Cisco IP phone (2002) located in Cluster 2 is calling a Cisco IP phone (1001) located in Cluster 1. Remember that you can follow a device through the trace by looking at the TCP handle value, time stamp, or name of the device. The TCP handle value for the device remains the same until the device is rebooted or goes offline.

In the following traces, the Cisco IP phone (2002) has gone off-hook. The trace shows the unique messages, TCP handle, and the calling number, which is displayed on the Cisco IP phone. The called number (1001), H.225 connect, and H.245 confirm messages can be seen in the following debug output. The codec type is G.711 µ-law.

```
16:06:13.921 CCM|StationInit - InboundStim - OffHookMessageID tcpHandle=0x1c64310
16:06:13.953 CCM|Out Message -- H225ConnectMsg -- Protocol= H225Protocol
16:06:13.953 CCM|Ie - H225UserUserIe IEData= 7E 00 37 05 02 C0 06
16:06:13.953 CCM|StationD - stationOutputCallInfo CallingPartyName=, CallingParty=2002,
CalledPartyName=1001, CalledParty=1001, tcpHandle=0x1c64310
16:06:14.015 CCM|H245Interface(2) OLC indication chan number = 2
16:06:14.015 CCM|StationD - stationOutputOpenReceiveChannel tcpHandle=0x1c64310 myIP:
e74610ac (172.16.70.231)
16:06:14.015 CCM|StationD - ConferenceID: 0 msecPacketSize: 20
compressionType:(4)Media_Payload_G711Ulaw64k
16:06:14.062 CCM|StationInit - InboundStim - StationOpenReceiveChannelAckID
tcpHandle=0x1c64310, Status=0, IpAddr=0xe74610ac, Port=20444, PartyID=2
16:06:14.062 CCM|H245Interface(2) paths established ip = e74610ac, port = 20444
16:06:14.187 CCM|H245Interface(2) OLC outgoing confirm ip = fc4610ac, port = 29626
```

The calling and called party number, which is associated with an IP address and a hex value, can be seen in the following traces.

```
16:06:14.187 CCM|StationD - stationOutputStartMediaTransmission tcpHandle=0x1c64310 myIP:
e74610ac (172.16.70.231)
16:06:14.187 CCM|StationD - RemoteIpAddr: fc4610ac (172.16.70.252)
```

The following traces show the packet sizes and the MAC address of the Cisco IP phone (2002). These traces are followed by the disconnect, then on-hook messages.

```
RemoteRtpPortNumber: 29626 msecPacketSize: 20 compressionType:(4)Media_Payload_G711Ulaw64k
16:06:16.515 CCM| Device SEP003094C26105 , UnRegisters with SDL Link to monitor NodeID= 1
16:06:16.515 CCM|StationD - stationOutputCloseReceiveChannel tcpHandle=0x1c64310 myIP:
e74610ac (172.16.70.231)
16:06:16.515 CCM|StationD - stationOutputStopMediaTransmission tcpHandle=0x1c64310 myIP:
e74610ac (172.16.70.231)
16:06:16.531 CCM|In Message -- H225ReleaseCompleteMsg -- Protocol= H225Protocol
16:06:16.531 CCM|Ie - Q931CauseIe -- IEData= 08 02 80 90
16:06:16.531 CCM|Ie - H225UserUserIe -- IEData= 7E 00 1D 05 05 80 06
16:06:16.531 CCM|Locations:Orig=1 BW=64 Dest=0 BW=-1 (-1 implies infinite bw available)
16:06:16.531 CCM|MediaManager - wait_AuDisconnectRequest - StopSession sending disconnect
to (64,2) and remove connection from list
16:06:16.531 CCM|MediaManager - wait_AuDisconnectReply - received all disconnect replies,
forwarding a reply for party1(16777219) and party2(16777220)
16:06:16.531 CCM|MediaCoordinator - wait_AuDisconnectReply - removing MediaManager(2) from
connection list
16:06:16.734 CCM|StationInit - InboundStim - OnHookMessageID tcpHandle=0x1c64310
```

Failed Call Flow

The following section describes an unsuccessful inter-cluster call flow, as seen in the CallManager trace. In the traces below, the Cisco IP phone (1001) has gone off-hook. A TCP handle is assigned to the Cisco IP phone.

```
16:05:33.468 CCM|StationInit - InboundStim - OffHookMessageID tcpHandle=0x4fbbc30
16:05:33.468 CCM|StationD - stationOutputDisplayText tcpHandle=0x4fbbc30, Display= 1001
16:05:33.484 CCM|StationD - stationOutputSetLamp stim: 9=Line instance=1 lampMode=LampOn
tcpHandle=0x4fbbc30
```

In the following traces the user is dialing the called number (2000) of the Cisco IP phone, and the process of digit analysis is trying to match the number.

```
16:05:33.484 CCM|Digit analysis: match(fqcn="", cn="1001", pss="", dd="")
16:05:33.484 CCM|Digit analysis: potentialMatches=PotentialMatchesExist
16:05:35.921 CCM|Digit analysis: match(fqcn="", cn="1001", pss="", dd="2")
16:05:35.921 CCM|Digit analysis:potentialMatches=ExclusivelyOffnetPotentialMatchesExist
16:05:36.437 CCM|Digit analysis: match(fqcn="", cn="1001", pss="", dd="20")
16:05:36.437 CCM|Digit analysis:potentialMatches=ExclusivelyOffnetPotentialMatchesExist
16:05:36.656 CCM|Digit analysis: match(fqcn="", cn="1001", pss="", dd="200")
16:05:36.656 CCM|Digit analysis:potentialMatches=ExclusivelyOffnetPotentialMatchesExist
16:05:36.812 CCM|Digit analysis: match(fqcn="", cn="1001", pss="", dd="2000")
```

The digit analysis has now been completed and the results are shown in the following traces. It is important to note that the `PotentialMatches=NoPotentialMatchesExist` reference below indicates that the Cisco CallManager is unable to match this directory number. Finally, a reorder tone is sent to the calling party (1001), which is followed by an on-hook message.

```
16:05:36.812 CCM|Digit analysis: analysis results
16:05:36.812 CCM||PretransformCallingPartyNumber=1001
|CallingPartyNumber=1001
|DialingPattern=2XXX
|DialingRoutePatternRegularExpression=(2XXX)
|PotentialMatches=NoPotentialMatchesExist
|CollectedDigits=2000
16:05:36.828 CCM|StationD - stationOutputCallInfo CallingPartyName=1001,
CallingParty=1001, CalledPartyName=, CalledParty=2000, tcpHandle=0x4fbbc30
16:05:36.828 CCM|StationD - stationOutputStartTone: 37=ReorderTone tcpHandle=0x4fbbc30
16:05:37.953 CCM|StationInit - InboundStim - OnHookMessageID tcpHandle=0x4fbbc30
```

Call Detail Records

This section provides detailed information about Call Detail Records (CDRs) and Call Management Records (CMRs, also known as Diagnostic CDRs).

CDR records are written to a database for use in post-processing activities. These activities include many functions, but will primarily be billing and network analysis.

The database is a Microsoft SQL Server 7.0 database. Access to the database can be made via Open DataBase Connectivity (ODBC).

Access is provided to all tables in the database in a read-only fashion and to the CDR and CMR tables in a read/write fashion.

To use CDR record data, you may want to read other tables in the database in an effort to obtain information about the type of device the CDR is about. This correlation between devices in the Device table and the IP address listed in the CDR record is not straightforward and is listed as a known issue later in this section.

Specific descriptions presented in this section include:

- Writing Records
- Reading Records
- Removing Records
- Table Schema
- Known Issues
- Fields in a Call Detail Record
- Call Records Logged By Call Type
- Call Management Records Logged By Call Type
- Codec Types (Compression / Payload Types)
- Cause Codes
- Alarms

Writing Records

Cisco CallManager writes CDR records to the SQL database as calls are made in a manner consistent with the configuration of each individual Cisco CallManager. This configuration is made via the Service Parameters screen in Cisco CallManager Administration.

All records are eventually written to the primary database for the cluster. Subscribers write the CDR data to their local database, and then replicate the data to the publisher database every 60 seconds by default. If the primary database is not available, the records will be written to any of the other backup databases. Once the primary database becomes available, then writing new records will continue on the primary database and the locally-written records will be moved to the primary.

Reading Records

The easiest way to read data from the SQL database may be to use ODBC. A good connection string would look like:

```
DRIVER={SQL Server};SERVER=machineX;DATABASE=CCM0300
```

Be sure to use the correct database name. If a Cisco CallManager Release 3.0(1) version of the software is installed over an existing installation, the database might be migrated if called for by the new installation. In this case the old database will still exist, and the new database will also exist. The names will differ by adding one to the number of the name. For instance, the original name is CCM0300. After a migration, the newer database name will be CCM0301. The highest number database should be used.

The primary database (machine and name) currently in use by the cluster can be found by clicking on the Details button of Cisco CallManager Administration (click Help to reach the Welcome screen where the Details button is located). The registry on machines hosting a database can also be checked. Look at the registry key: HKEY_LOCAL_MACHINE/Software/Cisco Systems Inc/DBL for the item called DBConnection0. This string item contains a connection string similar to that shown above with the machine name and database name of the primary database.

Access is controlled by use of SQL Users. Table 6-35 specifies the User ID and password that should be used when accessing the Cisco CallManager database.

Table 6-35 Example User ID and Password for Cisco CallManager Database

Tables	SQl UserID	Password	Capability
CallDetailRecord, CallDetailRecordDiagnostic	CiscoCCMCDR	dipsy	Read/Write
(Other)	CiscoCCMReader	cowboys	Read only

Removing Records

Since Cisco CallManager is relying on third party applications to post-process the CDR data, you should remove the CDR data when all applications are finished with the data. Since this involves modifying the database, the CiscoCCMCDR user should be used.

If CDR records accumulate to a configured maximum (10,000,000 CDR records), the oldest CDR records will be removed along with related CMR records once per day.

When removing CDR data after analysis, be sure to remove all related CMR records, as well.

Table Schema

Detailed information about the format and use of each field in the CDR is provided later in this section.

The primary tables used are the CallDetailRecord table (which holds CDR records) and the CallDetailRecordDiagnostic table (which holds CMR records). The CallDetailRecord table is related to the CallDetailRecordDiagnostic table via the two GlobalCallID columns, GlobalCallID_callManagerId, and GlobalCallID_callId. There may be more than one CMR per CDR.

The CallDetailRecord table holds information about the endpoints of the call and other call control/routing aspects of the call. The CallDetailRecordDiagnostic table holds information about the quality of the streamed audio of the call.

Known Issues

Cisco CallManager Release 3.0(1) has several known issues with the CDR data. A few of these are listed below.

IP to Device Name Translation

The CDR table lists IP addresses for the endpoints of a call. These IP addresses are not easily converted to device names so that the type of device can be determined.

On-net vs. Off-net

It is difficult to know if the call stayed completely on the IP network, or at least internal to the local system. One clue is to check the device type of both ends of the call. If both are phones, then you can assume that it stayed on-net. However, if one is a gateway, more assumptions must be made. If the gateway is an Analog Access type of device with a POTS or station port, the call might have just gone to a local analog phone, or might have gone out to the PSTN. Look at the number dialed and correlate this to the known dial plan to estimate if the call went off-net. Otherwise, the call probably went off-net.

Of-net Digits Dialer

If a call is placed out a gateway, the digits dialed to get to the gateway may not be the digits sent to the PSTN. The gateway may be intelligent and modify the directory number further. If this is the case, Cisco CallManager does not know, and the CDR will not reflect the actual digits sent Off-net.

Fields in a Call Detail Record

This section defines all of the fields in the current records. The field types are those used by Cisco CallManager, and not necessarily those defined in the CDR record in the database. The database field definitions are adequate to store the data, but the interpretation of the data should take into account the field types defined here.

All unsigned integers are 32bit unsigned integers.

Field Data Conversions

There are some fields that require conversion from decimal format to another format for displays. This section defines their values, and how to convert them or where to get information on how to convert them.

Time Values

All time values are represented as unsigned 32 bit integers. This unsigned integer value is displayed from the database as a signed integer.

This field is a `time_t` value that is obtained from the Windows NT (2000) system routines. The value is a coordinated universal time (UTC) value and represents the number of seconds since Midnight (00:00:00) Jan. 1, 1970.

Deciphering the Time Stamp

Using Microsoft Excel, you can write a formula to make converting this time stamp a little easier. If the value is in cell A1, you can make another cell:

```
=A1/86400+DATE(1970,1,1)
```

There are 86400 seconds in a day.

Then, format the resulting cell as a date/time field in Excel.

IP Addresses

All IP addresses are stored in the system as unsigned integers. The database displays them as signed integers. To convert the signed decimal value to an IP address, first convert the value to a Hex number (taking into consideration that it is really an unsigned number). The 32bit Hex value represents four bytes. The four bytes are in reverse order (Intel standard). To get the IP address, reverse the order of the bytes and convert each byte to a decimal number. The resulting four bytes represent the four-byte fields of the IP address in dotted notation.

Note The database displays it as a negative number when the low byte of the IP address has the most significant bit set.

Converting IP Addresses

Example 1:

- For example, IP Address 192.168.18.188 would be displayed as follows:
- Database Display = -1139627840.
- This converts to a Hex value of 0xBC12A8C0.
- Reverse the Hex bytes = C0A812BC
- CO A8 12 BC
- Bytes Converted from Hex to Decimal = 192 168 18 188, which would be displayed as 192.168.18.188.

Example 2:

- IP Address 192.168.18.59
- Database Display = 991078592
- This converts to a Hex value of 0x3B12A8C0
- Reverse Byte order = C0A8123B
- C0 A8 12 3B
- Bytes Converted from Hex to Decimal = 192 168 18 59 which would be displayed as 192.168.18.59.

CDR Field Definition

Table 6-36 provides field definitions for CDRs.

Table 6-36 CDR Field Definitions

Field	Definition
cdrRecordType	Type of this record Unsigned integer Specifies the type of this specific record. It could be a Start call record(0), End call r7ecord(1), or a CMR record(2).
globalCallIdentifier	Global Call Identifier The Global Call Identifier consists of two fields which are both unsigned integers. The values must be treated as unsigned integers. The two fields are: • Unsigned integer GlobalCallID_CallID • Unsigned integer GlobalCallID_CallManagerID This is the call identifier that is assigned to the entire call. All records associated with a standard call will have the same global call identifier.
origLegCallIdentifier	Origination leg call identifier Unsigned integer This is a unique identifier that is used to track the origination leg of a call. It is unique within a cluster.
dateTimeOrigination	Date/time of call origination Unsigned integer This represents the time that the call's originating device went off hook, or the time that an outside call was first recognized by the system (it received the Setup message). The value is a coordinated universal time (UTC) value, and represents the number of seconds since Midnight (00:00:00) Jan. 1, 1970.
origNodeId	Originator’s node ID Unsigned integer This field represents the node within the Cisco CallManager cluster where the call originator was registered at the time of this call.
origSpan	Originator’s span or port Unsigned integer This field contains the originator’s port or span number if the call originated through a gateway. If not, this field contains zero (0).
callingPartyNumber	Calling party number Up to 25 characters This is the directory number of the device from which the call originated.
origIpPort	Calling party’s IP port Unsigned integer This field contains the IP port of the device from which the call originated.

Table 6-36 CDR Field Definitions (continued)

Field	Definition
origIpAddr	Calling party's IP address Unsigned integer This field contains the IP address of the device from which the call originated.
originalCallingPartyNumberPartition	Calling party's partition Up to 50 characters This field contains the partition associated with the calling party.
origCause_Location	ISDN location value Unsigned integer This field contains the location value from the cause information element.
origCause_Value	Calling party cause Of call termination Unsigned integer This cause represents the reason the call to the originating device was terminated. In the case of transfers, forwards, and so on, the cause of call termination may be different for the originating device and the termination device. Thus, there are two cause fields associated with each call. Usually they will be the same.
origMediaTransportAddress_IP	The IP address for the originator's media connection Unsigned integer This is the destination IP Address to which the Media Stream from the originator was connected.
origMediaTransportAddress_Port	The port for the originator's media connection Unsigned integer This is the destination port to which the Media Stream from the originator was connected.
origMediaCap_payloadCapability	The codec type used by the originator Unsigned integer This field contains the Codec type (compression or payload type) that the originator used on the sending side during this call. It may be different than the codec type used on its receiving side.
origMediaCap_maxFramesPerPacket	The number of milliseconds of data per packet Unsigned integer This field contains the number of milliseconds of data per packet sent to the destination, by the originator of this call. The actual data size depends on the codec type being used to generate the data.
origMediaCap_g723BitRate	The bit rate to be used by G.723 Unsigned integer Defines the bit rate to be used by G.723. There are two bit rate values: 1 =5.3K bit rate and 2 = 6.3K bit rate.

Table 6-36 CDR Field Definitions (continued)

Field	Definition
lastRedirectDn	Directory number of the party that last redirected this call Up to 25 characters This is the directory number of the last device that redirected this call. This field applies only to calls that were redirected, such as conference calls, call forwarded calls, and so on.
lastRedirectDnPartition	Partition of the phone that last redirected this call Up to 50 characters This is the Partition of the last device that redirected this call. This field applies only to calls that were redirected such as conference calls, call forwarded calls, and so on.
destLegIdentifier	The call identifier for the destination leg of the call Unsigned integer This is a unique identifier that is used to track the destination leg of this call. It is unique within a cluster.
destNodeId	The node identifier for the node where the destination of the call was registered Unsigned integer The node within the Cisco CallManager cluster where the destination device was registered at the time of this call.
dest Span	The destination span or port Unsigned integer This field contains the destination port or span number if the call was terminated through a gateway. If not, this field contains a (0) zero.
destIpAddr	The IP address to which the call was delivered Unsigned integer This field contains the IP address of the signaling connection on the device that terminated the call.
destIpPort	The IP port to which the call was delivered Unsigned integer This field contains the IP port of the signaling connection on the device that terminated the call.
originalCalledPartyNumber	The destination received from the call originator Up to 25 characters This field contains the Directory Number to which the call was originally extended based on the digits dialed by the originator of the call. If the call completes normally (meaning it was not forwarded), this Directory Number should always be the same as the `finalCalledPartyNumber`. If the call was forwarded, this field contains the original destination of the call before it was forwarded.
originalCalledPartyNumberPartition	Called party's partition Up to 50 characters This field contains the partition associated with the called party.

Table 6-36 CDR Field Definitions (continued)

Field	Definition
finalCalledPartyNumber	The destination to which the call was delivered Up to 25 characters This field contains the Directory Number to which the call was actually extended. If the call completes normally (meaning it was not forwarded), this Directory Number should always be the same as the `originalCalledPartyNumber`. If the call was forwarded, this field contains the Directory Number of the final destination of the call after all forwards were completed.
finalCalledPartyNumberPartition	The partition associated with the final destination of the call. Up to 50 characters This field contains the partition associated with the destination to which the call was actually extended. In a normal call, this field should be the same as `originalCalledPartyNumberPartition`. If the call was forwarded, this field contains the partition of the final destination of the call after all forwards were completed.
destCause_location	Called party cause location Unsigned integer This is the ISDN Location value from the Cause Information Element.
destCause_value	Called party cause of call termination Unsigned integer This cause represents why the call to the termination device was terminated. In the case of transfers, forwards, and so on, the cause of call termination may be different for the recipient of the call and the originator of the call. Thus, there are two cause fields associated with each call. Usually they will be the same. When an attempt is made to extend a call to a busy device that is forwarded, the cause code will reflect `Busy` even though the call was connected to a forward destination.
destMediaTransportAddress_IP	The IP address for the destination outgoing media connection Unsigned integer This is the origination IP Address from which the Media Stream from the destination was connected.
origMediaTransportAddress_Port	The port for the destination outgoing media connection Unsigned integer This is the originator's port from which the Media Stream from the destination was connected.
destMediaCap_payloadCapability	The codec type used by the destination on sending side Unsigned integer This field contains the Codec type (compression or payload type) that the destination used on its sending side during this call. It may be different than the codec type used on its receiving side.

Table 6-36 CDR Field Definitions (continued)

Field	Definition
destMediaCap_maxFramesPerPacket	The number of milliseconds of data per packet Unsigned integer This field contains the number of milliseconds of data per packet sent to the originator by the destination of this call. The actual data size depends on the codec type being used to generate the data.
destMediaCap_g723BitRate	The bit rate to be used by G.723 Unsigned integer Defines the bit rate to be used by G.723. There are two bit rate values: 1 =5.3K bit rate and 2 = 6.3K bit rate.
dateTimeConnect	Date/time of connect Unsigned integer This is the date and time that the call was connected between the originating and terminating devices. The value is a coordinated universal time (UTC) value, and represents the number of seconds since Midnight (00:00:00) Jan. 1, 1970.
dateTimeDisconnect	Date/time of disconnect Unsigned integer This is the time that the call was disconnected between the originating and terminating devices, or when the call was torn down even if it was never connected. The value is a coordinated universal time (UTC) value, and represents the number of seconds since Midnight (00:00:00) Jan. 1, 1970.
duration	Call duration This is the number of seconds that the call was connected. It is the difference between the date/time of connect and the date/time of disconnect.

CMR Field Definitions

Table 6-37 provides field definitions for CMRs (diagnostic CDRs).

Table 6-37 Field Definitions

Field	Definition
cdrRecordType	Type of this record Unsigned integer Specifies the type of this specific record. It will be set to CMR record.
globalCallIdentifier	Global Call Identifier for this call The Global Call Identifier consists of two fields which are both unsigned integers. The values must be treated as unsigned integers. The two fields are: This is the call identifier that is assigned to the entire call. All records associated with a standard call will have the same global call identifier.

Table 6-37 Field Definitions

Field	Definition
nodeID	The Cisco CallManager node identifier The node within the Cisco CallManager cluster where this record was generated.
callIdentifier	Call Identifier Unsigned integer This is a call leg identifier that identifies to which call leg this record pertains.
directoryNum	Directory number used on this call This is the directory number of the device from which these diagnostics were collected.
directoryNumPartition	The partition associated with the directory number This is the partition of the directory number in this record.
dateTimeStamp	Date/time of call termination This represents the approximate time that the device went on hook. The time is put into the record when the phone responds to a request for diagnostic information. This is a `time_t` value.
numberPacketsSent	Number of packets sent The total number of RTP data packets transmitted by the device since starting transmission on this connection. The value is zero if the connection was set in *receive only* mode.
numberOctetsSent	Number of Octets (bytes) of data sent to the other party The total number of payload octets (that is, not including header or padding) transmitted in RTP data packets by the device since starting transmission on this connection. The value is zero if the connection was set in *receive only* mode.
numberPacketsReceived	The number of data packets received during this call The total number of RTP data packets received by the device since starting reception on this connection. The count includes packets received from different sources if this is a multicast call. The value is zero if the connection was set in *send only* mode.
numberOctetsReceived	The number of octets (bytes) of data received during this call The total number of payload octets (that is, not including header or padding) received in RTP data packets by the device since starting reception on this connection. The count includes packets received from different sources, if this is a multicast call. The value is zero if the connection was set in *send only* mode.
numberPacketsLost	Lost RTP packets during this connection The total number of RTP data packets that have been lost since the beginning of reception. This number is defined as the number of packets expected, less the number of packets actually received, where the number of packets received includes any that are late or duplicates. Thus, packets that arrive late are not counted as lost, and the loss may be negative if there are duplicates. The number of packets expected is defined to be the extended last sequence number received, as defined next, less the initial sequence number received. The value is zero if the connection was set in *send only* mode. (For details, see RFC 1889)

Table 6-37 Field Definitions

Field	Definition
jitter	The inter-arrival jitter during this connection An estimate of the statistical variance of the RTP data packet inter-arrival time, measured in milliseconds and expressed as an unsigned integer. The inter-arrival jitter J is defined to be the mean deviation (smoothed absolute value) of the difference D in packet spacing at the receiver compared to the sender for a pair of packets. Detailed computation algorithms are found in RFC 1889. The value is zero if the connection was set in *send only* mode.
latency	The latency experienced during this connection The value is an estimate of the network latency, expressed in milliseconds. This is the average value of the difference between the Network Time Protocol (NTP) timestamp indicated by the senders of the RTP Control Protocol (RTCP) messages and the NTP timestamp of the receivers, measured when these messages are received. The average is obtained by summing all the estimates, then dividing by the number of RTCP messages that have been received. For details refer to Request For Comments (RFC) 1889.

Call Records Logged By Call Type

Each normal call between two parties logs one CDR End Call record. Each End Call record contains all fields identified above, but some fields may not be used. If a field is not used, it will be blank if it is an ASCII string field, or `0` (zero) if it is a numeric field. When supplementary services are involved in a call, more End Call records may be written.

In addition to the CDR End Call record, there may be up to one CMR record per endpoint involved in a call. In a normal call between two parties each using a Cisco IP phone, there will be two CMR records written: one for the originator and one for the destination of the call.

This section describes the records written for different call types in the system.

Normal Calls (Cisco IP Phone-to-Cisco IP Phone)

Normal calls log three records per call. They are EndCall plus two diagnostic records, one for each endpoint. In the EndCall record, all fields may contain valid information. The duration will always be non-zero unless the `CdrLogCallsWithZeroDurationFlag` flag is enabled (set to true). The `originalCalledPartyNumber` field will contain the same directory number as the `finalCalledPartyNumber` field.

Abandoned Calls

The logging of calls with zero duration is optional Normally, these records are not logged. If logging calls with zero duration is enabled, the following things should be noted:

- If the call was abandoned (such as when a phone is taken off hook and placed back on hook), various fields will not contain data. In this case, the `originalCalledPartyNumber`, `finalCalledPartyNumber`, the partitions associated with them, `destIpAddr`, and the `dateTimeConnect` fields will be blank. All calls that were not connected will have a *duration* of zero seconds. When a call is abandoned, the cause code is `0` (zero).
- If the user dialed a directory number and then abandoned the call before it was connected, the `First Dest` and `Final Dest` fields and their associated partitions will contain the directory number and partition to which the call would have been extended. The `Dest IP` field will be blank, and the duration will be zero.

Forwarded or Redirected Calls

The call records for forwarded calls will be the same as those for normal calls except for the `originalCalledPartyNumber` field and the `originalCalledPartyNumberPartition` fields. These fields will contain the directory number and partition for the destination that was originally dialed by the originator of the call. If the call was forwarded, the `finalCalledPartyNumber` and `finalCalledpartyNumberPartition` fields will be different and will contain the directory number and partition of the final destination of the call. Also, when a call is forwarded, the `lastRedirectDn` and `lastRedirectDnPartition` fields will contain the directory number and partition of the last phone that forwarded or redirected this call.

Calls With Busy or Bad Destinations

These calls will be logged as a normal call with all relevant fields containing data. The Called Party Cause field will contain a cause code indicating why the call was not connected, and the Called Party IP and Date/Time Connect fields will be blank. If the originator abandoned the call, the cause will be `NO_ERROR (0)`. The duration will always be zero seconds. These calls will not be logged unless `CdrLogCallsWithZeroDurationFlag` is enabled.

Call Management Records Logged By Call Type

Each normal call between two Cisco IP phones logs exactly two CMR records. Each call CMR record contains all fields identified above. When supplementary services are involved in a call, more than one record may be written. This section describes when diagnostic records are written for different call types in the system.

Normal Calls

Normal calls log exactly two CMR records per call, one for each phone involved in the call. Currently, only Cisco IP phones and MGCP gateways are capable of responding to the diagnostic information request. All fields will contain valid information.

Abandoned Calls

If the call was abandoned (such as when a phone is taken off-hook and placed back on hook), all fields related to streaming data will be blank (zero). This is because no streaming connection was established, and therefore no data was transferred. No records with blank fields will be logged if the `CdrLogCallsWithZeroDurationFlag` is disabled.

Forwarded Calls

The call records for forwarded calls will be the same as those for normal calls.

Calls With Busy or Bad Destinations

In the normal case, only records that represent calls that were actually connected will be logged. In order to log calls with bad destinations, you must enable `CdrLogCallsWithZeroDurationFlag`. If it is enabled, then all calls will be logged including the case where the user goes off-hook and then on-hook again.

If the calls are logged, they will be logged as normal calls with all relevant fields containing data. There will only be one record per call since the calls were never connected to a destination. The record will be for the originator of the call.

Codec Types (Compression / Payload Types)

Table 6-38 provides values and descriptions for codec types.

Table 6-38 Codec Description

Codec	Description
1	NonStandard
2	G711A-law 64k
3	G711A-law 56k
4	G711µ-law 64k
5	G711µ-law 56k
6	G722 64k
7	G722 56k
8	G722 48k
9	G7231
10	G728
11	G729
12	G729AnnexA
13	Is11172AudioCap
14	Is13818AudioCap
15	G729AnnexB
32	Data 64k
33	Data 56k
80	GSM
81	ActiveVoice
82	G726_32K
83	G726_24K
84	G726_16K

Cause Codes

Table 6-39 provides a list of cause codes that may appear in the Cause fields.

Table 6-39 Cause Code Descriptions

Cause Code	Description
0	No error
1	Unallocated (unassigned) number
2	No route to specified transit network (national use)
3	No route to destination
4	Send special information tone

Table 6-39 Cause Code Descriptions (continued)

Cause Code	Description
5	Misdialed trunk prefix (national use)
6	Channel unacceptable
7	Call awarded and being delivered in an established channel
8	Preemption
9	Preemption - circuit reserved for reuse
16	Normal call clearing
17	User busy
18	No user responding
19	No answer from user (user alerted)
20	Subscriber absent
21	Call rejected
22	Number changed
26	Non-selected user clearing
27	Destination out of order
28	Invalid number format (address incomplete)
29	Facility rejected
30	Response to STATUS ENQUIRY
31	Normal, unspecified
34	No circuit/channel available
38	Network out of order
39	Permanent frame mode connection out of service
40	Permanent frame mode connection operational
41	Temporary failure
42	Switching equipment congestion
43	Access information discarded
44	Requested circuit/channel not available
46	Precedence call blocked
47	Resource unavailable, unspecified
49	Quality of Service not available
50	Requested facility not subscribed
53	Service operation violated
54	Incoming calls barred
55	Incoming calls barred within Closed User Group (CUG)
57	Bearer capability not authorized
58	Bearer capability not presently available
62	Inconsistency in designated outgoing access information and subscriber class

Table 6-39 Cause Code Descriptions (continued)

Cause Code	Description
5	Misdialed trunk prefix (national use)
6	Channel unacceptable
7	Call awarded and being delivered in an established channel
8	Preemption
9	Preemption - circuit reserved for reuse
16	Normal call clearing
17	User busy
18	No user responding
19	No answer from user (user alerted)
20	Subscriber absent
21	Call rejected
22	Number changed
26	Non-selected user clearing
27	Destination out of order
28	Invalid number format (address incomplete)
29	Facility rejected
30	Response to STATUS ENQUIRY
31	Normal, unspecified
34	No circuit/channel available
38	Network out of order
39	Permanent frame mode connection out of service
40	Permanent frame mode connection operational
41	Temporary failure
42	Switching equipment congestion
43	Access information discarded
44	Requested circuit/channel not available
46	Precedence call blocked
47	Resource unavailable, unspecified
49	Quality of Service not available
50	Requested facility not subscribed
53	Service operation violated
54	Incoming calls barred
55	Incoming calls barred within Closed User Group (CUG)
57	Bearer capability not authorized
58	Bearer capability not presently available
62	Inconsistency in designated outgoing access information and subscriber class

Table 6-39 Cause Code Descriptions (continued)

Cause Code	Description
63	Service or option not available, unspecified
65	Bearer capability not implemented
66	Channel type not implemented
69	Requested facility not implemented
70	Only restricted digital information bearer capability is available (national use)
79	Service or option not implemented, unspecified
81	Invalid call reference value
82	Identified channel does not exist
83	A suspended call exists, but this call identity does not
84	Call identity in use
85	No call suspended
86	Call having the requested call identity has been cleared
87	User not member of Closed User Group (CUG)
88	Incompatible destination
90	Destination number missing and DC not subscribed
91	Invalid transit network selection (national use)
95	Invalid message, unspecified
96	Mandatory information element is missing
97	Message type non-existent or not implemented
98	Message is not compatible with the call state, or the message type is non-existent or not implemented
99	An information element or parameter does not exist or is not implemented
100	Invalid information element contents
101	The message is not compatible with the call state
102	The call was terminated when a timer expired and a recovery routine was executed to recover from the error
103	Parameter non-existent or not implemented - passed on (national use)
110	Message with unrecognized parameter discarded
111	Protocol error, unspecified
126	Call split. This is a Cisco-specific code. It is used when a call is terminated during a transfer operation because it was split off and terminated (was not part of the final transferred call). This can help determine which calls were terminated as part of a transfer operation.
127	Interworking, unspecified

Alarms

An alarm is issued when CDR or diagnostic data is enabled, and the system is unable to write the data into the database.

Unable to Write CDR data. (Alarm # 1711 - Major Alarm)

The system attempted to open the database, and was unsuccessful. Probable causes include:

- Cisco CallManager does not have sufficient privileges to open the file for writing to the database. Make sure Cisco CallManager has privileges that will permit write operations.
- The path is not set up, or the database server is down.